Investigating methods and materials which can more accurately model and reduce a gravitational wave detector's coating thermal noise

Mark Fletcher

School of Physics and Astronomy

University of Glasgow

Presented as a thesis for the degree of Ph.D. in the University of Glasgow, University Avenue, Glasgow G12 8QQ

November 2018

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to first thank my supervisor Dr Iain Martin for his help, guidance and support throughout my PhD. This wouldn't have been possible without you. Next, I would like to thank Dr Jessica Steinlechner as she was instrumental in helping me to complete this body of work. I learned a lot from the both of you. More generally, I am grateful to the Institute for Gravitational Research (IGR) and the Science and Technology Facilities Council for providing me with a fully funded PhD position.

Throughout my time in the IGR, I have crossed paths with many people who not only contributed to my work but also made my experience enjoyable. I will never forget you guys: Dr Angus Bell, Dr Liam Cunningham, Dr Alan Cumming, Dr Marielle Van Veggel, Dr Karen Hamilton, Dr Peter Murray, Dr Kieran Craig, Dr Richard Middlemiss, Dr Martin Hart, Dr Raymond Robie, Dr Borja Sorazu, Dr Jamie Scott, Mr Russell Jones, Ms Sheena Barclay, Mr Mervyn Chan, Mr Alastair Doye, Mr Joseph Briggs, Mr Hunter Gabbard, Ms Hafizah Isa, Mr Brynley Pearlstone, Mr Simon Tait, Ms Jennifer Wright, Mr Zeno Tornasi and Mr Colin Craig.

I would also like to thank those outside of work that have helped me along the way. Firstly, my oldest friend Tom, you were always there when I needed a spot of advice or a welcomed break from work. Secondly, my family. You have always been there for me and helped tremendously in everything I do. And lastly my girlfriend, Indre. When I didn't want to continue, you pushed me to go further, work a little bit harder and stay positive. I will always be indebted to you. Aciu labai!

Abbreviations

IGR – Institute of Gravitational Research

ECR – Electron Cyclotron Resonance

IBD – Ion Beam deposition

RLVIP – Reactive Low Voltage Ion Plating

GeNS – Gentle Nodal Support

FEM – Finite Element Modelling

FEA – Finite Element Analysis

PCI – Photothermal Commonpath Interferometry

UWS – University of West of Scotland

DAQ – Data Acquisition

CoTE- Coefficient of Thermal Expansion

LR – Liftshitz and Roukes

LIGO – Laser Interferometry Gravitational wave Observatory

ET – Einstein Telescope

KAGRA – Kamioka Gravitational Wave Detector

LISA – Laser Interferometer Space Antenna

VBA – Visual Basic Analysis

SQL – Standard Quantum Limit

Preface

This thesis is an account of work carried out in the IGR between October 2014 and September 2018. The work focuses upon improving the processes currently used to determine the coating thermal noise in a gravitational wave detector as well as identifying materials which can reduce coating thermal noise.

In Chapter 1, gravitational waves are introduced: what they are, how they are formed and how they can be detected. A summary of all of the detected gravitational wave signals is also presented.

In Chapter 2, mechanisms of coating thermal noise are discussed. The relationship between Brownian noise and mechanical loss is described in depth.

In Chapter 3, the theory of bulk and shear loss is presented and is used to design two mathematical approaches to determine a coating's bulk and shear loss. Using these methods, the bulk and shear losses of three different coatings were calculated (Electron Cyclotron Resonance Ion beam deposited amorphous silicon, Reactive Low Voltage Ion Plating amorphous silicon and Reactive Low Voltage Ion Plating tantala). Both Dr Iain Martin and the author developed the methods used to determine a coating's bulk and shear losses. However, it was the author who implemented the theory into Excel VBA code which was then ran by the author to determine all of the bulk and shear losses presented. The DEKTAK thickness measurement of the ion-beam deposited (IBD) amorphous silicon coating was carried out by Mr David Vine at the University of West of Scotland. The mechanical losses of the IBD amorphous silicon coating were measured by Dr Iain Martin, Dr Jessica Steinlechner and the author. Dr Jessica Steinlechner measured the coated losses of the reactive low voltage ion plating (RLVIP) amorphous silicon coating and also estimated the maximum and minimum values of uncoated loss. Mr Simon Tait along with the author measured the coated losses of the tantala coated silica disc. Dr Raymond Robie measured the uncoated loss of the disc. All finite element analysis (FEA) presented in this chapter was carried out by the author, however guidance was sometimes sought from Dr Liam Cunningham, Dr Matthew Abernathy and Dr Iain Martin. All of the analysis of the coatings (that is not explicitly referred to above) was carried out by the author under the supervision of Dr Iain Martin and Dr Jessica Steinlechner.

In Chapter 4, the effect of stress and curvature upon the thermoelastic loss of a cantilever and disc was investigated. Dr Jessica Steinlechner measured the mechanical losses of the 59×10^{-6} m thick, amorphous silicon coated, silicon cantilever. The author in conjunction with Liam Cunningham and ANSYS support wrote a FEA code to calculate and plot the thermoelastic losses of a cantilever. COMSOL support advised in how best to stress and curve a cantilever within a thermoelastic loss analysis. The experimental in Section 4.6.3 were all carried out by the author with the exception of the coated silica disc which was measured by Mr Simon Tait. The experimental and FEA calculated uncoated losses of the silicon disc, presented in Section 4.7.1, were determined by Dr Matteo Lorrenzini and the University of Jena respectively. All other measurements and FEA analysis in this chapter, that was not explicitly described above, was carried out by the author under the guidance and supervision of Dr Iain Martin.

In Chapter 5, the effect of stress and temperature on the optical properties of silicon nitride membranes at 1550 nm was investigated. Dr Jessica Steinlechner and Dr Iain Martin measured the change in the product of the membrane's absorption and thermo-refractive coefficient as a function of 1550 nm and 532 nm laser power. Dr Iain Martin, Dt Jessica Steinlechner and the author measured the change in temperature of the membrane as a function of laser power. Mr Simon Tait measured the change in frequency of the membrane's resonant modes as it's frame was heated with a 532 nm laser. The data analysis was carried out by Dr Iain Martin, Dr Jessica Steinlechner, Dr Angus Bell and the author. All of the FEA modelling was conducted by the author, however guidance was given by Dr Liam Cunningham, Dr Iain Martin and COMSOL support.

In Chapter 6, the performance of using silicon nitride membranes as a substrate in mechanical loss experiments is compared to other commonly used substrates. The loss measurements of the uncoated membranes and silicon cantilevers were made solely by the author. However, the silica cantilever losses were measured by Dr Jessica Steinlechner whilst the losses of the silica discs were measured by both the author and Mr Simon Tait. The loss of the membrane coated with silica was measured by the author. The loss of the silica cantilever coated with silica was measured by both Mr Simon Tait and the author. All of the data analysis and FEA presented in this chapter was carried out by the author under the supervision of Dr Iain Martin. The electric field intensity profiles of the coatings were calculated by Dr Jessica Steinlechner and Dr Iain Martin.

Summary

Gravitational waves were predicted by Einstein's General Theory of Relativity which described gravity as arising from the curvature of space-time due to the presence of mass. Gravitational waves are caused by the acceleration of an asymmetrical mass distribution which perturbs space-time causing a ripple-like effect that travels at the speed of light. These ripples are known as gravitational waves.

The first directly detected gravitational waves were observed on the 14th of September 2015. These waves were formed from the inspiral and merger of two black holes. Since that day, gravitational waves have been observed (from four more binary black hole systems and from a binary neutron star system). These detections have demonstrated the power and potential of gravitational wave astronomy as they have provided vast new information regarding the binary systems which produce them (e.g. the mass of black holes, that merging neutron stars are a source of short gamma-ray bursts etc).

These detections were made using ground-based interferometric detectors in which a laser beam is split and passed along two perpendicular arms. At the end of these arms, the laser light is reflected back towards the beam splitter by test masses coated with a highly reflective mirror coating. At the beam splitter, the two laser beams recombine and the intensity of the signal is monitored. Differential changes in the arm-length result in changes to the interference pattern. Since gravitational waves are only expected to change a 1 km arm length detector by approximately 1×10^{-19} m, it is essential for all sources of noise to be exceedingly low.

Coating thermal noise - arising from thermally induced motion in the interferometer mirror coatings - forms a major limit to sensitivity of current gravitational wave detectors at their most sensitive frequencies. The magnitude of this noise source is proportional to the detector's operation temperature, laser beam radius and coating mechanical loss (also known as internal friction).

The research presented in this thesis focuses upon improving the processes currently used to determine the coating thermal noise in a detector as well as identifying materials which can reduce coating thermal noise.

Chapter 1 contains an introduction to the theory of gravitational waves, describes the sources used to produce them, the experimental methods used to detect them and summarises all of the signals already observed.

Chapter 2 presents a detailed description of thermal noise and outlines how the thermal noise of a highly reflective coating in a gravitational wave detector has been historically calculated.

Chapter 3 introduces the theory of a new and more accurate approach to determining the thermal noise of a gravitational wave detector coating by using the coating's bulk and shear mechanical losses. Two methods to determine the bulk and shear losses of a coating are developed and applied to a range of coatings (ECR IBD amorphous silicon and RLVIP amorphous silicon) that were deposited upon silicon cantilevers. It was found that, the bulk and shear losses of a coating could be significantly different to each other. The ECR IBD amorphous silicon coating had a bulk loss of $(3.4 \pm 0.8) \times 10^{-4}$ and shear loss of $(1.4 \pm 0.2) \times 10^{-4}$. Whilst the room temperature deposited RLVIP amorphous silicon coating had a bulk loss of $0 \pm 1 \times 10^{-9}$ and a shear loss of $(1.53 \pm 0.09) \times 10^{-4}$.This thesis is believed to present the first analysis which shows that the bulk and shear losses of a coating can indeed be significantly different from each other. Furthermore, the results also showed that it is possible for the loss of a coating to be purely shear loss. Since the thermal noise of a gravitational wave detector is less sensitive to shear loss than bulk loss, this result indicates the possibility of being able to further reduce coating thermal noise by using materials which are dominated by shear loss. Shear loss is explained in this chapter to have less effect than bulk loss upon the thermal noise of a gravitational wave detector, as a result of the shear motion of the coating changing the arm length of the detector much less than the coating's bulk motion. Chapter 3 lastly investigated whether the same coating deposited upon two different substrate geometries can be determined to have the same bulk and shear losses. This was an interesting test as it would help to illustrate the robustness of the methods developed to determine bulk and shear loss. A tantala coating was deposited upon a cantilever and disc substrate. The coating's bulk and shear losses were determined to be $(1.90 \pm 0.30) \times 10^{-3}$ and $(6.7 \pm 0.5) \times 10^{-4}$ upon the cantilever and $(1.28 \pm 0.22) \times 10^{-3}$ and $(7.8 \pm 0.5) \times 10^{-4}$ upon the disc. Whilst good qualitative agreement can be observed between the results (the bulk losses for the two geometries are greater than their

shear loss), the different substrate's bulk losses are significantly different from one another. It is expected that this could be due to limited experimental data.

To determine the thermal noise of a coating in a gravitational wave detector, the mechanical loss of the coating must first be known. Typically, the coating's mechanical loss is obtained using a combination of ring down experiments and mathematical equations. However, recent research and observations have called into question the accuracy of coating losses which have been determined using this approach when silicon cantilever substrates are used to deposit the coating on. It is speculated that the inaccuracy of these losses could be due to the effect of the coating causing the coated cantilever to curve and stress which in turn changes the silicon cantilever's thermoelastic loss. Chapter 4 presents an investigation into the effect of stress and curvature upon a silicon cantilever's thermoelastic loss. The results from finite element analysis (FEA) indicated that the thermoelastic loss of a silicon cantilever (which has been coated) is unaffected by stress but is affected by curvature. Experimental evidence appears to confirm these conclusions and further imply that either the FEA does not fully account for the effect of curvature upon thermoelastic loss or an unmodelled effect is also affecting a silicon cantilever's thermoelastic loss. FEA was also used to investigate the effect of stress and curvature upon a silicon disc's thermoelastic loss. A disc geometry was also investigated as discs along with cantilevers are the two most commonly used substrates in mechanical loss experiments due to their low substrate loss. It was observed that the magnitude of stress and curvature caused by depositing a coating onto a silicon disc in the lab had no effect upon the substrate's thermoelastic loss.

Some future gravitational wave detectors, such as the Einstein Telescope (ET), are being designed to operate at cryogenic temperatures. To achieve these detector's thermal noise requirements, new mirror coatings with low mechanical loss and low absorption (to ensure the mirror coating remains cool) at cryogenic temperatures will have to be found. One possible material which could be potentially used in a highly reflective coating is silicon nitride. Silicon nitride has already been shown to have a low mechanical loss which decreases as its stress increases. However, its absorption is currently too high to meet the ET's design requirements if it is to be used as one of two materials in a simple coating bi-layer structure. Chapter 5 outlines an investigation into whether the absorption of a silicon nitride membrane can be reduced by changing its stress. Due to experimental limitations, the direct effect of stress upon absorption could not be quantified, instead the effect of stress upon the product of the membrane's absorption and its thermo-refractive coefficient was determined. The product of the membrane's absorption and its thermo-refractive index was found to be independent of stress. Whilst this result implies that absorption is independent of stress, this should be explicitly checked. It was also determined from this experiment that the thermal conductivity and thermal expansion of a Norcada fabricated low stress silicon nitride membrane is (23 ± 3) W/mK and $(1.4 \pm 0.2) \times 10^{-6}\ 1/\mathrm{K}$ respectively.

The Institute for Gravitational Research (IGR) has recently gained access to an operational coating chamber. The usual substrate geometries which coatings are deposited upon for mechanical loss experiments are not compatible with the coating chamber due to their size. In this chapter, the possibility of using a thin, low loss silicon nitride membrane as a substrate for this chamber is presented. If silicon nitride membranes can be successfully used as substrates it will enable full control over the deposition process, meaning that the effect of individual deposition parameters upon a coating's mechanical loss can be investigated. Chapter 6 presents an investigation into whether silicon nitride membranes can be used as substrates in mechanical loss experiments and compares their performance to other commonly used substrates (silica cantilevers, silicon cantilevers and silica discs). Out of these substrates, membranes were shown to be the second most sensitive substrate to clamping effects, exhibit the second least variation of loss for nominally identical substrates and have the second lowest mechanical loss. Whilst none of these results excludes silicon nitride membranes being used in mechanical loss experiments, they indicate a preference should be shown in selecting silica discs before membranes. However, it was also observed that the bulk and shear losses of the same coating deposited upon a membrane and cantilever had significantly different values. Further work is required to fully understand these differences in coating loss before membranes can be reliably used for these experiments. The effect of heat-treatment upon a silicon nitride membrane's absorption was investigated at both 1064 and 1550 nm wavelengths. No degradation in performance was observed for annealing temperatures up to 900 °C, indicating that silicon nitride could be used as a partner material for amorphous silicon (which requires heat-treatment at 450 °C to reduce it's absorption to a usable level). The effect of using a silicon nitride and amorphous silicon coating in the (ET) was investigated. It was shown that such a coating can significantly reduce the low-temperature coating thermal noise when compared to an Advanced LIGO (silica/ titania doped tantala) coating. However, this coating does not have a suitably low optical absorption and so "multi-material" coating designs were investigated and a coating design proposed which meets the ET's absorption requirement. Whilst this "multi-material" coating would still not enable the ET to meet thermal noise design requirements, it represents the best coating which can currently be made from silica, titania doped tantala, silicon nitride and amorphous silicon materials.

1 Gravitational wave detection

1.1 Introduction

Gravitational waves – first detected in 2015 - were predicted by Einstein's General Theory of Relativity which described gravity as arising from the curvature of space-time due to the presence of mass [1]. Gravitational waves are caused by the acceleration of an asymmetrical mass distribution which perturbs space-time causing a ripple like effect, travelling at the speed of light. These ripples are known as gravitational waves.

Gravitational waves exhibit an extremely weak interaction with matter. Astrophysical bodies with large masses and large accelerations are therefore required to create waves which are detectable on Earth. Examples of these astrophysical sources include supernovae explosions, inspiraling binary neutron stars and merging black holes. Whilst the weak interaction of gravitational waves with matter makes them difficult to detect, it also makes them an incredibly useful tool in astronomy as they are less susceptible to scattering and absorption as they travel through space. Gravitational wave astronomy has already begun to enable the testing of cosmological theories [2], increase current knowledge of black holes [3] and neutron stars [4] and may lead to new unexpected phenomena being detected.

Gravitational waves were indirectly proven to exist by Hulse and Taylor in 1974 when they discovered a binary pulsar system and showed that its orbital decay could be explained by the loss of energy as gravitational wave radiation [5]. The first direct detection of a gravitational wave occurred on the 14th of September 2015 and was made by the Advanced Laser Interferometric Gravitational Wave Observatory (Advanced LIGO) [6]. This signal (called GW150914) was produced by an inspiral and merger of two black holes. Since this measurement has been made, a further 5 signals have been measured, some of which were created by different sources (binary neutron star system [7]), observed by a different detector (Advanced Virgo) [8] as well as confirmed through non-gravitational means (electromagnetically) [7]. To enable the detection of weaker signals and improve the signal to noise ratio, the sensitivity of current gravitational wave detectors must be improved.

The first gravitational wave detector was built by Joseph Weber in the 1960's. Weber's detector consisted of an aluminium bar designed to resonant when interacting with a passing gravitational wave with a specific frequency [9]. Since the days of this initial detector, significant advancements in detector technology has been made, most noticeable the shift from using resonant bar detectors to interferometers [10] [11]. Whilst current detectors have proved their ability in measuring gravitational wave signals, their sensitivity must still be improved to fully reap the rewards of gravitational wave astronomy. This can be achieved by reducing the current detectors' dominant noise sources which include: shot noise and radiation pressure associated with the laser light, seismic noise caused by the Earth's crust shaking, thermally induced motion of the mirrors and gravity gradient noise caused by the direct gravitational interactions between the interferometer mirrors and the surrounding environment. Current research is directed at tackling these noise sources and it is hoped that future detectors will show a 10 times sensitivity improvement as a result [12].

1.2 Gravitational radiation

Gravitational radiation is caused by the acceleration of mass. As with the production of electromagnetic radiation, a time varying moment must exist. The monopole moment of electromagnetic radiation and gravitational radiation is simply the total charge and total mass of a system respectively. Due to the conservation of charge and mass laws in a closed system, neither electromagnetic or gravitational monopole moments are time varying. The electromagnetic or gravitational dipole moments describe the centre of charge or mass in a system respectively. No conservation law precludes the existence of a time varying electric dipole moment and as a result this moment can produce electromagnetic radiation. However, a gravitational dipole moment cannot be time varying as a result of the conservation of momentum. All higher order odd-number polarities cannot produce gravitational and electromagnetic radiation as a result of the conservation of charge and mass laws. However, no conservation law prevents the existence of higher even moments producing radiation. A quadrupolar moment is the first moment capable of producing gravitational wave radiation. To produce a quadrupole moment capable of producing gravitational radiation, an acceleration of an asymmetric mass distribution is required. Physical examples of this distribution are the revolution of a non-spherical object or a binary star system where the individual stars have different masses.

Gravitational waves are transverse waves. They exert differential strains perpendicular to their direction of motion. The direction of the strains which the wave exerts is specified by the polarisation of its quadrupole moment. Figure 1-1 illustrates the two different quadrupole polarisations of a gravitational wave travelling perpendicularly into the page and through a ring of masses with diameter L.

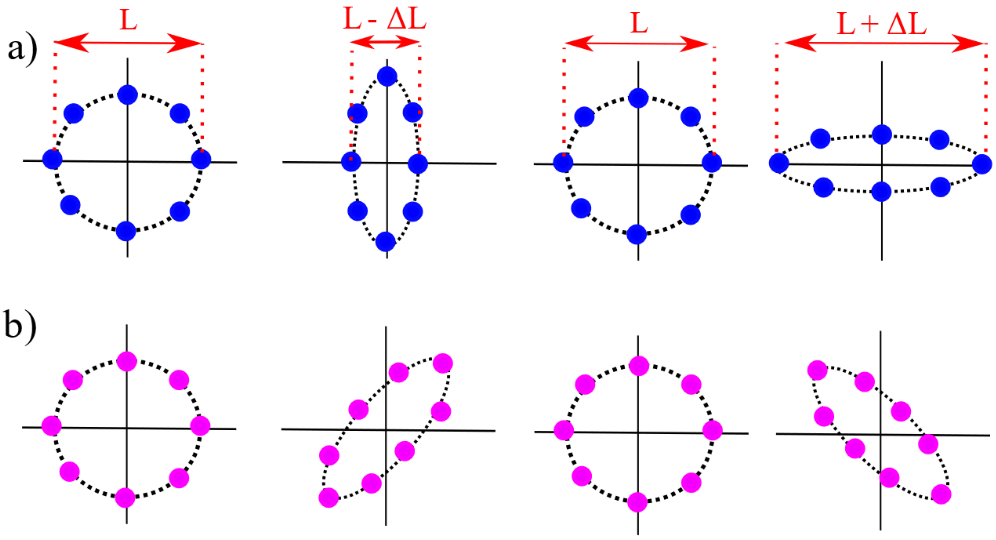

Figure 1-1: Effect of a passing gravitational wave upon a ring of masses. The gravitational wave is travelling perpendicularly into the page. a) illustrates the effect caused by an h_+ polarised wave whilst b) $h_\times$.

It can be observed that as the wave passes through the page, the perpendicular axes of the ring simultaneously experience an extension and compression of magnitude ΔL. The strain amplitude (h) of the wave is given by:

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

Current gravitational wave detectors aim to measure this strain caused by a passing wave. The most sensitive, currently operating gravitational wave detectors are the Advanced LIGO detectors. When it is performing at its technical specifications it is hoped that these detectors will exhibit a peak strain sensitivity of 3.5×10^{-24} Hz$^{-1/2}$ [13].

1.3 Gravitational wave signals

Gravitational wave sources are commonly divided into four categories: compact binary coalescence, burst, continuous and stochastic.

1.3.1 Compact binary coalescence

Gravitational waves are emitted by compact binary systems made from any combination of neutron stars and black holes. In a binary system the two bodies orbit one another around a common centre of mass. Over time the orbital period decreases as a result of the system losing energy in the form of gravitational waves. This loss in energy and decrease in distance between the two orbiting bodies results in an increase of the frequency and amplitude of the emitted gravitational waves. In the final few seconds before the bodies coalesce, the frequency of the signal rises from approximately 10 Hz to 300 Hz [14], resulting in the characteristic "chirp" signal as measured by a detector. The strain amplitude (h) of the emitted gravitational waves at a distance d, can be approximated as [2]:

$$h \approx 10^{-23} \left(\frac{100\ Mpc}{d}\right) \left(\frac{M_b}{1.2 M_s}\right)^{\frac{5}{3}} \left(\frac{f}{200\ Hz}\right)^{\frac{2}{3}}, \tag{1.2}$$

where M_s is the mass of the sun, f is the frequency of the gravitational wave and M_b is the mass parameter of the binary system which can be calculated using:

$$M_b = \frac{(M_1 M_2)^{\frac{3}{5}}}{(M_1 + M_2)^{\frac{1}{5}}}, \tag{1.3}$$

where M_1 and M_2 are the individual masses of the two bodies in the binary system. Whilst the population of binary black hole systems is expected to be smaller than that of neutron stars, Equations (1.2) and (1.3) tell us that binary black holes should be easier to observe due to their larger masses. A binary black hole system formed of two black holes of 36 and 29 solar masses, located 410 Mpc away from a detector would result in a strain amplitude of 1×10^{-21} at coalescence.

Detecting gravitational wave signals from inspiraling binary systems has enriched astrophysical knowledge. By combining gravitational wave detection with electromagnetic observations, astronomers have been able to observe that coalescing neutron stars result in the production of other forms of radiation such as gamma ray bursts [7]. Furthermore, binary signals have also been used to make an independent measurement of the Hubble constant [15].

1.3.2 Burst signal

A burst signal is one which typically lasts less than one second. A supernova explosion is an example of a burst source. There are two types of supernovae which can result in a gravitational wave, Type 1a and Type II [16].

A Type 1a supernova occurs when the initial star collapsing has a mass of less than 5 solar masses [17]. If the white dwarf formed from the collapse is able to accrete mass from a companion star and exceeds 1.4 solar masses (and therefore surpasses the Chandrasekhar limit [18]), a Type 1a supernova occurs. A Type 1a supernova explosion which is asymmetric will result in gravitational waves. Stars larger than 5 solar masses can also become supernovae capable of creating gravitational waves. These are known as Type II supernovae. They occur when a star larger than 5 solar masses fuses a dense iron core that exceeds the Chandrasekhar limit. An asymmetric explosion will again produce gravitational waves. Recent investigations show that supernova explosions are frequently asymmetric as a result of the explosion interacting with orbiting star matter [19]. The strain amplitude from a supernova explosion is predicted to approximately be [16]:

$$h \approx 5 \times 10^{-22} \left(\frac{E}{10^{-3} M_s c^2}\right) \left(\frac{15\ Mpc}{d}\right) \left(\frac{1\ kHz}{f}\right) \left(\frac{1\ ms}{t}\right)^{\frac{1}{2}} , \tag{1.4}$$

where E is the total energy radiated, c is the speed of light and t is the time until core collapse. Whilst the amplitude of the wave can be loosely calculated, detecting burst signals from supernovae is going to be extremely difficult as there are no robust theoretical models which describe their characteristics [20]. Due to the lack of a theoretical model, matched filtering (the data analysis method used to detect burst and continuous signals) will not be able to be used to detect this type of event as it operates by correlating theoretical signals (calculated using a theoretical model) to the output of the detector. Instead of using this approach, broadband searches of the data will be conducted and any anomalies observed in multiple detector data streams will be investigated.

1.3.3 Continuous signal

A continuous gravitational wave signal is one which transmits for a long period of time at a near constant frequency. Although these signals are much weaker than both binary and burst signals, being able to integrate the signal over long time intervals increases their likelihood of observation [21]. The most likely sources of these signals are neutron stars which exhibit an axial asymmetric: shape, precession or internal oscillation [22].

Axial asymmetry can arise either at the birth of a star or over time, due to changing rotational speed. The strain amplitude of gravitational waves created from these sources can be approximated by [23]:

$$h \approx \frac{4\pi^2 G I_{zz} f^2 \epsilon}{c^4 d} \ , \tag{1.5}$$

where G is the gravitational constant, ϵ is the ellipticity of the star and I_{zz} is the star's principal moment of inertia. The ellipticity of neutron stars is highly uncertain but it is thought to be limited by the breaking strain of their crust. It is suspected that exotic stars (such as the solid strange-quark star) which have a solid core will be able to support considerably larger ellipticities than a neutron star (which would enable them to produce larger strain amplitudes).

An axial asymmetric precession of neutron stars occurs when a star's symmetry axis and rotational axis do not match resulting in the star wobbling as it rotates. The strain amplitude of a gravitational wave caused by this effect can be approximated as [23]:

$$h \approx 10^{-27} \left(\frac{\theta_\omega}{0.1}\right)\left(\frac{1\,kpc}{d}\right)\left(\frac{v}{500\,Hz}\right)^2 \ , \tag{1.6}$$

where v is the orbital frequency and θ_ω is the angle of the wobble in radians. It is suspected that neutron stars may develop their wobble by accreting matter from other neutron stars or nebula. This is a field of current investigation [24]. The strain amplitude of a freely precessing neutron star located at 400 Mpc away, spinning at 500 Hz and with an optimistically large angle of wobble 0.1 [25], is of the order 10^{-27}. This signal is approximately one million times weaker than that of a black hole binary coalescence (calculated in Section 1.3.1).

Axial asymmetric oscillations of a neutron star's interior fluids are also expected to cause gravitational waves. These oscillations can occur naturally in newly formed neutron stars or in stars which accrete matter from others. The strain amplitude of a gravitational wave caused by a rotating neutron star's fluid experiencing axial asymmetric oscillations can be estimated using [26]:

$$h \approx 4.4 \times 10^{-24} \alpha \left(\frac{\omega}{\sqrt{\pi G}}\right)^3 \left(\frac{20\ Mpc}{d}\right)\ , \tag{1.7}$$

where $\boldsymbol{\alpha}$ is the amplitude of the oscillation and $\boldsymbol{\omega}$ is the angular velocity.

1.3.4 Stochastic signal

The stochastic gravitational wave background is created from the superposition of many weak, unresolved astrophysical signals caused by, but not limited to: the expansion of the universe after the Big Bang [27], binary systems made from both neutron stars and black holes [28], supernovae explosions [29] and individual neutron stars [30]. To detect the stochastic gravitational wave background, multiple detectors are required to cross correlate their data streams which removes independent detector noise and enables the identification of mutually observed random signals [31].

Some of the signals which are proposed to be contained within the stochastic gravitational wave background are expected to date back as far as 10^{-35} seconds after the Big Bang [14]. To put this in perspective, the Cosmic Microwave Background dates from 10^5 years after the Big Bang [32]. By being able to detect these ancient signals within the stochastic gravitational wave background, our knowledge and understanding of the earliest stages of the universe will be improved.

1.4 Gravitational wave detections

To date, six gravitational wave signals have been observed. The masses and distances associated with these detections are listed in Table 1-1.

The first gravitational wave was detected on the 14th of September 2015. This signal was called GW150914 and was measured by the Advanced LIGO detectors [6]. It was produced by the inspiral and merger of two black holes. The detector at Livingston observed the signal 7×10^{-3} s before Hanford which enabled the location of the event to be determined in the Southern Celestial Hemisphere with a 90 % probability. The chirped signal (Section 1.3.1) indicative of an inspiral and merger can be observed in Figure 1-2 which shows a time-frequency representation of the measured strain in each detector. Immediately before the collision, the black holes were travelling at 1.8×10^5 km/s. The collision occurred at a luminosity distance (an estimate of a distance to the event based upon its luminosity) of approximately 410 Mpc. This signal also represented the first direct observation of a binary black hole system merging to form a single black hole and provided the first evidence of the existence of stellar black holes with a mass equal to a few tens of solar mass.

The second observed gravitational wave (GW151226) occurred on the 26th of December 2015 [33]. It was again caused by the inspiral and merger of two black holes and was only measured by Advanced LIGO. Although the black hole masses forming this binary system were smaller than those which caused GW150914, the signal lasted five times as long meaning scientists could get a test of general relativity which was twice as precise as that made during the first detection [34].

The third (GW170104) [35] and fourth (GW170608) [36] gravitational wave signals were also caused by the inspiral and merger of binary black holes and they too were only measured by Advanced LIGO. GW170104 is the furthest confirmed event to date and GW170608 observed the smallest black holes to date.

On the 1st of August 2017, Advanced Virgo joined Advanced LIGO in measuring data. On the 14th of August, both of the Advanced LIGO detectors as well as Advanced Virgo observed the inspiral and merger of a binary black hole system (GW170814) [8]. Observing the event using three detectors improved sky localisation by a factor of 20 [37]. It also enabled the first measurement of the polarisation of a gravitational wave.

The most recent gravitational wave detection occurred on the 17th of August 2017 (GW170817) [7]. This event was only observed by the Advanced LIGO detectors, however Advanced Virgo was able to help in determining the position of the event (as a result of knowing the direction in which the Advanced Virgo detector is least sensitive to passing gravitational waves). The signal was produced by the inspiral and merger of two neutron stars and lasted for over 100 seconds. 1.7 seconds after the gravitational wave detectors observed this event, both the Fermi and INTEGRAL space telescopes observed a short gamma ray burst in the direction of the event. After these detections, an alert to other astronomers was issued. An event with characteristics consistent with merging neutron stars was subsequently observed over the following days and weeks in both the radio and x-ray bands. The observations of this event via gravitational waves and multiple electromagnetic wavelength helped to provide evidence of merging neutron stars being responsible for short gamma ray bursts.

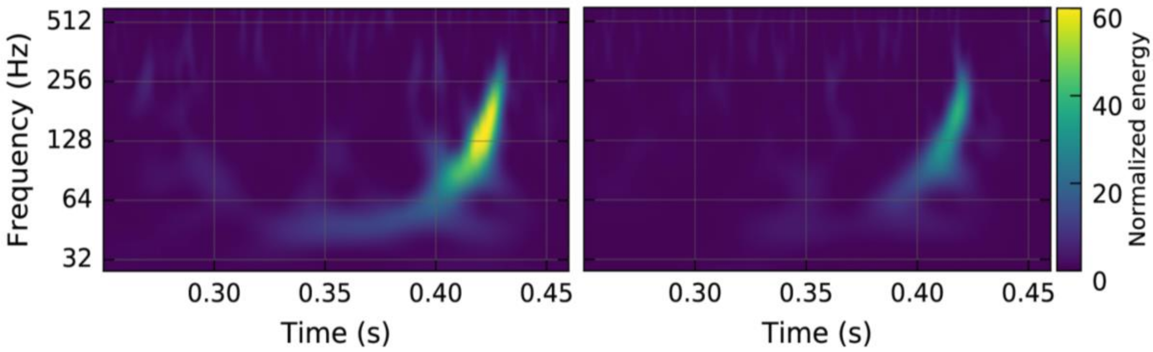

Figure 1-2: A time-frequency representation of the strain data measured by the two Advanced LIGO detectors: (a) Hanford and (b) Livingston for the first gravitational wave observed (GW150914). The strength of the gravitational wave signal (yellow) can be clearly seen through the noise (blue). The gravitational wave signal can be seen to increase in frequency over time, thus illustrating a "chirp" signal [6]. *The increase in frequency of the signal is due to the two black holes inspiralling into one another. The signal ends when the black holes have fully merged.*

Property, (units)	GW150914	GW151226	GW170104	GW170608	GW170814	GW170817
Source of Wave	Binary black hole	Binary black hole	Binary black hole	Binary black hole	Binary black hole	Binary neutron star
Primary mass, (solar masses)	$35.4^{+5.0}_{-3.4}$	$14.2^{+8.3}_{-3.7}$	$31.2^{+8.4}_{-6.0}$	$12.0^{+5.0}_{-3.4}$	$30.5^{+5.7}_{-3}$	1.36 – 1.60
Secondary mass, (solar masses)	$29.8^{+3.3}_{-4.3}$	$7.5^{+2.3}_{-2.3}$	$19.4^{+5.3}_{-5.9}$	7.0^{+2}_{-2}	$25.3^{+2.8}_{-4.2}$	1.17 – 1.36
Chirp mass, (solar masses)	$28.2^{+1.8}_{-1.7}$	$8.9^{+0.3}_{-0.3}$	$21.1^{+2.4}_{-2.7}$	$7.9^{+0.2}_{-0.2}$	$24.1^{+1.4}_{-1.1}$	$1.188^{+0.004}_{-0.002}$
Final mass, (solar masses)	$62.2^{+3.7}_{-3.4}$	$20.8^{+6.1}_{-1.7}$	$48.7^{+5.7}_{-4.6}$	$18.0^{+4.8}_{-0.9}$	$53.2^{+3.2}_{-2.5}$	$2.74^{+0.04}_{-0.01}$
Radiated GW energy, (c^2 × solar masses)	$3^{+0.5}_{-0.5}$	$1.0^{+0.1}_{-0.2}$	$2.0^{+0.6}_{-0.7}$	$0.85^{+0.07}_{-0.17}$	$2.7^{+0.4}_{-0.3}$	> 0.025
Luminosity distance (Mpc)	440^{+160}_{-180}	440^{+180}_{-190}	880^{+450}_{-390}	340^{+140}_{-140}	540^{+130}_{-210}	40^{+8}_{-14}

Table 1-1: List of the sources, masses, energies and distances associated with the different gravitational wave detections [6] [7] [8] [33] [36] [35].

1.5 Gravitational wave detectors

There are two different types of detector which have been used to detect gravitational waves. These are known as resonant bar detectors and laser interferometric detectors.

1.5.1 Resonant bar detectors

Joseph Weber built the first resonant bar detector in 1960 [9]. This detector was made from a large aluminium cylinder which had a resonant frequency of approximately 1660 Hz. The cylinder was isolated from ground and acoustic vibrations, operated at room temperature and its motion was monitored by using piezoelectric transducers and amplifiers. Weber designed his detector with the hope that a gravitational wave would pass through the detector exciting the resonance of the bar. The resonance of the bar was chosen to match the expected frequency of a wave caused by a supernovae explosion. To ensure the significance of his observations, Weber aimed to measure coincident signals from two detectors positioned at different positions on the Earth (initially they were separated by 2 km but this was later increased to 1000 km). In 1969 and 1970, Weber reported coincident events in both of his detectors and claimed that they were caused by gravitational waves [38, 39]. Numerous other research groups around the world tried and failed to duplicate his observations [40, 41, 42]. Scepticism grew regarding his claims, largely leading to them being discounted when calculations suggested that Weber's detectors strain sensitivity was approximately 7 orders of magnitude less sensitive than required to measure a gravitational wave.

Over the years, improvements were made to Weber's initial detector design to maximise the probability of detection. To reduce the effect of thermal noise, some detectors were operated at cryogenic temperatures [43]. To increase directional sensitivity, some detectors used a spherical resonant mass which could be excited from any direction [44]. Today, all resonant bars have been decommissioned. However, two cryogenic spherical detectors are still operational [44, 45]. The decrease in number of resonant detectors is due to them having narrower operational bandwidths (only sensitive at detector resonance) and less sensitivity than interferometric detectors.

1.5.2 Interferometric detectors

To improve upon the sensitivity and operational bandwidth of Weber's resonant bar detector, Gertsenshtein and Putovoit [46] proposed a new type of detector based on an interferometer.

In a Michelson interferometer, a laser beam is split using a beam splitter and passed along two perpendicular arms as shown in Figure 1-3. Mirrors positioned at the ends of these arms, reflect the laser light back towards the beam splitter where it recombines forming an interference pattern.

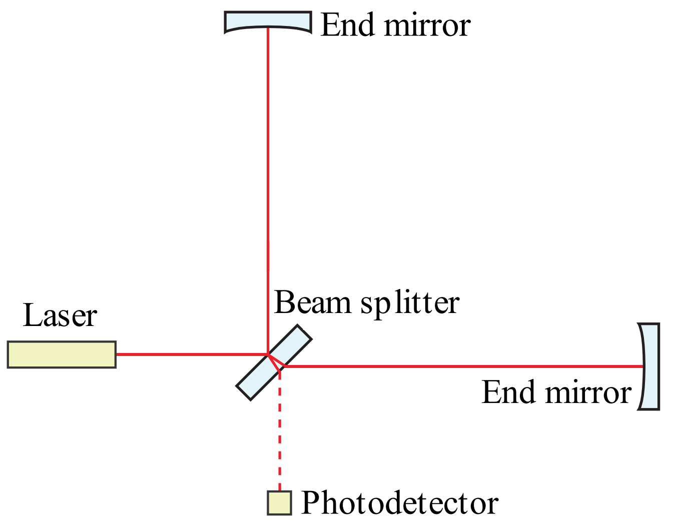

Figure 1-3: Michelson interferometer. A laser beam is split using a beam splitter. The split beams are reflected from the end mirrors (test masses coated with highly reflective coating) and recombined. The photodetector measures the recombined light signal.

In a gravitational wave detector, the mirrors are made by depositing a highly reflective coating upon a test mass. Furthermore, in a gravitational wave detector, the interferometer is designed such that when no effects of a gravitational wave are present, a dark fringe falls upon the photodetector. When a gravitational wave passes perpendicularly through the plane of the detector, the length of one of the arms increases by ΔL whilst the other decreases by ΔL, as shown in Figure 1-1. This relative change in arm length results in a phase shift between the two split laser beams causing the intensity of the interference pattern observed at the photodetector to change. Measuring this change in intensity enables the amplitude of the gravitational wave to be determined. Measuring the frequency of the change in intensity enables the frequency of the wave to determined.

The first working prototype of an interferometric gravitational wave detector was built in the 1970s by Robert Forward [47, 48]. It had an effective arm length of 8.5 m. In the following years, several other groups around the world built their own prototypes: Germany [49], USA [50] and UK [51, 52]. When the technology of these relatively small-scale interferometers (less than 40m) were shown to be reliable and scalable, longer baseline interferometers were built (between 300 m and 4 km) which had a more realistic probability of detecting gravitational waves. These detectors will be discussed in more detail in Section 1.8.

1.6 Limits to the sensitivity of gravitational wave detectors

A number of different noise sources limit the sensitivity of interferometric gravitational wave detectors. The following subsections will describe the dominant noise sources in more detail.

1.6.1 Seismic noise

The Earth's surface is continually moving. When this motion directly causes the detector's test masses to move, this is known as seismic noise. Seismic noise is a dominant noise source in ground based gravitational wave detectors below 10 Hz. Whilst wind and human activity are largely responsible for seismic noise at frequencies between 1 and 10 Hz, ocean waves travelling as surface waves in the Earth's crust are the predominant cause at frequencies below 1 Hz [53]. To mitigate the effects of seismic noise, gravitational wave detectors are purposefully built away from large population centres. However, even by building a detector at these locations which exhibit relatively low seismic noise, no gravitational waves could ever be observed as the effect of seismic noise upon the test masses is still too great (i.e. the effects of seismic noise upon the test masses at a relatively quiet location on the Earth is approximately 7 to 8 orders of magnitude larger than the effect caused by the first observed gravitational wave [54]). To reduce the effects of seismic noise upon the test masses a combination of passive and active isolation systems are used.

The passive isolation systems used in gravitational wave detectors incorporate pendulums and springs [55, 56]. Test mass mirrors are suspended as pendulums to isolate them from horizontally induced seismic motion. This is particularly effective at high frequencies as the mirrors horizontal displacement is proportional to $\frac{1}{f^2}$ at frequencies greater than the pendulum's resonant mode. Suspending the test masses using multiple pendulum stages further reduces these seismically induced horizontal displacements. However, one drawback of using pendulum stages is that they couple the vertical and horizontal motion of the test masses. Multiple-stage spring systems are therefore used in conjunction with the multiple stage pendulum systems to suspend the test mass in order to reduce any vertical motion. The currently operational gravitational wave detector called Advanced VIRGO employs a 7 stage pendulum system and a 6 stage spring system [56] to minimise seismic noise.

Another drawback of using pendulum systems to suppress the horizontal test mass motion is that these systems exhibit resonant modes which can limit detector sensitivity. To reduce the amplitude of the pendulum mode, active damping is used. This involves sensors to monitor the motion of the suspension stages and actuators to counteract their motion [57].

1.6.2 Gravitational gradient noise

Gravitational gradient noise is caused by local fluctuations in the gravitational field resulting in motion of the interferometer mirrors **[58, 59]**. This motion of the mirrors places a lower limit on ground-based detector sensitivity at frequencies less than 10 Hz. Numerous sources can be responsible for gravitational gradient noise, including: seismic waves which travel over the surface of the Earth, atmospheric changes, ocean dynamics and human activity **[60, 53]**. Unfortunately, the mirrors cannot be shielded from gravitational gradient noise. However, steps can be taken to reduce its effects. Locating detectors away from coastlines can mitigate the effects of micro seismic noise caused by ocean waves. Placing detectors underground can reduce the magnitude of human induced gravitational fluctuations. Positioning detectors in space will remove all earth based gravitational field fluctuations. Current research and development is investigating the use of seismometers to measure local ground vibrations, with the results being used to subtract the effect of gravitational gradient noise from the signal measured by the detector **[61]** .

1.6.3 Thermal noise

Thermal noise is caused by the thermally induced motion of atoms and molecules in the interferometer mirrors and their suspensions. This thermal motion causes the resonant modes of the mirrors and suspensions to be excited. Although the suspended mirror is designed such that the majority of its resonant modes are out with the detection band, the thermal noise caused by tails of the resonances still provide a major limit to a detector’s peak sensitivity. The magnitude of the mirrors and suspensions thermal noise is dependent upon the mechanical dissipation (or mechanical loss) of the materials used to form them. By choosing materials of low mechanical loss, the off-resonance thermal noise of the mirrors and suspensions will be reduced as more thermally induced motion will be concentrated near the resonant mode frequencies. Fused silica is a material which has been identified to have very low mechanical loss at room temperature. For this reason, the suspension fibres and test masses in Advanced LIGO and Advanced Virgo are made from silica [62]. A more detailed discussion of detector thermal noise and mechanical loss is discussed in Chapter 2.

1.6.4 Shot noise

Shot noise is caused by the statistical variation in the number of laser light photons incident upon the photodetector resulting in a fluctuation in the photoelectric current. If on average N photons are incident on a photodetector over a particular time, then the uncertainty in the number of photoelectrons is equal to $\sqrt{N}$ (as defined by Poisson statistics). This uncertainty places a limit on detector sensitivity as it hides the small intensity changes due to a gravitational wave. The smallest measurable strain due to shot noise is [16]:

$$h = \frac{1}{L}\left(\frac{\mu c\lambda}{4\pi^2 P}\right)^{\frac{1}{2}} , \tag{1.8}$$

where λ is the wavelength of laser light, μ is Planck's constant, P is the laser power and L is the interferometer arm length. Shot noise becomes larger at high frequencies and can be reduced by increasing laser power. However, increasing laser power has the adverse impact of increasing radiation pressure noise.

1.6.5 Radiation pressure noise

Radiation pressure noise is caused by fluctuations in the number of photons which are reflected off of the test mass mirrors. These reflected photons transfer momentum to the test mass mirrors which results in a differential displacement forming between them. This displacement places a lower limit upon detector strain sensitivity. This can be calculated at a particular frequency using [16]:

$$h = \frac{G}{Lmf^2}\left(\frac{\mu P}{4\pi^4 \lambda c}\right)^{\frac{1}{2}} , \tag{1.9}$$

where G represents the number of reflections off a test mass mirror and m is the mass of the mirror. The effects of radiation pressure noise become greater as frequency decreases while increasing the mass of the mirrors and decreasing the laser power can reduce these effects.

1.6.6 Standard quantum limit

The Standard Quantum Limit (SQL) describes the limit of detector sensitivity caused by the combination of radiation pressure noise and shot noise. Radiation pressure noise (Equation (1.9)) and shot noise (Equation (1.8)) are complementary noise sources which are affected oppositely by changes in laser power. At each frequency, there exists a laser power where this combination of noise sources is minimised. This frequency-dependent minimum noise level is known as the SQL. The SQL occurs when it is assumed that shot noise and radiation pressure noise are uncorrelated from one another. To surpass the SQL, gravitational wave detectors use signal recycling [63] and squeezing [64] to introduce correlations between these two noise sources.

1.7 Interferometric techniques

In practice gravitational wave detectors incorporate many advanced techniques to improve sensitivity, resulting in more complex detector designs than the Michelson interferometer shown in Figure 1-3.

1.7.1 Delay line configurations and Fabry–Perot cavities

Longer detector arms are beneficial as they increase the absolute displacement caused by gravitational waves. However, detector arm lengths are limited in length due to problems associated with the Earth's curvature and practical considerations such as cost. Both delay line configurations and Fabry–Perot cavities aim to improve detector sensitivity by increasing the apparent arm lengths of the detectors without increasing them physically.

In a delay line interferometer the optical path is folded [65], as illustrated in Figure 1-4. Figure 1-4 (a) illustrates a method to increase the standard Michelson interferometer's apparent arm length by using mirrors located near the beam splitter to reflect the laser light back towards the photodetector. The gravitational wave detector called GEO600 uses this configuration to increase its actual arm length of 600 m to an apparent length of 1200 m [66]. In Figure 1-4 (b), laser light passes through a small hole in the input mirror and reflects multiple times at different positions off of the input and end mirrors before exiting through the small hole. This technique was first proposed to be used in a gravitational wave detector by Weiss in 1972 [65], however due to scattering problems it was never adopted.

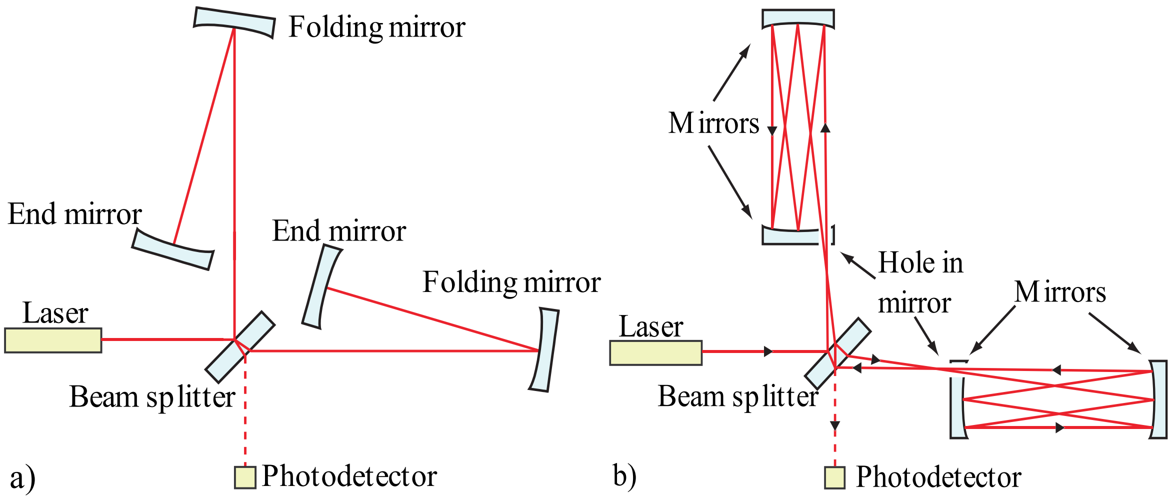

Figure 1-4: Schematics of different delay line interferometer configurations. These interferometers increase their apparent arm length by: a) using folded arms and b) reflecting the laser light along each arm multiple times.

Advanced LIGO and Advanced Virgo use Fabry-Perot cavities to increase the arm length [67, 68, 11]. Figure 1-5 illustrates an interferometer which uses Fabry-Perot cavities. The light from the laser passes through a partially transparent mirror where it then reflects back and forth along itself, increasing the laser's power in the cavity before exiting the same way it entered. In Advanced LIGO, the light reflects approximately 280 times before exiting, thus increasing the actual arm length of 4km to an apparent arm length of 1120 km long [69].

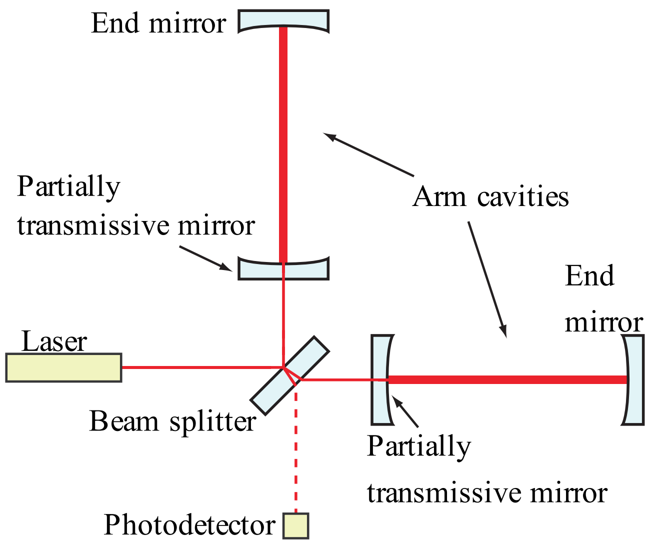

Figure 1-5: Fabry-Perot interferometer. The interferometer has a Fabry-Perot cavity in each arm.

1.7.2 Power and signal recycling mirrors

A gravitational wave detector is designed such that when no effects of a gravitational wave are present, a dark fringe is formed at the photodetector. Due to energy conservation laws, this means the interferometer's laser light is thus being lost in the direction of the laser (assuming no scattering losses). By placing a partially transparent mirror between the laser and the beam splitter, laser light which would have been otherwise lost from the interferometer can be reflected back into the detector. The addition of this mirror effectively creates a cavity between the power recycling mirror and the interferometer. As a result, the power in the interferometer increases and detector sensitivity improves as shot noise decreases. A schematic of an interferometer which uses a power recycling mirror is illustrated in Figure 1-6.

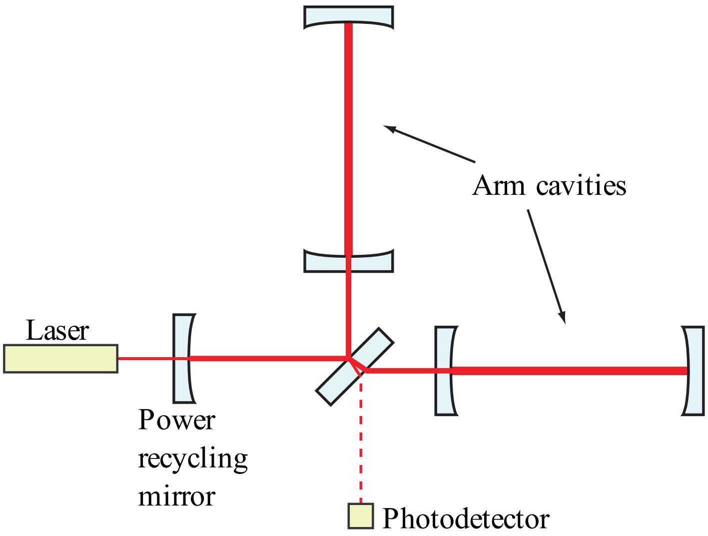

Figure 1-6: Schematic of an interferometer which incorporates a power recycling mirror. The position of the mirror forms a cavity between itself and the rest of the interferometer.

The size of a gravitational wave signal can also be amplified at a particular frequency by placing a 'signal recycling' between the beamsplitter and the photodiode to reflect any out-going light (signals) back into the interferometer where it resonates. See Figure 1-7. Positioning a mirror at this location in the detector creates a cavity between the mirror and the interferometer. By changing the position of this mirror, the length of the cavity will change as well as the frequency at which the interferometer can amplify a signal [70].

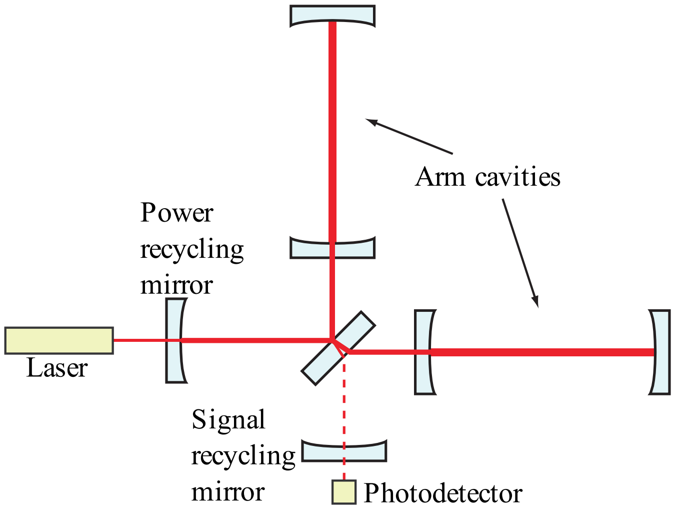

Figure 1-7: Schematic of an interferometer which incorporates a signal and power recycling mirror. The position of the signal recycling mirror forms a cavity between itself and the rest of the interferometer. Varying the length of this cavity enables the interferometer's sensitivity to be improved at specific frequencies.

1.7.3 Squeezed light

A classical description of shot noise and radiation pressure noise was presented in Sections 1.6.4 and 1.6.5. However, in order to understand how the SQL limit can be surpassed using squeezed light, a quantum mechanical treatment of these noise sources is required.

Quantum mechanics describes the origin of both the shot noise and radiation pressure noise to be due to vacuum fluctuations (photons randomly being created and destroyed in the vacuum field) entering the interferometer at the output port [64]. Vacuum fluctuations in the phase quadrature (imaginary part of the detected signal) results in shot noise whilst fluctuations in the amplitude quadrature (real part of the detected signal) results in radiation pressure noise.

Light is governed by the probabilistic rules of quantum physics. Heisenberg's uncertainty principle states that if two commutable physical observables of light are measured, then there is a fundamental limit to the precision in which these observables can be known. If the phase and amplitude (commutable observables) of a detector's laser light is measured and plotted then this would result in the values being scattered as shown in Figure 1-8 (a), where the diameter of the circle represents the quantum uncertainty in the laser light's phase and amplitude (which is determined by the magnitude of the vacuum fluctuations). The opaqueness / transparency of the blue colour highlights the probability of observation (where opaqueness represents a high probability). This equal uncertainty in phase and amplitude represents the SQL. By injecting squeezed laser light into a gravitational wave detector at the dark port of the beam splitter, the SQL can be reduced.

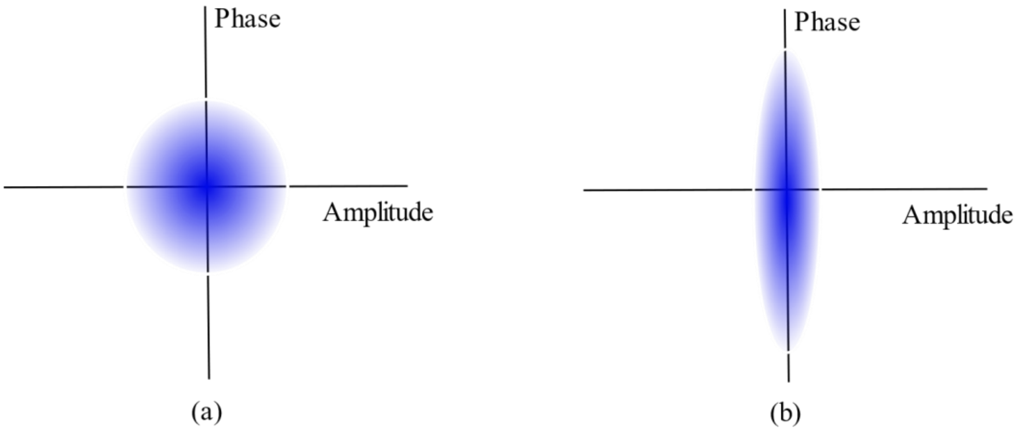

Figure 1-8: Spread in a laser light's phase and amplitude values. (a) shows the spread in values of unsqueezed light whilst (b) shows the spread using squeezed light. The larger the spread in values, the larger the measurement uncertainty. It can be observed that amplitude squeezed light (b), reduces the uncertainty in the unsqueezed laser light's amplitude but not it's phase. The opaqueness/transparency of the blue colour illustrates the probability of observation (where opaqueness represents a high probability).

Squeezed laser light describes light whose optical field has been manipulated such that there are more uncertainties in one quadrature (amplitude or phase) and less in the other. An example of a squeezed state is shown in Figure 1-8(b). Gravitational wave detectors inject an amplitude squeezed state of light into the dark port and measure the amplitude of the gravitational wave signal in order to improve detector sensitivity at low frequencies (as it reduces the dominant radiation pressure noise (amplitude fluctuations)). In doing this, the detector sensitivity has surpassed the SQL (at low frequencies). Likewise, a gravitational wave detector will surpass the SQL at high frequencies using a phase squeezed state of light and measuring the phase of the gravitational wave signal [64].

1.8 Current state of gravitational wave detectors

The first long baseline gravitational wave detectors were built in the 1990's. Four of these first-generation detectors were built: three LIGO detectors in the USA, Virgo in Italy, GEO600 in Germany and TAMA in Japan [71]. The detectors that are currently used are known as second generation detectors. They are first generation detectors which have been upgraded. These detectors are called Advanced LIGO and Advanced Virgo. The first-generation Japanese detector was decommissioned and, in its place, a second generation detector is currently being built, called KAGRA. In the future it is planned that the second-generation detectors will be upgraded again and a new Indian ground-based detector and European third-generation detector will be built. A space-based detector is also planned to be launched into orbit as well. Figure 1-9 shows the strain sensitivity of both current and future detectors. The following subsections describe the history, present and future of these individual detectors in more detail.

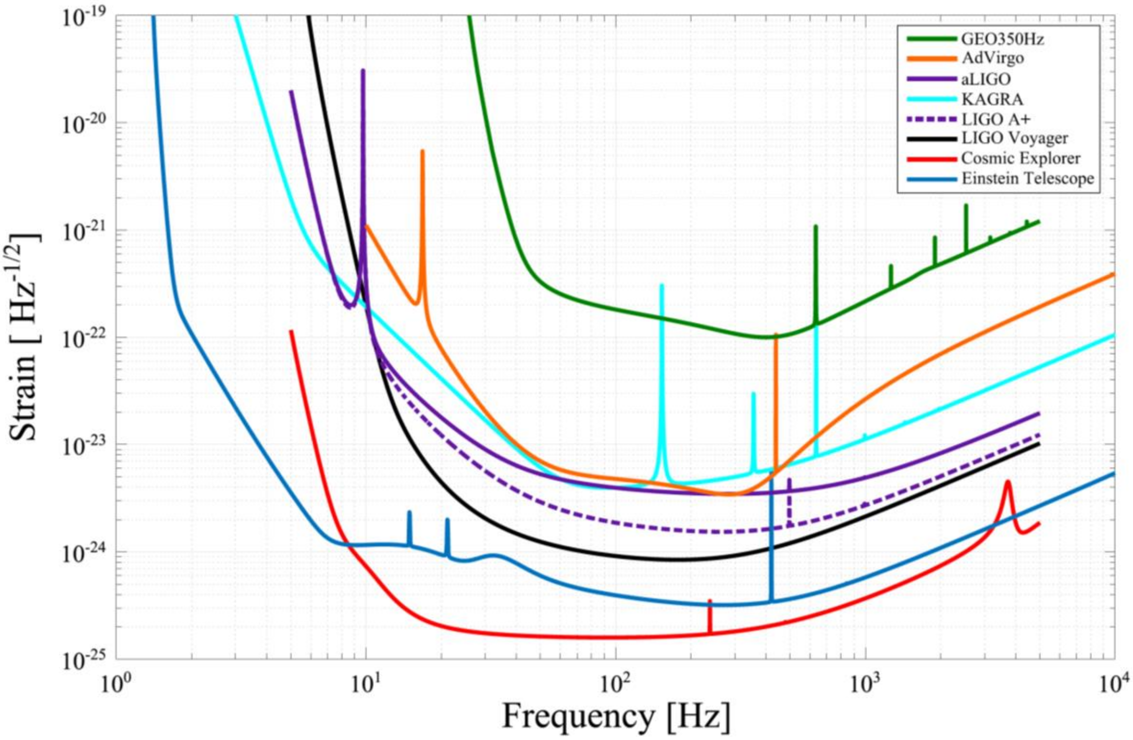

Figure 1-9: Strain sensitivity of current and future ground based gravitational wave detectors [54].

1.8.1 GEO600

GEO600 is a 600 m arm-length, ground-based gravitational wave detector near Hannover in Germany. The project was born from the collaboration of a British university and a German research institute (University of Glasgow and the Max Planck Institute for Quantum Optics) who had originally planned to build their own detectors but due to financial constraints pooled their resources together.

At the time of construction, GEO600 used a myriad of new technologies in order to maximise performance. These included a delay line interferometer to increase arm length, a triple pendulum suspension system to minimise seismic noise, power and signal recycling mirrors to minimise shot noise and hydroxy-catalysis bonding between suspension fibres and test masses to minimise thermal noise [72, 66]. The detector underwent a program of upgrades in early 2009 called GEO-HF [73]. Due to the design of GEO600 and the site which it is built on, no major changes to the detector infrastructure were made. Instead upgrades were limited to incorporating advanced technologies into the detector with the primary aim of increasing detector sensitivity above 500 Hz by decreasing shot noise. Some of these technological improvements included: increasing the input laser power to 35 W, an output mode cleaner and the introduction of squeezed light. The improvement in the detector's strain sensitivity can be observed in Figure 1-10.

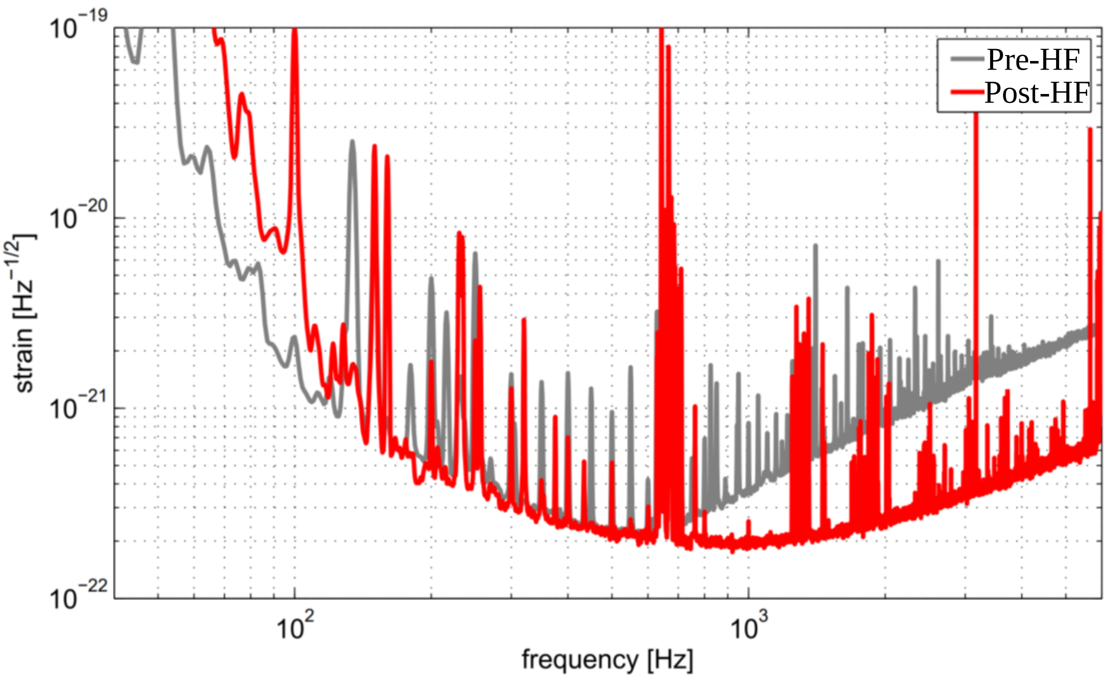

Figure 1-10: Strain sensitivity of the GEO600 detector pre and post GEO HF upgrades. It can be observed that above 500 Hz, the sensitivity of the detector was greatly increased post HF upgrades [74].

From 2009 to present day, GEO600 has not been able to measure any gravitational waves. However, GEO600 can claim to have helped other more sensitive detectors to achieve this goal by providing a long baseline interferometer where advanced technologies could be tested and finessed before being incorporated into their setups. As a result of infrastructure limitations, it is unlikely that GEO600 will be able to keep pace with the sensitivity of

Advanced LIGO or Advanced VIRGO and their future upgrades. For this reason, it is most likely that GEO600 will remain for the time being, a testing site for advanced technologies.

1.8.2 Advanced LIGO

Advanced LIGO consists of two gravitational wave detectors that are located near Baton Rouge, Louisiana and Hanford, Washington. The detectors and their sites are shown in Figure 1-11.

(a) (b)

Figure 1-11: Pictures of the Advanced LIGO (a) Hanford and (b) Livingston detectors [75] [76].

Caltech and MIT carried out the initial design and construction of these detectors. However, today more than 100 research institutes contribute to the development, maintenance and running of Advanced LIGO. The two detectors were initially designed such that they could be easily upgraded at different points in the future.

These initial detectors were retrospectively named iLIGO (initial LIGO) [77]. The iLIGO Hanford Observatory was built with a 4 km and 2 km interferometer contained within the same vacuum system. The iLIGO Livingston Observatory had only one 4 km length interferometer. Both iLIGO Hanford and Livingston Observatories used a 10 W laser and a power recycling mirror. Unlike GEO600, the iLIGO detectors did not use a delay line interferometer but instead used a Fabry Perot cavity to increase its optical path length. To suspend their highly reflective test masses, iLIGO used a single loop of steel wire. Construction on iLIGO began in 1994 and the detector became operational in the early 2000's, reaching peak design sensitivity of 2×10^{-23} Hz$^{-1/2}$ at 110 Hz in 2005 [78]. However, from the time of its first data run until 2007, iLIGO did not observe any gravitational waves. In 2007, the detector was taken offline and a small number of upgrades were made to it resulting in its name being changed to eLIGO (enahanced LIGO) [79]. The most significant enhancements made to the detectors were to increase the input laser power to 35 W as well as adding an output mode cleaning cavity. Both of these changes predominantly helped in reducing the shot noise of the detector at high frequencies, although the peak strain sensitivity was also improved by a factor of 2. After two years in observation mode, eLIGO failed in measuring any gravitational waves and in 2007 the detector underwent major upgrades leading to approximately a factor of 3 improvement in peak strain sensitivity when compared to eLIGO. This detector became known as Advanced LIGO [10]. Advanced LIGO had a 125 W laser input power to reduce the effects of shot noise, larger test masses to reduce radiation pressure noise and silica suspension systems and a new mirror coating to reduce the effects of thermal noise. The different Advanced LIGO noise contributions can be observed in Figure 1-12. Advanced LIGO detected its first gravitational wave in 2015 [6]. To date, Advanced LIGO has detected 6 gravitational waves (these are discussed in more detail in Section 1.4).

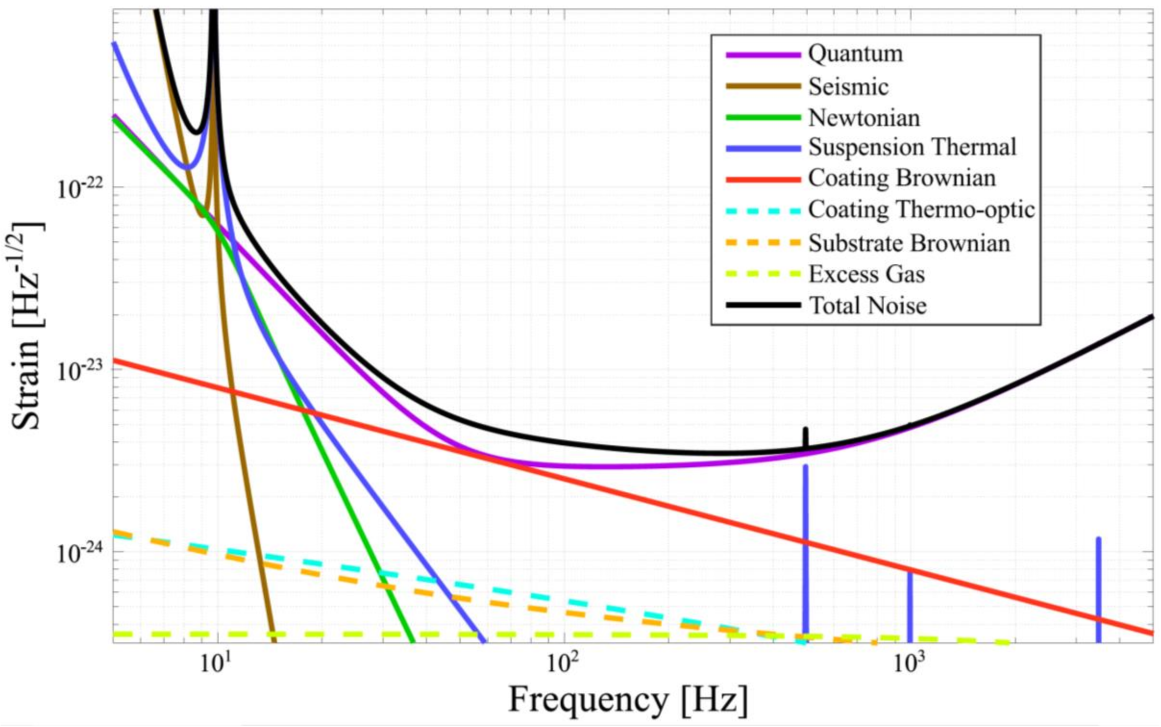

Figure 1-12: Advanced LIGO's noise sources. It can be observed that both quantum noise and coating Brownian noise are two dominant noise sources at peak sensitivity [10].

1.8.3 Advanced Virgo

Virgo was a ground based gravitational wave detector built in Cascina, Italy. Construction began on Virgo in 1996 and finished in 2003. The detector had 3 km arms, a 20 W input laser power, a power recycling mirror, Fabry-Perot cavities and pendulum and cantilever suspension systems made of steel [80]. Virgo began to measure data in 2004 and showed a peak sensitivity of 6×10^{-22} $Hz^{-1/2}$ at 300 Hz [16]. In 2011, after not observing any gravitational waves, minor upgrades were made to the detector. Post-upgrade the detector became known as Virgo+. The most notable changes to the initial Virgo design was to increase the laser input power to 50 W, use silica suspension fibres and increase the finesse of the Fabry-Perot cavities [80]. After again not observing any gravitational waves, the detector was upgraded another time with the aim of a factor ten improvement in strain sensitivity compared to the initial Virgo. This detector was called Advanced Virgo. Advanced Virgo had a 200 W laser input power, improved vacuum system, larger test masses and new optics which corrected for previous spherical aberrations (which affected Virgo and Virgo+) [11]. Advanced VIRGO joined Advanced LIGO in measuring data on the 1st of August. On the 15th of August, both Advanced Virgo and Advanced LIGO observed the same gravitational wave [8]. To date, Advanced Virgo has measured 1 gravitational wave [8] and contributed to the discovery of another [7] (these will be discussed in more detail in Section 1.8).

1.8.4 KAGRA

KAGRA is an underground, cryogenic detector. It is currently being built in the Kamioka mine in Japan. Construction of the detector began in 2012 and it is expected to finish in 2020. Much of the technology and design of this detector was based on previously built prototypes called TAMA 300 and CLIO [81] [82]. KAGRA is built more than 200 m underground where the seismic noise is 100 times less than that measured on the surface [83]. Furthermore, KAGRA's test masses and suspensions will be held at a constant temperature of 20 K to reduce the effects of thermal noise. Sapphire was used to construct the test masses and suspensions in KAGRA instead of silica (what is used in other detectors), as it exhibits a lower mechanical loss at lower temperatures. KAGRA has a 3 km long arm length and a 2 W input laser power. Similarly to other detectors, KAGRA incorporates a multi-stage suspension system, power and signal recycling mirrors as well as Fabry Perot cavities.

1.8.5 Future detectors

1.8.5.1 Evolution of Advanced LIGO

There are currently three further upgrades planned for Advanced LIGO. Once these upgrades are completed the detector will be called, in sequential order: A+, Voyager and Cosmic Explorer [84].

It is currently estimated that the Advanced LIGO detectors will be upgraded from 2019 to 2022 [85]. These upgraded detectors will be called A+ and will utilize squeezed light (to reduce quantum noise), have stronger test mass suspensions in order to support larger test masses (reducing radiation pressure noise) and will use a new mirror coating (to reduce coating thermal noise). These upgrades should improve sensitivity by a factor of 1.7 compared to Advanced LIGO.

It is proposed that A+ will be upgraded in 2025, with the upgraded detector (referred to as LIGO Voyager) having a factor of 2 better sensitivity. This major upgrade will operate at 120 K [84]. As a result of this change in temperature, all test masses and the beam splitter will be replaced with silicon, as it exhibits a lower mechanical loss at this temperature [86] [87]. Changing to silicon optics also necessitates a change to a 1550 nm input laser wavelength (as silicon is not transmissive at 1064 nm). Current coatings show an increase in mechanical loss as they are cooled, and therefore improved coatings are required to reach desired sensitivity at 120 K.

There are longer-term plans for a completely new LIGO detector, known as Cosmic Explorer [84]. This detector is forecasted to be fully operational in 2035. The details of this detector are not finalised as much of its design will depend on the performance of the new technologies incorporated within A+ and Voyager. However, one currently discussed design is to increase the arm length to 40 km, increase the input laser power to 300 W and maintain the silicon optics, cryogenic operating temperatures as well as the 1550 nm laser wavelength used in Voyager. It is believed that by making these upgrades, Cosmic Explorer should exhibit a factor of 4 better strain sensitivity than Voyager.

1.8.5.2 LIGO India

There are plans to build an Advanced LIGO detector in India, which would become fully operational by 2020 [88]. This detector will be built by a collaboration involving Indian research institutes and LIGO Labs. LIGO Labs will supply the schematics, parts and expertise to build an identical detector to one of the Advanced LIGO detectors and the Indian research institutes will be responsible for its operation and maintenance. Building a detector in India is extremely useful for the gravitational wave detector community as the geographical location relative to the other detectors will enable a factor of 3 improvement in sky localisation [71].

1.7.5.3 Einstein Telescope

The Einstein telescope (ET) is a proposed underground, European gravitational wave observatory which is still in its early stages of design and not expected to become operational until 2030 at the earliest [89]. The design comprises three separate detectors positioned in a triangular configuration as shown in Figure 1-13.

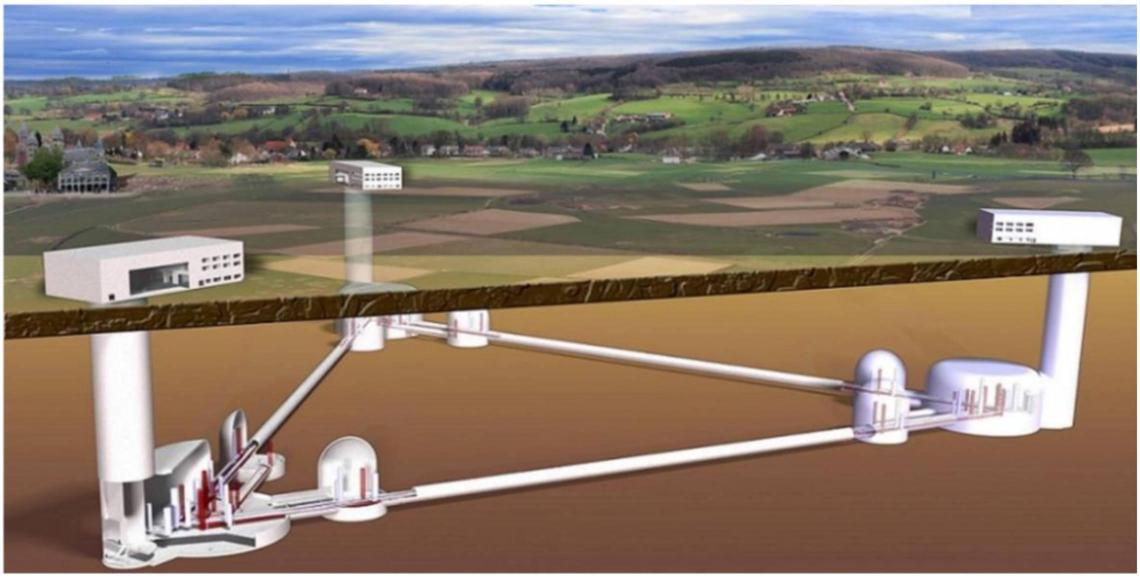

Figure 1-13: Graphic illustrating the planned design of ET. Three different detectors orientated in a triangular configuration buried under the ground [90].

Each detector has access to two interferometers which have 10 km Fabry-Perot arm lengths. One of these interferometers is designed to have peak sensitivity at low frequencies (1.5 Hz to 30 Hz) whilst the other is designed for high frequencies (30 Hz to 10 kHz). The low frequency detector operates at cryogenic temperatures whilst the high frequency detector operates at room temperature. The design of ET not only increases detector sensitivity at a broad range of frequencies, but it also enables the observatory to more accurately identify the direction of origin of gravitational waves as well as measure the effects of passing waves which travel in the same plane as the detector. Significant research and development is still required to realise ET.

1.8.5.4 Laser Interferometric Space Antenna

The Laser Interferometric Space Antenna, known as LISA, is a space based gravitational wave detector planned to be launched in 2030 by the European Space Agency. LISA consists of three identical spacecraft positioned at the corners of an equilateral triangle with length 2.5×10^6 km [91]. This triangular configuration of satellites is to revolve around the sun in an Earth like orbit, approximately 20 degrees behind Earth as shown in Figure 1-14. The plane of the triangle is to be inclined 60 degrees relative to the ecliptic to minimise the relative change in distance between the different spacecraft [91].

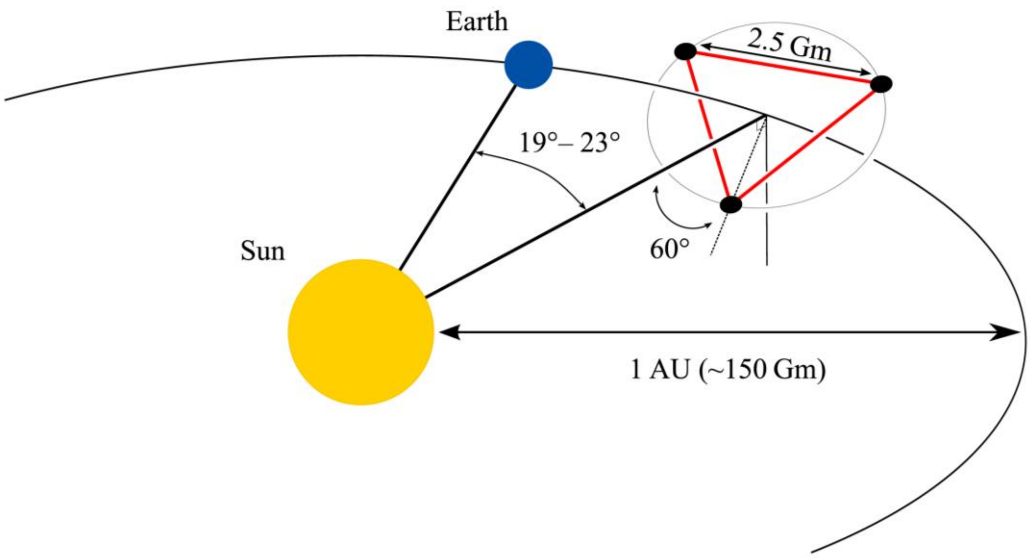

Figure 1-14: Orbit of LISA relative to the Earth and Sun. The three spacecraft are orientated in a triangular configuration and the detectors arms are outlined in red [91].

Each spacecraft contains two 2 W 1064 nm lasers, two free floating test masses as well as monitoring and control systems which are able to detect and realign the position of the test masses. LISA is designed such that each satellite in conjunction with the other two is able to form its own Michelson interferometer with an arm length of 2.5×10^6 km. Due to the immense length of this detector's arms, LISA does not reflect laser beams from one satellite back to the other as the laser experiences significant divergence along its path. Instead LISA's spacecraft phase lock a laser to the incoming detected light and transmit the phase locked light back.

LISA complements current ground-based detectors as it enables gravitational waves to be observed in a frequency region where ground-based detectors cannot. LISA is designed to operate between 20 uHz and 1 Hz. In this range of frequencies LISA should be able to see gravitational waves emitted by supermassive binary black holes as well as extreme mass ratio inspirals.

1.9 Motivation

The recent detections of gravitational waves has demonstrated the power and potential of gravitational wave astronomy, provided further proof of General Relativity, proof of the existence of black holes with a mass equal to a few tens of solar mass and proof of a gamma-ray burst originating from merging neutron stars. As new detectors come online and the sensitivity of all detectors improves, the potential of gravitational wave astronomy will only become greater as more signals from different sources will be detected. To maximise the benefits of gravitational wave astronomy, further research and development is required to tackle current limiting noise sources. Coating thermal noise is one of the dominant noise sources in current detectors at their most sensitive frequencies and will be discussed in detail in Chapter 2. Whilst chapters 3 - 6 will focus upon presenting the authors research and contribution to the field of gravitational wave coatings.

2 Coating thermal noise

2.1 Introduction

The mirrors in a gravitational wave detector are formed by depositing a highly reflective coating upon a test mass. The thermal noise of the coating is one of the main limits to the sensitivity of currently operating gravitational wave detectors.

Coating thermal noise describes the combination of three different noises: Brownian noise, thermoelastic noise and thermo-refractive noise. Coating Brownian thermal noise limits the sensitivity of ground-based gravitational wave detectors at their most sensitive frequencies. According to the Equipartition Theorem, every mechanical system has a mean value of $\frac{1}{2}k_bT$ thermal energy associated with each degree of freedom. This thermal energy causes random, thermally driven molecular vibrations in the coating at a microscopic level and leads to the excitation of the resonant modes of the coated test masses at a macroscopic level, thus changing the detector's arm-length. This random change in arm-length introduces noise into the detector. Thermoelastic noise and thermo-refractive noise are caused by random temperature fluctuations which also have the effect of changing the actual and optical arm length of the detector. Current research is aimed at minimizing the effects of these noise sources in order to improve current detector sensitivity as well as enable the design sensitivity of third-generation detectors to be achieved. This chapter will outline the general theory and mathematics of thermal noise before individually describing coating Brownian noise, thermoelastic noise and thermo-refractive noise.

2.1.1 Brownian noise

Brownian thermal noise is caused by random, thermally driven motion of the mirror test-masses, coatings and the mirror suspensions. Brownian motion was first observed by Robert Brown in 1828 when he observed the random, thermally driven motion of pollen grains suspended in water [92]. Einstein explained the origin of this random motion as a result of collisions between the pollen grains and the surrounding water molecules [93], during which the pollen grains lose their initial kinetic energy. Einstein therefore connected the fluctuation the pollen grains with the dissipation caused by the intrinsic internal friction of the water molecules.

2.1.2 The Fluctuation-Dissipation Theorem

The Fluctuation-Dissipation Theorem, derived by Callen et al in a series of papers published between 1951 and 1952 [94, 95, 96] enables the thermal noise of a system to be determined. The theorem states that any linear system in thermal equilibrium will exhibit random, thermally driven fluctuations whose magnitude and angular frequency (ω) is related to the dissipative (i.e. real) part of the system's impedance ($\Re[Z(\omega)]$):

$$S_f(\omega) = 4k_bT\Re[Z(\omega)]\,, \tag{2.1}$$

where $S_f(\omega)$ represents the power spectral density of the fluctuating thermal force, k_b is the Boltzmann constant and T is the temperature of the system. $Z(\omega)$ represents the system's impedance which is equal to:

$$Z(\omega) = \frac{F(\omega)}{v(\omega)}\,, \tag{2.2}$$

where $F(\omega)$ is the force applied to the system and $v(\omega)$ is the system's velocity caused by the applied force. Alternatively, the Fluctuation-Dissipation Theorem can be expressed in the form:

$$S_x(\omega) = \frac{4k_bT}{\omega^2}\Re[Y(\omega)]\,, \tag{2.3}$$

where $S_x(\omega)$ represents the power spectral density of the system's fluctuating displacement (thermal noise) and $Y(\omega)$ is the mechanical admittance of the system (which is equal to the inverse of impedance). To summarise, the Fluctuation-Dissipation Theorem states that thermal noise is related to energy dissipation and the magnitude of noise is related to the magnitude of dissipation.

2.2 Dissipation

Most broadly speaking, there are two types of dissipation mechanisms in a detector: external and internal.

2.2.1 External Dissipation

External dissipation describes the process of energy being lost to the environment. In a gravitational wave detector, examples of external dissipation are:

- Gas damping - energy is lost from the test mass mirrors and suspensions to residual gas molecules (where the suspensions are described in Section 1.5.1).
- Recoil damping - energy is lost from the pendulum into the surrounding support structure.
- Frictional damping - energy is lost at the pendulum suspension points.

By designing detectors with these dissipative mechanisms in mind, the magnitude of their effect on detector sensitivity can be minimised. Therefore the dominant source of thermal noise in a detector is due to internal dissipation.

2.2.2 Internal dissipation

Internal dissipation is caused by a material's anelasticity [97]. When a stress (σ) acts upon an ideal elastic material, a strain (ε) of magnitude:

$$\varepsilon = \frac{\sigma}{Y}, \tag{2.4}$$

instantaneously develops within the material, where Y represents the material's Youngs modulus. However, when a stress is applied to an anelastic material the strain response is not instantaneous and instead develops over time. Applying a periodic stress of angular frequency (ω) of the form:

$$\sigma = \sigma_0 e^{i\omega t}, \tag{2.5}$$

to an anelastic material will result in a strain of the form:

$$\epsilon = \epsilon_0 e^{i(\omega t - \phi)}, \tag{2.6}$$

where σ_0 and ϵ_0 are stress and strain amplitudes and ϕ is the phase by which the strain lags behind the stress. ϕ is known as the mechanical loss angle and is a measure of the ratio of the energy dissipated per oscillation ($E_{\mathrm{lostpercycle}}$) of the system relative to the total energy stored in the system (E_{stored}):

$$\phi(f) = \frac{E_{\mathrm{lostpercycle}}}{2\pi E_{\mathrm{stored}}}\,, \tag{2.7}$$

where f is the frequency of oscillation. The mechanical loss of a system can be most easily measured at a system's mode frequencies.

The internal dissipation of anelastic materials can arise from the rearrangement of molecular structures or defects such as interstitial impurities, point defects, dislocations and grain boundaries in response to an applied stress [97].

Knowing a system's mechanical loss enables the systems thermal noise to be determined.

2.2.2.1 Thermal noise associated with a single resonant mode

The relationship between the mechanical loss and thermal noise for a single resonant mode can be illustrated by considering the example of a harmonic oscillator.

An anelastic spring supporting a mass m can be described by the modified version of Hooke's law:

$$F(\omega) = -k\big(1 + i\phi(\omega)\big)x\,, \tag{2.8}$$

where F is the restoring force of the spring and $k(1 + i\phi(\omega))$ is the complex spring constant. Assuming that the system experiences an internal thermal driving force $\mathrm{A}(\omega)$, then the system's equation of motion takes on the form:

$$m\ddot{x} = -k\big(1 + i\phi(\omega)\big)x + A(\omega)\,. \tag{2.9}$$

The displacement (x), velocity ($\dot{x}$) and acceleration ($\ddot{x}$) of a harmonic oscillator are related as follows:

$$x = \frac{\dot{x}}{i\omega}\,, \tag{2.10}$$

$$\ddot{x} = i\omega\dot{x}\,. \tag{2.11}$$

Equation (2.9) can then be written in the form:

$$A(\omega) = i\omega m\dot{x} + \frac{k}{i\omega}\big(1 + i\phi(\omega)\big)\dot{x}\,. \tag{2.12}$$

Manipulating Equation (2.12) into the form of Equation (2.2) enables the impedance of the system to be calculated as:

$$Z = \frac{k + i\phi(\omega)k - \omega^2 m}{i\omega}\ . \tag{2.13}$$

The admittance is equal to:

$$Y = \frac{1}{Z} = \frac{i\omega}{k + i\phi(\omega)k - \omega^2 m}\ , \tag{2.14}$$

and the real part of the admittance is equal to:

$$\Re[Y(\omega)] = \frac{\omega\phi(\omega)k}{(k - \omega^2 m)^2 + \phi^2(\omega)k^2}\ . \tag{2.15}$$

The power spectral density of the system's displacement noise (i.e. thermal noise) can then be calculated using Equation (2.3):

$$S_x(\omega) = \frac{4k_bTk\phi(\omega)}{\omega(k - m\omega^2)^2 + k^2\phi^2(\omega)}\ , \tag{2.16}$$

Substituting $k = \omega_0^2 m$ into Equation (2.16), the thermal noise spectral density of the oscillator with a resonant angular frequency (ω_0) can be written as:

$$S_x(\omega) = \frac{4k_bT\omega_0^2\phi(\omega)}{\omega m[\phi^2(\omega)\omega_0^4 + (\omega_0^2 - \omega^2)^2]}\ , \tag{2.17}$$

It is interesting to note that at frequencies equal to the systems resonance ($\omega = \omega_0$), the thermal noise becomes:

$$S_x(\omega) = \frac{4k_bT}{\omega_0^3 m\phi(\omega_0)}\ , \tag{2.18}$$

which shows an inversely proportional relationship between thermal noise and mechanical loss. At frequencies far below the resonance, where ω<<ω_0 (assuming $\phi^2(\omega)$<<1), Equation (2.17) can be approximated by:

$$S_x(\omega) = \frac{4k_bT\phi(\omega)}{\omega_0^2 m\omega} \ , \tag{2.19}$$

whilst when $\omega >> \omega_0$, Equation (2.17) can be approximated by:

$$S_x(\omega) = \frac{4k_bT\omega_0^2\phi(\omega)}{m\omega^5} \ , \tag{2.20}$$

Figure 2-1 shows the thermal noise of two otherwise identical oscillators with different mechanical loss. Both oscillators exhibit the same total amount of thermal noise (when integrated over all frequencies), however it can be observed that the oscillator made from a lower mechanical loss material has a lower thermal noise off-resonance and a higher thermal noise on resonance than the oscillator made from a higher mechanical loss material. Gravitational wave detectors are purposefully built such that the resonant modes of their mirrors are at frequencies much higher than the most sensitive operating frequencies of the detector (approximately 300 Hz for the most sensitive operating frequency). By designing the detectors this way and constructing the mirrors from materials which exhibit as low a mechanical loss as possible, the magnitude of thermal noise at the detector's most sensitive operating frequencies can be minimised.

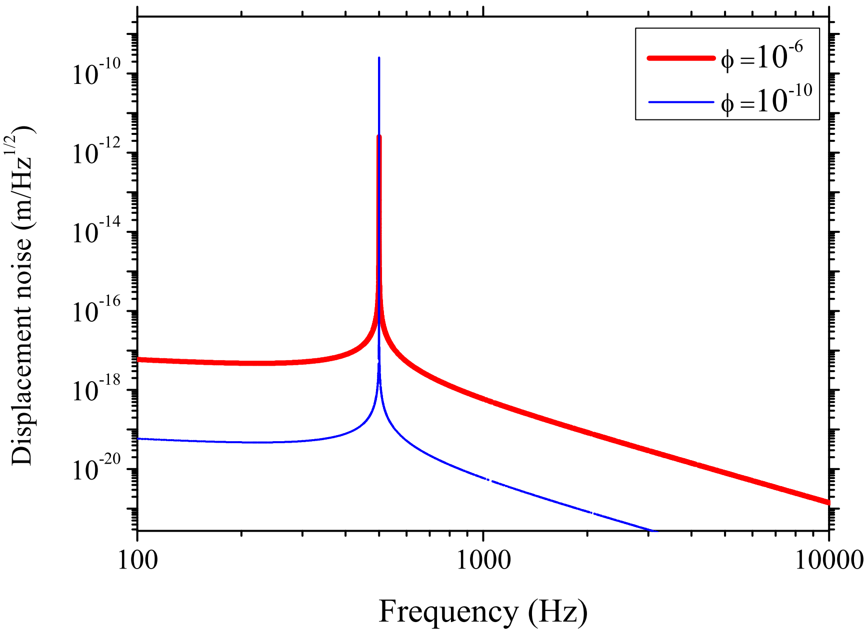

Figure 2-1: Thermal noise spectral densities of two oscillators. One of the oscillators has a mechanical loss of 1×10^{-6} *whilst the other has* 1×10^{-10}*. Both oscillators have a temperature of 290 K, a mass of 20 kg and a resonant mode frequency of 500 Hz.*

2.3 Brownian noise in a detector mirror

A gravitational wave detector mirror is a much more complex system than a simple harmonic oscillator. Thermally induced vibrations at the microscopic level will ultimately lead to the various mechanical resonances of the system to be excited. When the modes of the mirror are excited, the arm length of the interferometer changes resulting in a displacement noise. The dissipation mechanism associated with Brownian noise is the intrinsic internal friction of the materials used to construct the test mass and mirror coatings [98, 99].

Initial attempts to calculate the Brownian noise of a mirror assumed that the motion of each of the mirror's resonant modes as independent from one another (known as the normal mode expansion method) [100, 101]. This enabled the mirror's Brownian noise to be determined by summing the noise contributions from each mode. However, the assumption that a body's resonant modes are independent from one another is only true when the mechanical dissipation is homogeneously distributed throughout the body. A gravitational wave detector mirror is formed from a large uniform test mass which is coated on one side with a highly reflective coating. The dissipation of these coatings is several orders of magnitude greater than that of the test mass, introducing an inhomogeneous spatial distribution of loss into the system. Treating the noise contributions from the mirror's resonant modes independently from one another is therefore not accurate as the inhomogeneous dissipation introduces correlations between the Brownian noise of different modes [102]. An alternative approach to the *normal mode expansion method* was developed to more precisely calculate the Brownian noise of the detector mirrors.

The *advanced mode expansion method* enables the mirror's Brownian noise to be calculated as it uses cross mode coupling terms to account for the correlations developed between different mode's Brownian noise [103]. However, this method is extremely time consuming as a large number of modal solutions are required in order to obtain an accurate value of the mirror's Brownian noise.

In 1998, Levin outlined a new method to calculate a mirror's Brownian noise [102] which assumes that a notional pressure (with the same spatial profile as the detector's laser beam) is applied to the front face of the mirror. By calculating the power dissipated in the mirror (W_{diss}) from this pressure and using the Fluctuation-Dissipation Theorem, the Brownian noise power spectral density of the mirror can then be calculated using [102]:

$$S_x(f) = \frac{2k_bT}{\pi^2 f^2}\left(\frac{W_{\text{diss}}}{F_0^2}\right) , \tag{2.21}$$

where F_0 is the peak amplitude of the oscillating force applied to the mirror surface. If the mirror experiences inhomogeneously distributed dissipation then W_{diss} is equal to:

$$W_{diss} = 2\pi f \int_{vol} \epsilon(x,y,z)\phi(x,y,z,f)dV \ , \tag{2.22}$$

whilst for homgeneously distributed dissipation:

$$W_{diss} = 2\pi f U_{\max}\phi(f) \ , \tag{2.23}$$

where ϵ is the energy density of elastic deformation when the mirror is maximally deformed by the notional pressure, U_{max} is the mirror's energy at peak elastic deformation, V is the mirrors volume and x, y and z are directions in space. Bondu first applied Levin's method to determine the Brownian noise of an uncoated mirror (i.e. test mass). In his calculation Bondu assumed that the diameter of the test mass was much greater than the laser diameter as it enabled him to approximate the test mass as being half infinite. Bondu determined the Brownian thermal noise to be equal to [104]:

$$S_x(f) = \frac{2k_bT}{f\sqrt{\pi^3}}\left(\frac{1-\nu^2}{Yw_o}\right)\phi_{\text{substrate}}(f)\,, \tag{2.24}$$

where w_o is the distance over which the electric field amplitude of the laser beam falls to $\frac{1}{e^2}$ of its maximum value, $\phi_{\text{substrate}}$ is the mechanical loss of the test mass and Y and ν represent the Young's modulus and Poisson ratio of the test mass's material respectively.

Nakagawa et al [105] were the first to use Levin's method to determine the Brownian noise of a mirror (i.e. coated test mass). Nakagawa et al argued that the mirror's thermal noise could not be well-approximated by an uncoated test mass since the detector's laser light (which is used to sense for fluctuations) is reflected directly by the coating. Their logic (as well as Levins) was that a source of dissipation closer to the reflecting surface of the mirror would contribute more to a mirror's Brownian noise than if it were further away. Since the laser light's first point of contact is the coating, the effect of the coating would have to be considered. In their derivation, Nakagawa et al approximated the coating to be a uniform layer which has the same material properties as the test mass but with a different mechanical loss (ϕ_{coating}). Their expression to determine the Brownian noise of a mirror was [105]:

$$S_x(f) = \frac{2k_bT}{\pi^{\frac{3}{2}}f}\left(\frac{1-\nu^2}{w_0Y}\right)\left(\phi_{\text{substrate}} + \frac{2}{\sqrt{\pi}}\left(\frac{1-2\nu}{1-\nu}\right)\left(\frac{d}{w_0}\right)\phi_{\text{coating}}\right)\,, \tag{2.25}$$

where d is the thickness of the coating. Harry et al [106] derived a more detailed expression for a mirror's Brownian noise. Their expression explicitly accounted for the layer structure of the coating applied to the test masses. In a gravitational wave detector, coatings are formed from alternating layers of two materials e.g. in Advanced LIGO these materials are silica and titania doped tantala. Incorporating this layer structure leads to anisotropy of the coating's mechanical loss. This can be observed in their expression of a mirror's Brownian noise:

$$S_x(f) = \frac{2k_bT}{\pi^{\frac{3}{2}}f}\left(\frac{1-\nu^2}{w_0Y}\right)\left(\begin{matrix}\phi_{\text{substrate}} + \frac{1}{\sqrt{\pi}}\left(\frac{d}{w_0}\right)\left(\frac{1}{YY'(1-\nu'^2)(1-\nu^2)}\right)\times \\ \left[\begin{matrix} Y'^2(1+\nu)^2(1-2\nu)^2\phi_{\parallel} + \\ YY'\nu'(1+\nu)(1+\nu')(1-2\nu)(\phi_{\parallel}-\phi_{\perp}) \\ +Y^2(1+\nu')^2(1-2\nu')^2\phi_{\perp}\end{matrix}\right]\end{matrix}\right), \qquad (2.26)$$

where $\phi_{\parallel}$ and $\phi_{\perp}$ are the coating's mechanical losses associated with strains parallel and perpendicular to the coating surface respectively. In their derivation, Harry et al also accounted for the coating exhibiting different Young's modulus (Y') and Poisson ratio (ν') values from the test mass. When $Y' = Y$, $\nu' = \nu$ and $\phi_{\parallel} = \phi_{\perp}$, Equation (2.26) simplifies to Equation (2.25).

2.3.1 Coating Brownian noise

In an ideal world, a coating's Brownian noise is determined using Equation (2.26) when $\phi_{\text{substrate}} = 0$. However, in practice a coating's Brownian noise is calculated using Equation (2.26) whilst assuming $\phi_{\parallel} = \phi_{\perp}$, This assumption is required as whilst ring down experiments (Section 3.3) estimate $\phi_{\parallel}$, there is no current approach able to estimate $\phi_{\perp}$. In order to estimate a value of thermal noise, it is therefore assumed in the literature that $\phi_{\parallel} = \phi_{\perp}$ [106].

2.3.2 Hong's interpretation of coating's Brownian noise

Presently, the thermal noise of a coating is estimated using Equation (2.26) (and assuming that $\phi_{\text{substrate}} = 0$ and $\phi_{\parallel} = \phi_{\perp}$). However, Hong et al argues that separating the mechanical loss of a gravitational wave detector coating into its parallel and perpendicular components is fundamentally flawed as it can lead to unphysical situations where the energy stored in parallel and perpendicular motion is negative [107]. Hong therefore presents a new and alternative method of estimating a coating's thermal noise which is dependent upon the coating's bulk and shear losses (see Chapter 3).

The parallel and perpendicular losses of a gravitational wave detector coating arises from its layer structure (as it was thought that the loss of a coating made from alternating layers of isotropic, amorphous materials may be different as it moves parallel and perpendicularly to its surface). Hong's theory arises from the requirement of amorphous, isotropic materials needing two independent elastic constants to fully describe their mechanical properties.

Since each elastic constant represents a source of loss (imaginary component of elastic constant) then there exist two sources of loss for each isotropic amorphous material. Therefore whilst there is no real physical reason to suspect a single silica coating layer having different parallel and perpendicular losses, there is a real physical reason as to why it may have different bulk and shear losses.

2.4 Coating thermo-optic noise

Coating thermo-optic noise describes the combination of coating thermoelastic noise and coating thermo-refractive noise. These two noise sources are often considered under the umbrella term of coating thermo-optic noise as they are both caused by the same fundamental mechanism (random temperature fluctuations) and both exhibit the same dissipation mechanism (heat transfer) [99].

2.4.1 Coating thermoelastic noise

Coating thermoelastic noise arises from random temperature fluctuations in the coating which cause changes in the dimensions of the coating as a result of the coating's thermal expansion coefficient. When the coating's dimensions change, the length of the interferometer's arm length changes as well, resulting in displacement noise. The power spectral density of a coating's thermoelastic noise is given by [108, 109]:

$$S_x(f) \approx \frac{8k_bT^2d_c^2}{w_0^2\sqrt{\pi^3 f}}(1+\nu_s)^2\left(\frac{C_{avg}^2}{C_s^2}\right)\left(\frac{\alpha_s^2}{\sqrt{K_sC_s}}\right)\bar{\bar{\Delta}}^2 \ , \tag{2.27}$$

where C is the specific heat capacity, α is the thermal expansion coefficient and K is the thermal conductivity. $\bar{\bar{\Delta}}^2$ is equal to:

$$\bar{\bar{\Delta}}^2 = \left(\frac{C_s}{2\alpha_sC_{avg}}\left[\frac{\alpha_{avg}}{1-\nu_{avg}}\left(\frac{1+\nu_{avg}}{1+\nu_s}+\frac{(1-2\nu_s)Y_{avg}}{Y_s}\right)\right] -1\right)^2 , \tag{2.28}$$

where subscript s denotes substrate material properties and subscript avg represents coating material properties which have been calculated by performing a weighted average of the two alternating coating layer's material properties:

$$C_{avg} = C_a\frac{d_a}{d_a+d_b}+C_b\frac{d_b}{d_a+d_b} \ , \tag{2.29}$$

where subscripts a and b represent the two different alternating layers of the coating.

2.4.2 Coating thermo-refractive noise

In a gravitational wave detector, laser light is reflected by mirrors positioned at the ends of the detector's arms. This laser light slightly penetrates the coating of the mirror before being reflected back into the detector. Coating thermo-refractive noise is caused by random temperature fluctuations in the coating which change the refractive index of the coating materials and therefore also changes the optical path length of the laser beam in the coating. The magnitude of a coating's thermo-refractive noise is dependent upon the coating's thermo-optic coefficient (β):

$$\beta = \frac{dn}{dT} \ , \tag{2.30}$$

where n is the refractive index. The power spectral density of a coating's thermo-refractive noise can be calculated using [110]:

$$S_x(\omega) = \frac{\sqrt{2}\left(\beta_{eff}^2 \lambda^2 k_b T^2\right)}{\pi w_0^2 \sqrt{\omega \rho_c C_c K_c}} \ , \tag{2.31}$$

where λ is the wavelength of the laser, ρ_c is the density of the coating, K_c is the thermal conductivity of the coating and C_c is the specific heat capacity of the coating. β_{eff} is the effective temperature dependence of the high and low refractive index layers in the coating:

$$\beta_{eff} = \frac{n_L^2 \beta_L + n_H^2 \beta_H}{4(n_L^2 - n_H^2)} \ , \tag{2.32}$$

where the subscripts H and L represent the high and low index materials in the coating.

2.4.3 Coating thermo-optic noise

Since both coating thermoelastic noise and coating thermo-refractive noise are caused by the same random temperature fluctuations in the coating, they can be treated as one noise source known as thermo-optic noise. Using Levin's methodology, Evans et al calculated the coating's thermo-optic power spectral density to be equal to [111]:

$$S_x(f) = \frac{2k_b T^2}{w_0^2 \sqrt{\pi^3 K_s C_s f}} \left(\bar{\alpha}_c d - \bar{\beta}\lambda - \frac{\bar{\alpha}_s d C_c}{C_s} \right)^2 \ , \tag{2.33}$$

where d is the total coating thickness and $\bar{\alpha}$ is an effective thermal expansion coefficient. $\bar{\alpha}$ can be calculated using:

$$\bar{\alpha}_X = 2\alpha_X(1 + \nu_X) \ , \tag{2.34}$$

where $X \in (C, S)$, subscript C represents coating material properties, subscript S represents test mass material properties, $\bar{\beta}$ is the effective temperature dependence of the high and low refractive index layers in the coating:

$$\bar{\beta} = \frac{B_H + B_L\left(2\left(\frac{n_H}{n_L}\right)^2 - 1\right)}{4(n_H^2 - n_L^2)} \ , \tag{2.35}$$

$$B_Y = \beta_Y + \bar{\alpha}_Y n_Y \ , \tag{2.36}$$

and $Y \in (H, L)$. Interestingly, Equation (2.33) shows that the coating thermoelastic noise ($\bar{\alpha}_c d$) and coating thermo-refractive noise ($\bar{\beta}\lambda$ and $\frac{\bar{\alpha}_s d C_c}{C_s}$) can act to cancel one another out. Intuitively this makes sense and can be best observed by way of an example. Consider a random increase in coating temperature. This temperature change will result in the coating expanding (if the coating's thermal expansion coefficient is positive), thus reducing the interferometer arm length. However, this same increase in coating temperature will also result in an increase in path length and optical path length within the coating (assuming that $\bar{\beta}$ is positive). In this example, it can therefore be observed that the displacement noise of these combined sources will act to cancel one another out. Carefully selecting a coating's materials within a gravitational wave detector enables the extinction of a coating's thermo-optic noise.

2.5 Effect of coating thermal noise upon detector sensitivity

Coating thermal noise was not a limiting noise source in first-generation gravitational wave detectors. However, due to detector upgrades (and the magnitude of other noise sources being reduced) coating thermal noise limits the sensitivity of second-generation gravitational wave detectors at their most sensitive operating frequencies. This can be observed in Figure

1-11, which shows the different noise sources present in the Advanced LIGO detectors. The effects of coating Brownian noise is much larger than coating thermo-optic noise in second generation gravitational wave detectors. In the future, the magnitude of coating thermal noise will be reduced by operating detectors at cryogenic temperatures (KAGRA and ET). However, if different coating materials are used in future upgrades then coating thermo-optic noise could prove to be larger than coating Brownian noise.

2.6 Finite Element Analysis

In order to accurately model the coatings in this thesis, FEA will be used. FEA is a computational method used to investigate how a component will react under a variety of conditions and allows many properties, such as stress, strain, temperature and response to applied forces, to be modelled. The mathematical algorithm determines these properties of the component by first dividing the geometry into many small individual elements known as a mesh. Specific partial differential equations are then allocated to each element in the mesh to describe the components' behaviour at that exact location. Each equation is then solved and the solutions are combined to give a solution for the entire component. The main challenge facing all finite element modellers is to ensure that the output results are trustworthy. This is achieved by ensuring that the specified analysis and model inputs are correct, and that the mesh shows "convergence".

In general, increasing the number of mesh elements in a model will lead to more accurate solutions as the mathematical resolution of the body becomes greater. However, densifying the mesh will also lead to increased computational time. A compromise between these two competing effects is therefore required. The solution is to generate a mesh which will be dense enough to give sufficiently accurate results within an acceptable time. This is achieved by first creating a mesh with few elements and solving the analysis. The number of elements in the mesh is then increased and the model re-solved. This process is continued until the results of the model are found to have converged. Performing a mesh convergence is essential in all FEA models. To obtain "convergence" it is essential to specify correct boundary conditions as well as the model's initial conditions as otherwise the FEA model will try to solve an unrealistic problem which has no converged solution.

The most common FEA software packages are called COMSOL and ANSYS.

2.7 Gravitational wave detector coating materials

The materials used to make the coatings in a gravitational wave detector must be carefully selected in order to minimise the detector's coating thermal noise and absorption. The current coating used in the Advanced LIGO and VIRGO detectors is a bilayer stack, where one layer is made from silica and the other titania doped tantala. The mechanical loss of this detector coating is 2.3×10^{-4} [112] whilst its absorption is 0.3 ppm [68]. At room temperature, this is the best coating which can be used to minimise both detector thermal noise and absorption. However, at cryogenic temperatures, the mechanical loss of this coating is high due to silica exhibiting a mechanical loss peak. New detectors which will be built to operate at cryogenic temperatures must therefore use a different coating which can ensure that the detector behaves optimally.

Due to its low mechanical loss at cryogenic temperatures, silicon nitride is currently being considered as a replacement material for silica in gravitational wave detector coatings. Amorphous silicon is also being considered as a possible partner material to silicon nitride due to its very low mechanical loss ($< 2 \times 10^{-5}$ e.g. lower than tantala) at temperatures below 30 K [113]. However, the draw back of a silicon nitride/ amorphous silicon coating is that it has high absorption (due to the amorphous silicon layers having high absorption at 1064 nm which is the laser wavelength currently used in gravitational wave detectors). One way of solving this absorption issue which is being investigated is to use a "multi-material" coatings [114, 115]. In these coating designs, a few layers of low-absorption silica and titania doped tantala are deposited on top of a amorphous silicon/silicon nitride coating stack to reflect the majority of laser power (thus reducing the power in the more highly absorbing amorphous silicon and silicon nitride layers below). This represents a trade-off between absorption and mechanical loss, as the high loss of the silica and titania doped tantala layers will increase the thermal noise of the coating.

The last type of coating that is currently being investigated for future cryogenic gravitational wave detectors are crystalline coatings. Unlike a lot of amorphous materials which show mechanical loss peaks at low temperatures, many crystalline coatings do not. Four crystalline materials which are currently being investigated are Gallium Arsenic, Aluminium Gallium Arsenic, Gallium Phosphorous and Aluminium Gallium Phosphorous. Current research into these materials centres around: reducing their absorption (which is higher than a room temperature silica/tantala layer) and improving the fabrication process to increase the size and uniformity of the crystalline coatings.

It is essential that over the next few years developments are made in the field of coating research, so that next generation gravitational wave detectors can reach design sensitivity.

2.8 Other Applications

The development of highly reflective, low mechanical loss and low optical absorption coatings is of obvious interest to the gravitational wave community but also has benefits to other fields of industry and science. One industry which would benefit from these developments in coating research is the laser building industry. The laser building industry is constantly trying to improve their coatings as it will directly improve their laser performance. Using more highly reflective, lower loss and lower absorbing coatings would enable more precise measurements to be made using a laser, increase the laser's resolution and improve the laser's stability (thermal management). These improvements in laser performance would benefit high-precision meteorology which uses lasers to detect micro-movements in the Earth's crust and are currently limited in sensitivity by the thermal noise of their laser cavities.

2.9 Conclusion

Thermal noise is one limit to the sensitivity of second-generation ground-based gravitational wave detectors at their most sensitive frequencies. The thermal noise of a detector's coated test mass is dependent upon its spatial distribution of dissipation as well as the mechanical loss of the materials used to make the coated test masses. Since the detector's laser beam is directly reflected off the coating and the coating materials exhibit significantly higher mechanical loss than the test mass materials, the thermal noise of the mirror coatings will be a significant limit to detector sensitivity. Second-generation detector's coating thermal noise is dominated by it's Brownian thermal noise component at the detector's most sensitive frequencies. In order to improve the sensitivity of future detectors, a coating's Brownian thermal noise must be reduced. Current research therefore aims to identify coating materials with low mechanical loss and low Brownian thermal noise.

3 Bulk and shear loss

3.1 Introduction

The intrinsic mechanical loss of a material is thought to arise from the rearrangement of defects or structural units of atoms in response to an applied stress [97, 100]. These loss mechanisms can be represented by transitions between states in a double well potential (as shown in Figure 3-1) with stable configurations represented by a potential minimum (labelled a and b in Figure 3-1) [116].

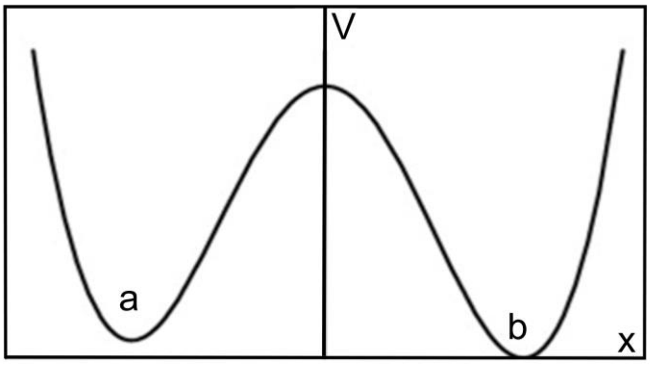

Figure 3-1: Image of a double well potential. The Y axis represents the potential (V) whilst the x axis represents position (x). a and b represent the stable configurations of the double well potential.

In principle, different types of motion (i.e. different types of stress) couple to the loss mechanism (or double well potentials) in different ways, resulting in a different level of loss [107]. To accurately estimate the thermal noise of a detector's coating, it is therefore necessary to correctly account for the type of motion which is used to calculate the mechanical loss of the coating materials.

Mechanical loss is usually measured using a 'ring-down' technique, in which the loss of a coating is measured at a particular vibrational mode of a coated sample (Section 3.3). Many different samples and mode shapes can be used – depending on the materials and applications involved [117, 118, 119]. Since different types of motion couple to the loss mechanism in different ways, it is not always accurate to compare the loss of a coating measured at one resonant mode to another coating measured at a different resonant mode. Therefore care needs to be taken when comparing mechanical loss values measured using different sample geometries, and even different resonant modes for the same type of sample.

A solution to this problem was provided by Hong et al [107], who showed that the loss of a coating measured at a particular mode (ϕ_{coating}) can be deconstructed into two loss factors – one associated with bulk motion and the other shear motion. These bulk (ϕ_{B}) and shear (ϕ_{S}) loss angles are fundamental to the coating and can therefore be compared across different types of samples and resonant modes. By deconstructing the resonant mode dependent loss of a coating into its more fundamental and invariant bulk and shear loss quantities, Hong shows that it is possible to identify the coating with the lowest bulk and shear losses even if the coatings have been measured at different resonant mode shapes.

This chapter begins by introducing the theory of bulk and shear loss. It then outlines two new methods to determine the bulk and shear losses of a coating before applying these methods to analyse a range of coatings that have been deposited onto different substrate geometries which oscillate with different mode shapes. The main contribution of this work is the development of these two new methods of calculation.

3.2 Theory of bulk and shear loss

To mathematically describe any isotropic amorphous coating, two independent elastic constants are required (e.g. Young's modulus and Poisson ratio or Young's modulus and shear modulus etc) [107]. Since each elastic constant contains an imaginary component which describes an individual loss mechanism, every isotropic amorphous coating therefore must have two independent loss mechanisms. Hong describes a coating using the bulk and shear elastic constants [107] and argues that the mode dependent mechanical loss (ϕ_{coating}) is therefore made from a combination of both bulk loss (ϕ_{B}) and shear loss (ϕ_{S}). Whilst the loss of a coating determined at different mode shapes will vary (since each mode shape samples different amounts of bulk and shear motion), the fundamental and independent bulk and shear losses are invariant. Mathematically, ϕ_{coating} can be written in the form [107]:

$$\phi_{\mathrm{coating}} = \frac{U_{\mathrm{Bulk}}}{U_{\mathrm{c}}}\phi_{\mathrm{B}} + \frac{U_{\mathrm{Shear}}}{U_{\mathrm{c}}}\phi_{\mathrm{S}}\,, \tag{3.1}$$

where the energy stored in bulk (U_{Bulk}) and shear motion (U_{Shear}) are defined as being:

$$U_{\mathrm{Bulk}} = \int \frac{1}{2} K\theta^2 dV\,, \tag{3.2}$$

$$U_{\text{Shear}} = \int \mu \xi_{ij} \xi_{ij} dV \,, \tag{3.3}$$

where:

$$\theta = S_{ii} \,, \tag{3.4}$$

$$\xi = \frac{1}{2}\left(S_{ij} + S_{ji}\right) - \frac{1}{3} g_{ij} S_{kk} \,, \tag{3.5}$$

K is the bulk modulus, μ is the shear modulus, g_{ij} represents a Kronecker delta function, U_c is the energy stored in the coating, V is the volume of the coating and S_{ii} represents the 3×3 strain tensor [107]. The bulk loss arises from motion which involves volume change, whilst the shear loss arises from shear motion where there is no change in volume but a change in shape [120]. Examples of both bulk and shear motion are illustrated in Figure 3-2 (a) and (b) respectively.

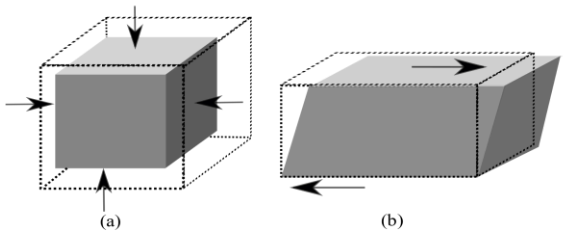

Figure 3-2: Two bodies undergoing bulk (a) and shear (b) motion. The dashed volumes illustrate the original shape of the body whilst the opaque structures show their end state. The arrows illustrate the direction of force applied to the original body to transform it into its end state.

Using Levin's approach [102] (described in Section 2.3) Hong derived the thermal noise power spectral density (S_ϵ) of a single coating layer (used as a mirror coating in a gravitational detector) as being [107]:

$$S_\epsilon = \frac{4k_b T}{\pi f}\left(\frac{\phi_B U_{\text{Bulk}}}{F_0^2} + \frac{\phi_s U_{\text{Shear}}}{F_0^2}\right), \tag{3.6}$$

where:

$$\frac{U_{\text{Bulk}}}{F_0^2} = \frac{(1-2\nu_c)d}{3\pi w_0^2}\left(\begin{array}{c} \frac{Y_c}{Y_s^2}\frac{(1-2\nu_s)^2(1+\nu_s)^2}{(1-\nu_c)^2} + \\ \frac{1}{Y_s}\frac{2(1-2\nu_s)(1+\nu_s)(1+\nu_c)}{(1-\nu_c)^2} + \frac{1}{Y_c}\frac{(1+\nu_c)^2}{(1-\nu_c)^2} \end{array}\right), \tag{3.7}$$

$$\frac{U_{\text{Shear}}}{F_0^2} = \frac{2d}{3\pi w_0^2}\left(\begin{array}{c}\dfrac{Y_c}{Y_s^2}\dfrac{(1-\nu_c+\nu_c^2)\quad(1+\nu_s)^2(1-2\nu_s)^2}{(1-\nu_c)^2(1+\nu_c)} \\ -\dfrac{(1-2\nu_c)(1+\nu_s)(1+\nu_c)(1-2\nu_s)}{Y_s(1-\nu_c)^2}+\dfrac{(1-2\nu_c)^2}{Y_c}\dfrac{(1+\nu_c)^2}{(1-\nu_c)^2}\end{array}\right), \quad (3.8)$$

w_0 is the field amplitude radius of the detectors laser beam applied to the coating surface (i.e. at radius w_0 from the centre of the beam, the light intensity is $\frac{1}{e^2}$), ν is the Poisson ratio, d is the coating thickness, k_b is the Boltzmann constant, f is the frequency, T is the temperature, Y is the Young's modulus and subscript s and c are used to define properties of the substrate and coating respectively. Applying the simplification ($Y_c = Y_s$ and $\nu_c = \nu_s$), a greater insight as to the effects of bulk and shear loss upon thermal noise can be observed:

$$S_\epsilon = \frac{8k_bT(1-\nu-2\nu^2)d}{3\pi^2 fYw_0^2}\left(2(1+\nu)\phi_B + (1-2\nu)\phi_S\right). \quad (3.9)$$

Thermal noise of a single coating layer can therefore be reduced by decreasing its temperature, the thickness of coating or the relative magnitudes of the coating's bulk and shear mechanical loss factors. It can also be decreased by increasing the laser diameter. It is interesting to note that the thermal noise of a coating is more sensitive to bulk loss than shear loss. Physically, this occurs because bulk motion within the coating layer causes a thickness fluctuation of the layer and a fluctuation of the coating-substrate interface which add constructively together to increase thermal noise (by changing the detector arm length) [107]. Shear motion has less affect upon thermal noise as when the coating moves in this way the change upon the detector arm length is smaller than would be for bulk motion [107].

The effect of a coating's bulk and shear loss upon thermal noise can be observed in Figure 3-3. The three lines represent the same measured mode-dependent coating loss (Equation (3.1)) but different ratios of bulk and shear loss. It can be seen that as the bulk/shear loss ratio decreases, the thermal noise decreases.

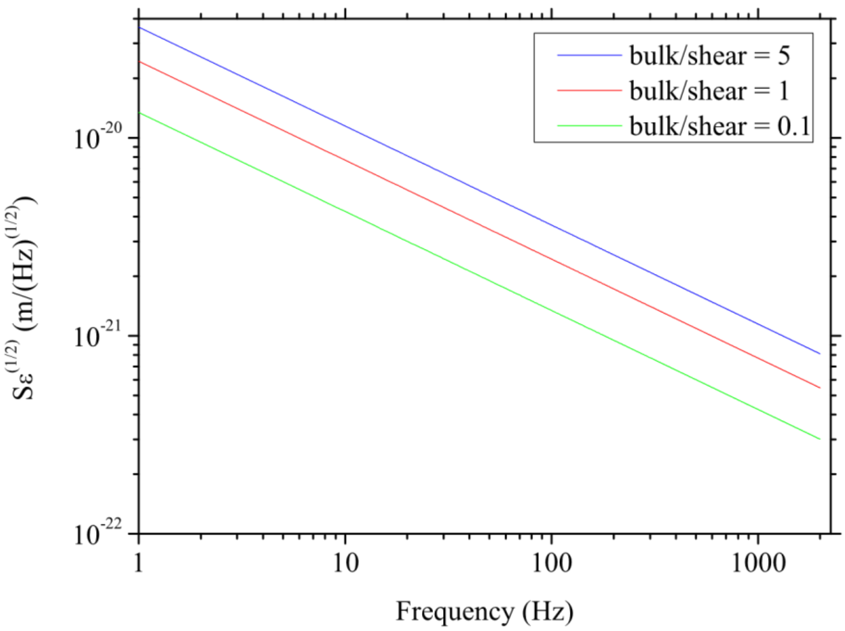

Figure 3-3: A single layer coating's thermal noise calculated using Equations (3.6) – (3.8). Each line was calculated using the same mode dependent coating loss but with different proportions of loss stored in bulk and shear motion. As convention dictates the square root of the thermal noise power spectral density is plotted.

The current method used to determine coating thermal noise in a gravitational wave detector (Equation (2.26) when $\phi_{\text{substrate}} = 0$ and assuming $\phi_{\parallel} = \phi_{\perp} = \phi_{\text{coating}}$) has a linear relationship between thermal noise and the coating's mode dependent mechanical loss. This means coatings which are measured to have the lowest mode dependent mechanical loss also have the lowest thermal noise. However, as shown in Figure 3-3, this is not the case when using Hong's theory as a coating which has been measured to have a higher mode dependent loss could theoretically have a lower thermal noise than another coating with a lower mode dependent coating loss due to its bulk and shear split. Based upon this new thermal noise theory, coatings that were dismissed in the past as being too noisy (as they had a higher mode dependent mechanical loss than others) should be reinvestigated as they may in fact improve current detector performance.

It should be emphasised that Figure 3-3 shows the thermal noise of a single layer coating as a simplified illustration and not the multilayer stack used in a gravitational wave detector (see Sections 2.3 and 6.7.3 for more details). A fuller analysis is required to account for the shear and bulk losses of the two materials used to make a highly reflective coating stack [107]. Here it is assumed for simplicity that both coating layers have identical thermal properties, mechanical properties and loss values, allowing the coating to be treated as a single layer of material.

3.3 Method to determine the bulk and shear loss of a coating

To determine the bulk and shear losses of a coating, the mode dependent losses of the coating must first be calculated. In this section, techniques for measuring the mode dependent mechanical loss of a coating will be described, and the analysis which was developed to extract the bulk and shear losses will be explained.

3.3.1 Mechanical loss experiments

A material's mechanical loss is usually measured by exciting a resonant mode of a sample (made from that material) and measuring the exponential decay of its resulting motion. When all external sources of energy loss are suitably minimised, this decay is proportional to the internal mechanical loss of the material. To determine the mechanical loss of a potential gravitational wave detector coating, the coating is first deposited onto a test-sample (substrate) of known loss. By measuring the mechanical loss of the coated sample and determining its difference in loss from the uncoated sample, the loss of the coating can be calculated.

3.3.1.1 Cantilever substrate geometries

Thin silicon [121, 122] or silica cantilevers [123] are commonly used as substrates in mechanical loss experiments. Approximate geometries of these substrates are illustrated in Figure 3-4.

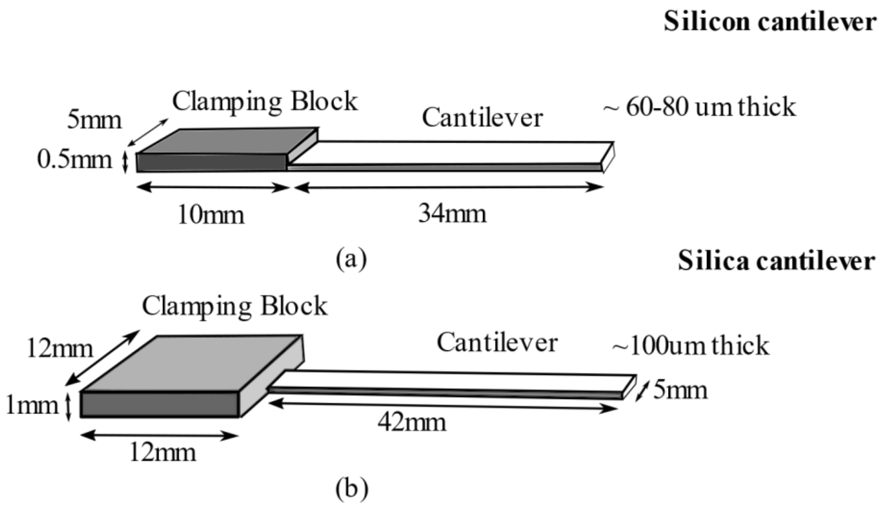

Figure 3-4: Typical geometries of a (a) silicon and (b) silica cantilever used for loss measurements.

To most accurately calculate a coating's bulk and shear loss (using the methods derived in this chapter), the mode dependent loss of the coating should ideally be determined at resonant mode shapes with different bulk/shear energy ratios. Bending and torsional modes are the most common resonant mode shapes of a cantilever. They also exhibit significantly different bulk and shear energy ratios from one another. For this reason, loss measurements were made at both the uncoated and coated cantilever's bending and torsional modes. Figure 3-5 illustrates the shape of the first (lowest frequency) bending mode and first (lowest frequency) torsional mode for a clamped silicon cantilever.

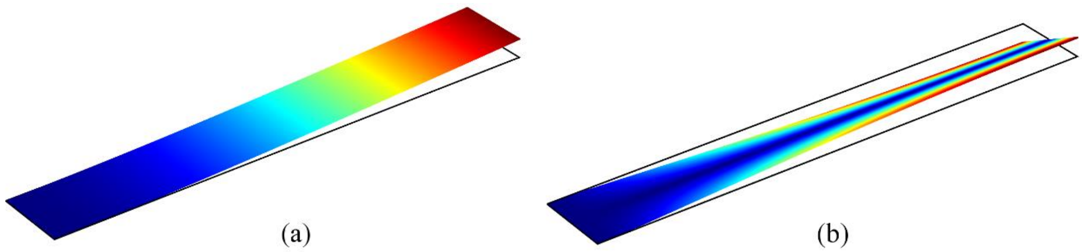

Figure 3-5: Resonant mode shapes of a clamped silicon cantilever. (a) First bending mode. (b) First torsional mode. The colour blue represents the parts of the cantilever with minimum deflection whilst red illustrates maximum deflection. The black outline illustrates the undeformed cantilever.

3.3.1.2 Measuring the mechanical loss of an uncoated and coated cantilever

A schematic of the apparatus used to measure the mechanical loss of cantilever samples is shown in Figure 3-6. The thick end of the cantilever (shown in Figure 3-4) was held horizontally in a stainless-steel clamp inside a vacuum tank. The clamping arrangement using a thicker part of the sample has been shown to be effective in reducing frictional losses associated with the clamp [124].

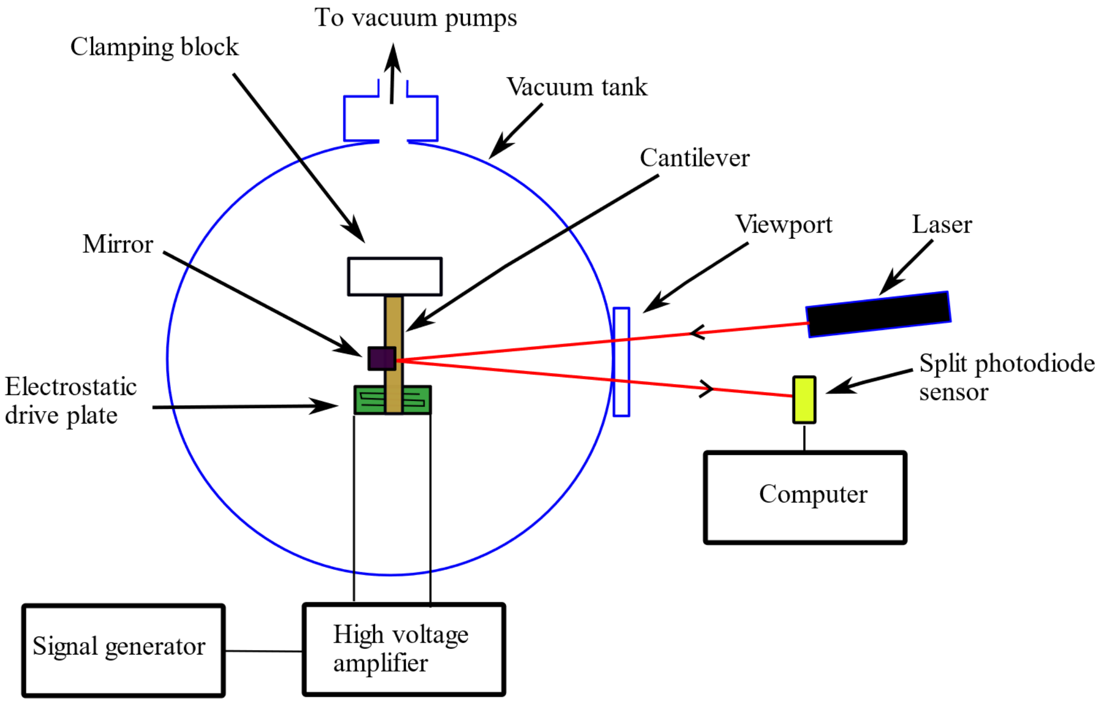

Figure 3-6: Experimental setup used to measure the loss of an uncoated or coated cantilever.

The vacuum tank was evacuated to approximately 2×10^{-6} mbar and an electrostatic drive plate (positioned approximately 5 mm below the cantilever) was used to excite the cantilever by applying an oscillating high-voltage signal at the frequency of the desired resonant mode. After excitation, the drive plate is turned off and the free amplitude decay $a(t)$ of the resonant mode is recorded using laser light reflected from the cantilever and directed onto a split photodiode sensor outside of the vacuum tank as shown in Figures 3-6 and 3-7 (for more information regarding the electric circuit and DAQ used to measure the current in the photodiode, please see [123]). To maximise the sensitivity of the experiment, the laser light was aligned to a position on the cantilever which experiences maximum displacement (the end of the cantilever, in the case of a bending mode, and close to one edge of the cantilever, in the case of a torsional mode).

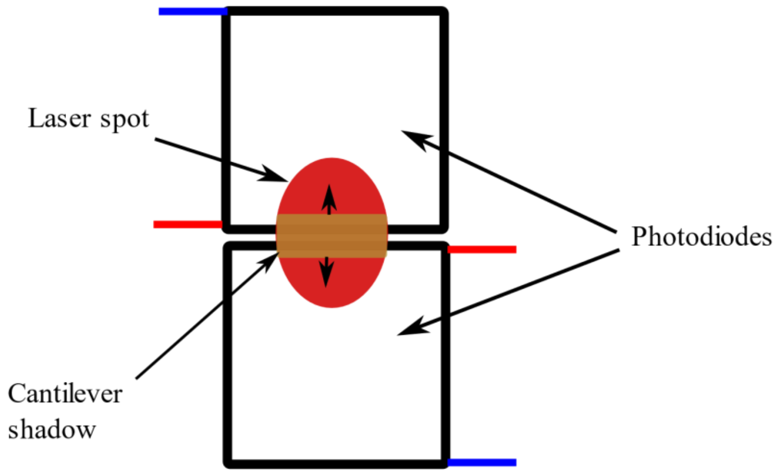

Figure 3-7: Split photodiode arrangement used to measure the amplitude of oscillation of a cantilever.

The amplitude decay, or 'ring down', obtained from monitoring the laser light is collected using a Labview program. The measured amplitude decay is related to the mechanical loss (assuming all external sources of damping have been sufficiently reduced) by:

$$a(t) = a_0 \exp(-\pi f_0 \phi t) , \tag{3.10}$$

where a_0 is the initial amplitude, f_0 is the resonant frequency of the cantilever and t is the time. The mechanical loss (ϕ) of the uncoated ($\phi_{\text{uncoatedloss}}$) or coated ($\phi_{\text{coatedloss}}$) cantilever is calculated by fitting Equation (3.10) to the measured amplitude decay. An example of a typical ring down and its corresponding fit can be observed in Figure 3-8.

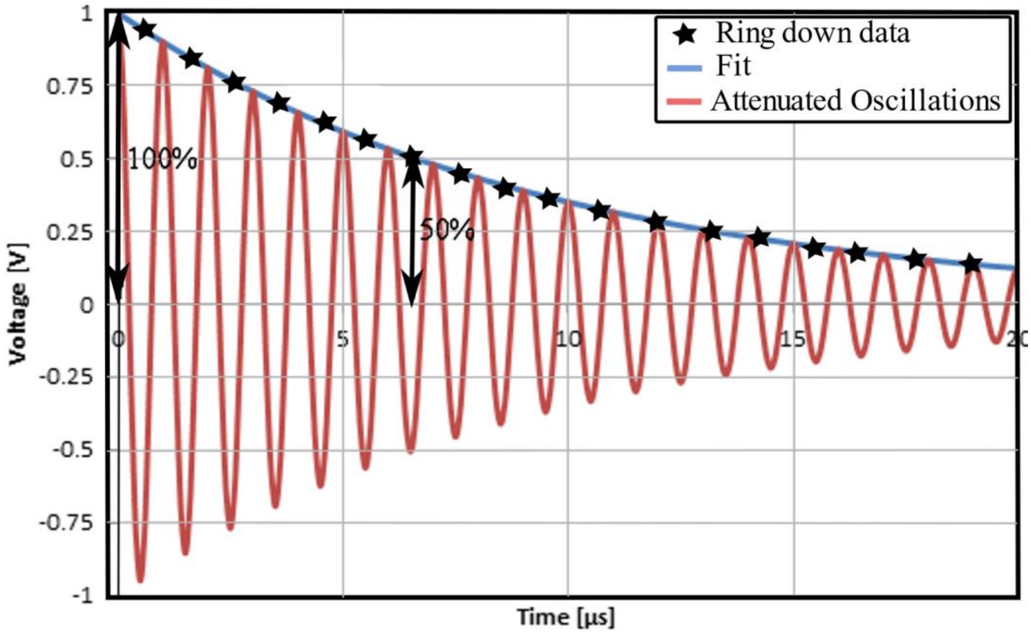

Figure 3-8: An example of a ring down measurement and its corresponding fit. The attenuated oscillations of the cantilever are also included in the image.

The frequency (f) of the cantilever's n^{th} excited bending mode is equal to [125]:

$$f_n = \frac{(k_n L)^2 a}{4\pi\sqrt{3}L^2}\left(\frac{Y}{\rho}\right)^{\frac{1}{2}}, \tag{3.11}$$

where a is the thickness of the cantilever, L is its length, Y is its Young's modulus, ρ its density and $k_n L$ takes on the values 1.875, 4.694, 7.853, 10.996, 14.137 for $n = 1$ to 5 and $k_n L = \frac{(2n-1)\pi}{2}$ for $n > 5$ [126]. Similarly, the frequency of the cantilever's n^{th} torsional mode can be calculated using [127]:

$$f_n = \frac{D_n}{L}\sqrt{\frac{4Ga^2}{b^2\rho}}, \tag{3.12}$$

where G is the cantilever's shear modulus, b is its width and $D_n = \frac{(2n-1)\pi}{2}$ for $n \geq 1$.

3.3.2 Mode dependent coating loss

This subsection describes the analysis required to calculate the mode-dependent loss of a coating for both a bending mode and a torsional mode of a cantilever substrate.

3.3.2.1 Bending mode coating loss

The mechanical loss of a coated cantilever at a bending mode ($\phi_{\mathrm{coatedloss,b}}$) has contributions from both the loss of the cantilever substrate ($\phi_{\mathrm{uncoated.b}}$) and the coating material ($\phi_{\mathrm{coating,b}}$). Since the coating is much thinner than the substrate, less elastic energy from the motion is stored in the coating than in the substrate and therefore only a small fraction of the energy associated with the resonant mode is able to be dissipated by the coating. Accounting for these differences in energies stored in the coating and substrate, the mechanical loss of the coated cantilever can be approximated by [128]:

$$\phi_{\mathrm{coatedloss,b}} \approx \phi_{\mathrm{uncoated,b}} + \frac{U_{\mathrm{c,b}}}{U_{\mathrm{s,b}}} \phi_{\mathrm{coating,b}} , \tag{3.13}$$

where $U_{\mathrm{c,b}}$ and $U_{\mathrm{s,b}}$ represent the elastic strain energies of the coating and substrate during a bending mode vibration. The energies stored in a coating and it's cantilever substrate can be calculated by considering a bar of length L, thickness a and width b which has a thin coating of thickness t on one surface. If the coated bar is bent into an arc of a circle with radius R (as shown in Figure 3-9), the energy stored in the coating layer can be determined as being:

$$U_{\mathrm{c,b}} = \frac{Y_{\mathrm{c}} t b}{2L} \Delta L^2 , \tag{3.14}$$

where ΔL is the change in length due to bending.

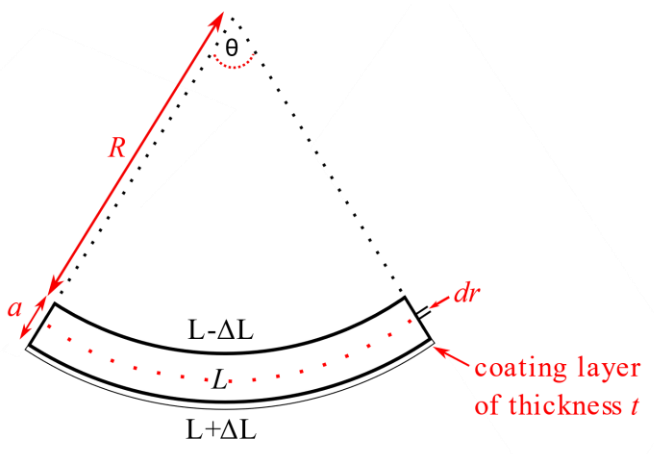

Figure 3-9: A bar of length L, thickness a and width b which has a thin coating of thickness t on one surface and is bent into an arc of a circle with radius R.

θ is the angle of bending and it can be assumed that when the angle of bending is small, the approximation:

$$\Delta L \approx \frac{a}{2}\theta\,, \tag{3.15}$$

holds. The energy stored in the bar can be calculated by splitting the bar into two parts, one which is experiencing compression and the other expansion. Both parts contribute equally to the total energy stored in the bar. The energy stored in the half of the beam undergoing compression ($U_{\mathrm{s,b,compression}}$) is equal to:

$$U_{\mathrm{s,b,compression}} = \int_{R}^{R+\frac{a}{2}} \frac{Y_s b\theta^2 (r-R)^2}{2L}\,dr = \frac{1}{48}\frac{Y_s b\theta^2 a^3}{L}\,. \tag{3.16}$$

Therefore the total energy stored in the bar is equal to:

$$U_{\mathrm{s,b}} = 2\times\frac{1}{48}\frac{Y_s b\theta^2 a^3}{L} = \frac{1}{24}\frac{Y_s b\theta^2 a^3}{L}\,, \tag{3.17}$$

and the ratio of energies stored in the substrate and coating for a bending mode can be approximated by:

$$\frac{U_{\mathrm{s,b}}}{U_{\mathrm{c,b}}} \approx \frac{Y_s a}{3Y_c t}\,, \tag{3.18}$$

when θ is small (Equation 3.15). Combining Equations (3.13) and (3.18) the mechanical loss of the coating at a bending mode can be determined using:

$$\phi_{\mathrm{coating,b}} \approx \frac{Y_s a}{3Y_c t}\left(\phi_{\mathrm{coated,b}} - \phi_{\mathrm{uncoated,b}}\right). \tag{3.19}$$

3.3.2.2 Torsional mode coating loss

The ratio of energy stored in the substrate ($U_{s,t}$) relative to the coating ($U_{c,t}$) for a torsional mode can be approximated by [129]:

$$\frac{U_{s,t}}{U_{c,t}} \approx \frac{G_s a}{3G_c t}\,, \tag{3.20}$$

where G_s and G_c are the shear moduli of the substrate and coating respectively. The mechanical loss of a coating at a torsional mode can therefore be calculated using:

$$\phi_{coating,t} \approx \frac{G_s a}{3G_c t}\left(\phi_{coated,t} - \phi_{uncoated,t}\right). \tag{3.21}$$

3.3.3 Verifying the accuracy of the bending and torsional mode energy ratio approximations using Finite Element Analysis

Equations (3.18) and (3.20) are used in the literature to approximate the energy stored in the coating relative to the substrate for bending and torsional modes [122, 129]. This section will look at verifying the accuracy in applying these equations to the coating and cantilever geometries used in loss measurements as well as identify their limitations by comparing their values to those obtained using finite element analysis (FEA).

FEA is a computational method used to investigate how a component will react under a variety of conditions and allows many properties, such as stress, strain, temperature and response to applied forces, to be modelled. The mathematical algorithm determines these properties of the component by first dividing the geometry into many small individual elements known as a mesh. Specific partial differential equations are then allocated to each element in the mesh to describe the components' behaviour at that exact location. Each equation is then solved and the solutions are combined to give a solution for the entire component. The main challenge facing all finite element modellers is to ensure that the output results are trustworthy. This is achieved by ensuring that the specified analysis and model inputs are correct, and that the mesh shows "convergence".

In general, increasing the number of mesh elements in a model will lead to more accurate solutions as the mathematical resolution of the body becomes greater. However, densifying the mesh will also lead to increased computational time. A compromise between these two competing effects is therefore required. The solution is to generate a mesh which will be dense enough to give sufficiently accurate results within an acceptable time. This is achieved by first creating a mesh with few elements and solving the analysis. The number of elements in the mesh is then increased and the model re-solved. This process is continued until the results of the model are found to have converged. Performing a mesh convergence is essential in all FEA models. To obtain "convergence" it is essential to specify correct boundary conditions as well as the model's initial conditions as otherwise COMSOL will try to solve an unrealistic problem which has no converged solution.

To verify the accuracy of the energy ratio equations, FEA was used to calculate the elastic strain energies in coated cantilevers of varying geometries and material properties. The FEA modelling package ANSYS was used and a *modal analysis* [130] was carried out to identify the resonant modes of the cantilevers and to determine the coating and substrate elastic strain energy for each mode. To ensure the models had converged, the mesh density was increased until the strain energy ratios were observed to stop changing with further increases in mesh density.

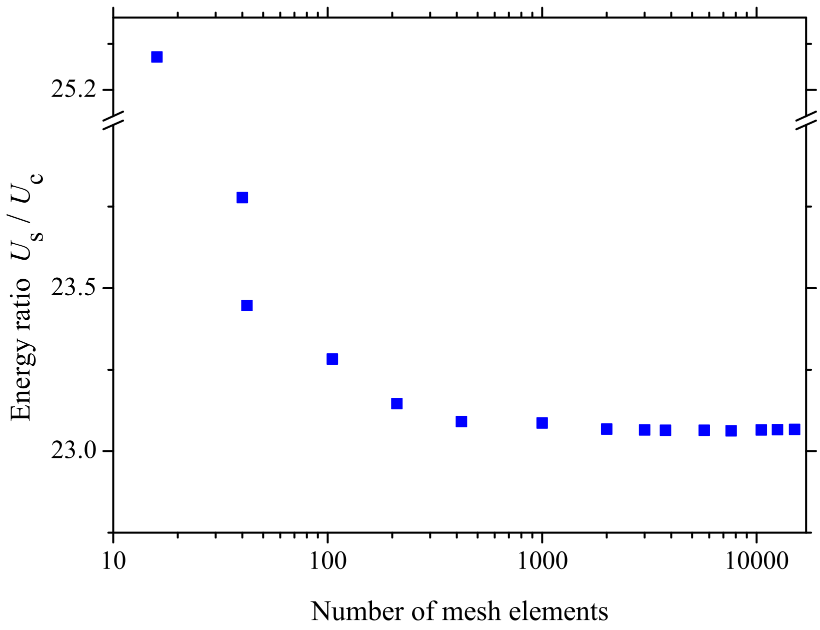

(a)

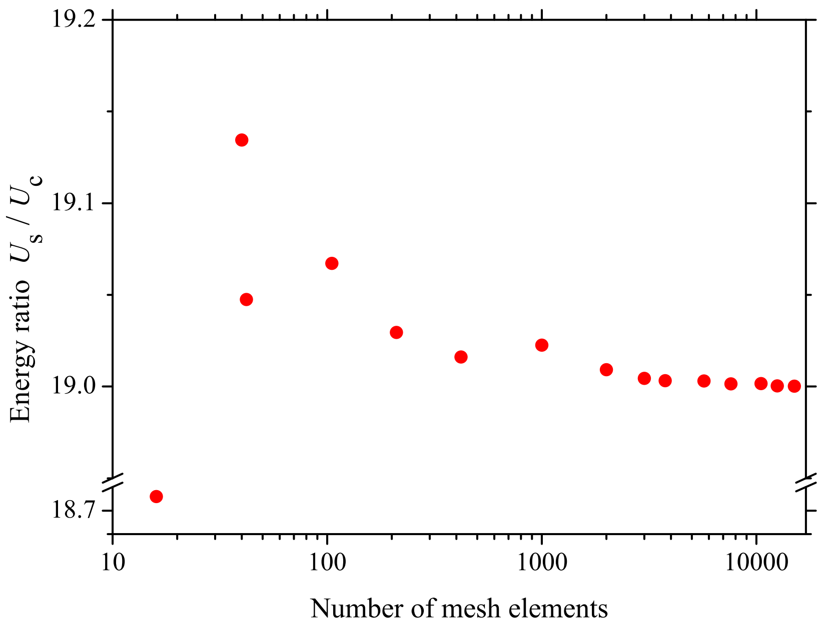

(b)

Figure 3-10: Mesh convergence for the first bending mode (a) and first torsional mode (b) of a coated silicon cantilever. The blue squares represent the energy ratio for the coated silicon cantilever's first bending mode when a particular number of mesh elements was used in the FEA model. A red circle similarly represents the silicon cantilever's first torsional mode energy ratio. The cantilever modelled was 34×10^{-3} *m in length,* 5×10^{-3} *m in width and* 63.4×10^{-6} *m in thickness. The coating was* 1×10^{-6} *m thick amorphous silicon.*

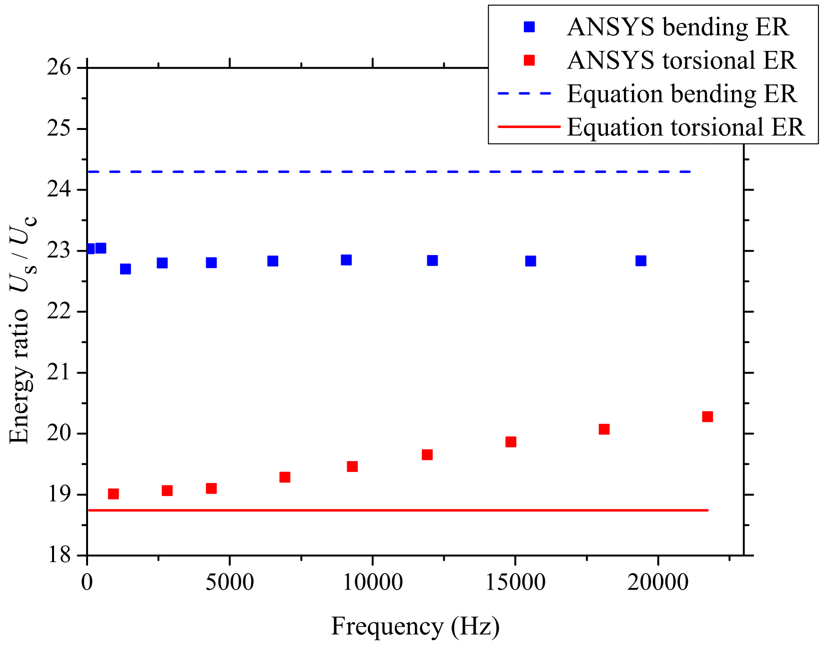

Figure 3-11: Energy ratios (ER) of a crystalline silicon cantilever that has been coated with a 1×10^{-6} *m thick amorphous silicon layer. The energy ratios were calculated using ANSYS as well as Equations (3.18) and (3.20). The ANSYS model used anisotropic silicon properties. The cantilever was* 34×10^{-3} *m in length,* 5×10^{-3} *m in width and* 63.4×10^{-6} *m in thickness.*

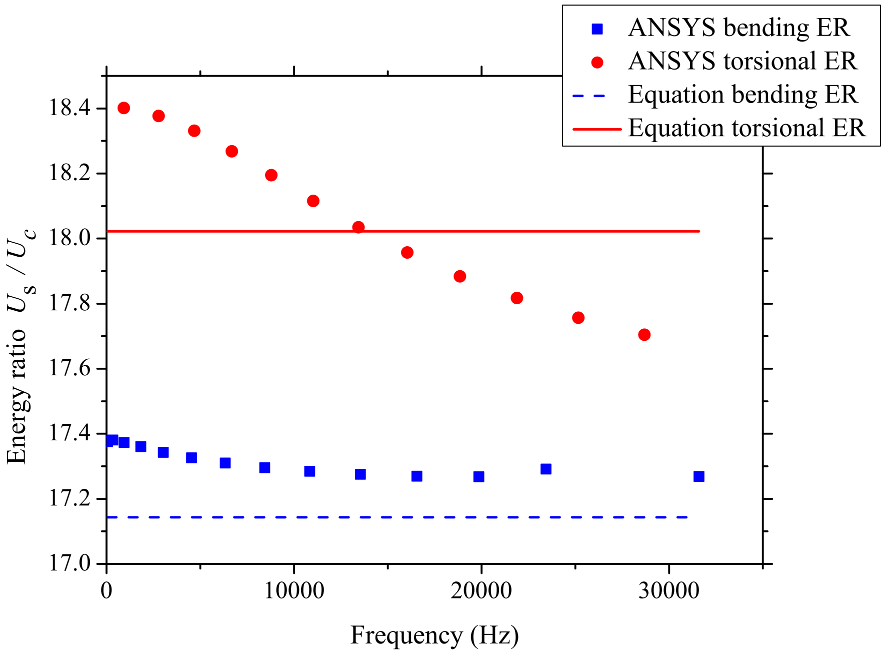

Figure 3-12: Energy ratios of a silica cantilever that has been coated with a 1×10^{-6} *m thick tantala layer. The energy ratios were calculated using ANSYS as well as Equations (3.18) and (3.20). The cantilever was* 42×10^{-3} *m in length,* 5×10^{-3} *m in width and* 100×10^{-6} *m in thickness.*

Figure 3-10 shows the mesh convergence of the energy ratios of a coated silicon cantilever for the first bending and torsional mode. Assuming that meshes which determine energy ratios with a deviation of less than 1 % can be described as converged, then it is observed in Figure 3-10, that approximately 100 elements are required for mesh convergence. Similar results were obtained for the cantilever's next 18 modes. Figure 3-11 compares the modelled energy ratio values determined by FEA to those predicted by the analytical equations for a coated silicon cantilever. A similar comparison for a coated silica cantilever is shown in Figure 3-12. In both cases a converged mesh of 12500 elements was used. It can be seen that the energy ratios calculated from Equations (3.18) and (3.20) give a reasonably good approximation to the FEA values over the frequency range studied. However, all of the FEA energy ratios display some frequency dependence, which is not predicted by the equations. For both types of cantilever studied, the FEA bending mode energy ratios appear to converge as the frequency (i.e. mode order) is increased. For the silica cantilever, there is approximately a 2% off-set between the converged value and the predicted energy ratio from the equation, while this off-set is -5% for the silicon cantilever.

The FEA results for the torsional modes show a greater frequency dependence than for the bending modes. For the silicon cantilever, the FEA ratios are close to the analytical values at low frequency, but appear to diverge from the analytical value at higher frequency. For the silica cantilever, the FEA ratios start out higher than the analytical value, but the ratio decreases at higher frequencies and drops below the analytical prediction. For both cantilevers it can be observed that the FEA torsional mode energy ratio tends to the bending mode energy ratio at high frequencies. It is hypothesised that this occurs as a result of the shape of higher order torsional modes approximating that of higher order bending modes.

It can be concluded, that although the bending and torsional energy ratio equations show similar accuracy in replicating the FEA values below 20 kHz, the equation for torsional modes starts to become significantly less accurate at frequencies above 20 kHz. Additional FEA models with different cantilever geometries were also examined and confirm this conclusion. The results from these models are presented in Figures 3-13 and 3-14.

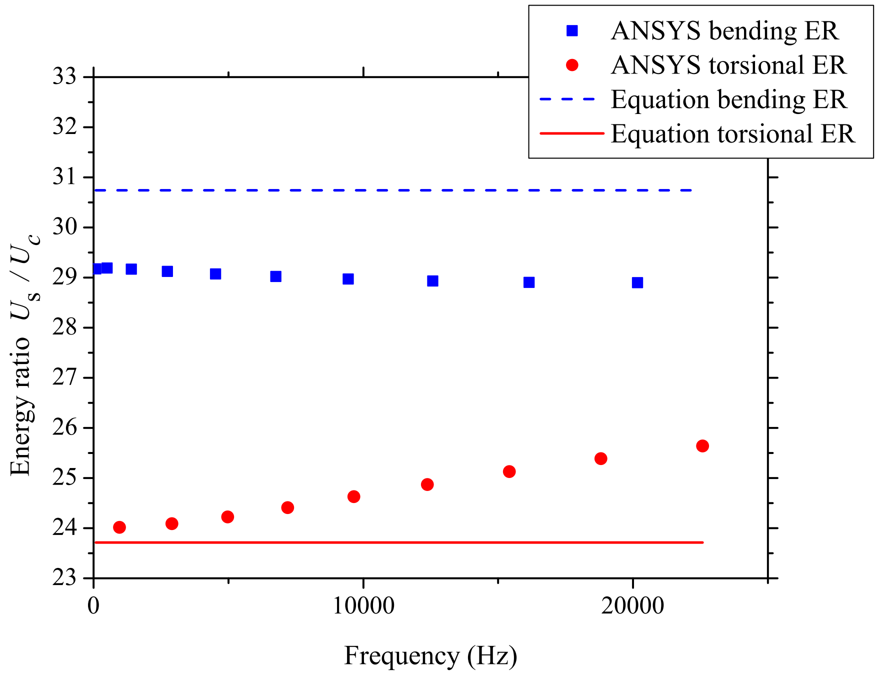

Figure 3-13: Energy ratios of a silicon cantilever that has been coated with a 8.24×10^{-7} *m thick amorphous silicon layer. The cantilever was* 34×10^{-3} *m in length,* 5×10^{-3} *m in width and* 66.1×10^{-6} *m in thickness.*

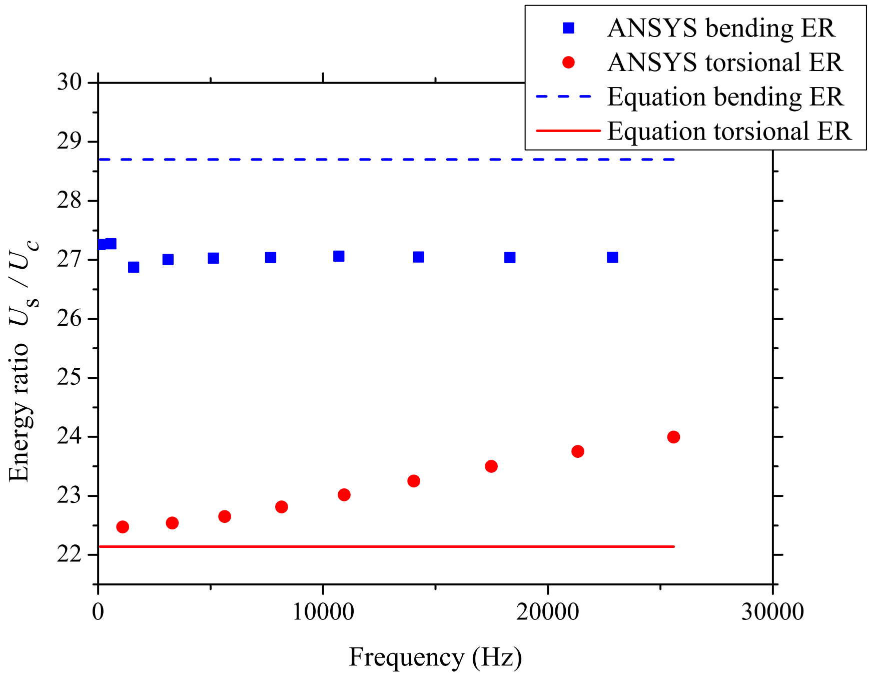

Figure 3-14: Energy ratios of a silicon cantilever that has been coated with a 1×10^{-6} *m thick amorphous silicon layer. The cantilever was* 34×10^{-3} *m in length,* 5×10^{-3} *m in width and* 74.9×10^{-6} *m in thickness.*

In the literature mode dependent coating losses are most commonly estimated using the approximate energy ratio equations (Equations (3.18) and (3.20)). However, in this thesis all mode dependent coating losses will be estimated using FEA. Whilst FEA analysis is more time consuming, it is more rigorous, particularly since it does not rely upon the small angle approximations required to derive Equations (3.18) and (3.20). Furthermore, FEA modelling treats each mode independently which enables it to identify frequency dependent relationships within the energy ratios.

It should also be noted that the results presented in this section were confirmed using a second FEA package called COMSOL (licence number 7075984).

3.3.4 Procedures to determine the bulk and shear losses of a coating

Two methods were developed to determine the bulk and shear losses of a coating. The first method – using simultaneous equations – was quicker to implement but gave cruder results. The second method – using the minimisation approach – was more time consuming but was found to give more accurate results.

3.3.4.1 Simultaneous equation method

This first method was developed to allow a relatively quick calculation of the bulk and shear losses of a coating. The more bending and torsional modes which have been measured, the better this method will perform. For each bending mode for which uncoated and coated loss measurements are available, an equation of the following form is constructed:

$$\phi_{\mathrm{coating,bx}} = \frac{U_{\mathrm{bulk,bx}}}{U_{\mathrm{c,bx}}}\phi_{\mathrm{bulk}} + \frac{U_{\mathrm{shear,bx}}}{U_{\mathrm{c,bx}}}\phi_{\mathrm{shear}}\,, \tag{3.22}$$

whilst for each torsional mode, an equation of the following form:

$$\phi_{\mathrm{coating,tx}} = \frac{U_{\mathrm{bulk,tx}}}{U_{\mathrm{c,tx}}}\phi_{\mathrm{bulk}} + \frac{U_{\mathrm{shear,tx}}}{U_{\mathrm{c,tx}}}\phi_{\mathrm{shear}}\,, \tag{3.23}$$

is built where bx and tx are used to represent the different bending and torsional modes e.g. $U_{\mathrm{bulk,b1}}$ and $U_{\mathrm{bulk,t1}}$ represent the coating elastic strain energies stored within bulk motion for the first bending and torsional modes. To determine the unknown shear and bulk losses in Equations (3.22) and (3.23), the bulk energy, shear energy, coating energy, substrate energy and mode dependent coating losses for all of the resonant modes (whose uncoated and coated losses have both been measured) must be determined. The various elastic energies were obtained using a *solid mechanics eigenfrequency analysis* within the FEA package, called COMSOL [131]. For each mode shape of the clamped cantilever, the bulk and shear energies in the coating were calculated using user defined volume integrals (Equations (3.2) and (3.3) respectively). The total energy stored in the coating and cantilever were calculated using COMSOL's predefined elastic energy equation. The energies were calculated repeatedly for a range of increasing mesh densities until the energies converged (i.e. they were observed to change by less than 1 %).

These coating and substrate energies were used in conjunction with the measured coated and uncoated losses to determine the coatings mode dependent losses. COMSOL's bulk, shear and coating energies were then input into Equations (3.22) and (3.23) along with the mode dependent coating losses and a series of simultaneous equations were constructed by pairing every possible combination of bending and torsional mode equation together e.g:

$$\phi_{\mathrm{coating,b4}} = \frac{U_{\mathrm{bulk,b4}}}{U_{c,b4}}\phi_{bulk} + \frac{U_{\mathrm{shear,b4}}}{U_{\mathrm{c,b4}}}\phi_{\mathrm{shear}}\,, \tag{3.24}$$

and

$$\phi_{\mathrm{coating,t2}} = \frac{U_{\mathrm{bulk,t2}}}{U_{\mathrm{c,t2}}}\phi_{\mathrm{bulk}} + \frac{U_{\mathrm{shear,t2}}}{U_{\mathrm{c,t2}}}\phi_{\mathrm{shear}}\,, \tag{3.25}$$

where b4 and t2 are used to represent the fourth bending and second torsional mode respectively. Each pair of simultaneous equations was then solved individually for the bulk and shear loss. For N bending modes and M torsional modes, $N \times M$ simultaneous equations were formed and $N \times M$ values of bulk and shear loss were calculated.

To account for the uncertainties in all of the inputs to this calculation, this process of determining the bulk and shear losses was repeated a further 99 times. Each time the bulk and shear losses were re-calculated, every input parameter (coating's Youngs modulus, coated loss, uncoated loss etc) was allowed to randomly vary within their associated error bounds. This allowed the uncertainties associated with all of the material properties, and the uncertainties associated with the mechanical loss measurements, to be taken into account. The result was a set of $100 \times N \times M$ bulk and shear loss values. These were then averaged to obtain a best estimate for the loss values, with the error in the bulk and shear losses taken to be the standard deviation of the set of losses.

3.3.4.2 Minimisation method

In the minimisation method, random values of bulk and shear loss are used to calculate a theoretical coating loss for each mode using Equations (3.22) and (3.23). The bulk and shear loss values are randomly varied until a minimum difference between these calculated mode-dependent coating losses and the experimental losses is obtained for all modes. Similarly to the simultaneous equation method, the energies used in Equations (3.22) and (3.23) are calculated using COMSOL.

Similarly to the simultaneous equation method, the minimisation method was also repeated a further 99 times using randomly varied input parameters within their associated error bounds. The final result was 100 bulk and shear losses being calculated. The best estimates and errors of the bulk and shear losses were determined by taking an average and standard deviation of the set of losses.

This minimisation method is believed to be more accurate (i.e. to have a smaller error associated with its method of estimation) than the simultaneous equation approach as the simultaneous equation approach only uses two data points (one bending and one torsional mode) to estimate a bulk and shear loss at any one time whereas the minimisation method uses all data points to estimate a value of bulk and shear loss. Using only two points of information is more likely to lead to greater outliers in the calculated bulk and shear loss whereas using all data points will act to constrain the magnitude of these extreme values. When a standard deviation of the bulk and shear losses is calculated to determine their error, the set of losses with the most outliers (simultaneous equation approach) will have a larger standard deviation (and therefore error).

3.4 Determining the bulk and shear losses of a coating

The bulk and shear losses of two amorphous silicon coatings, deposited using different techniques were determined using the methods outlined in Section 3.3.4. One coating was deposited using electron cyclotron resonance (ECR) ion beam deposition (IBD) whilst the other was deposited using reactive low voltage ion plating (RVLIP) deposition.

3.4.1 Bulk and shear loss of ECR IBD amorphous silicon

3.4.1.1 Sample

A silicon cantilever of 5×10^{-3} m in width, 34×10^{-3} m in length and 66×10^{-6} m in thickness was coated with amorphous silicon using ECR IBD [132]. During the deposition process, the cantilever was masked to ensure that only one face was coated. The thickness of the coating was measured at multiple points along the cantilever using a DEKTAK surface profiler. This DEKTAK measurement was performed by David Vine [133]. The results are shown in Figure 3-15.

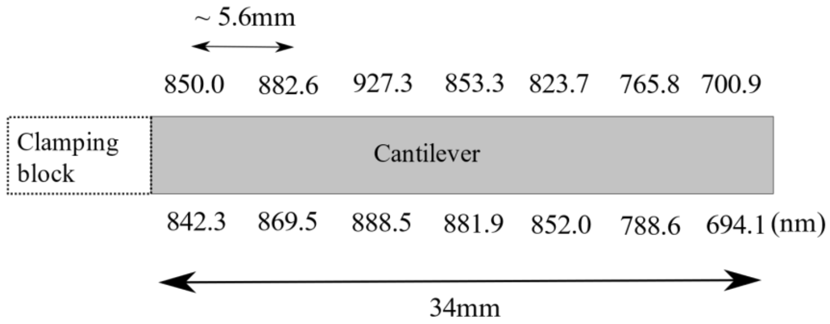

Figure 3-15: *Image of the thickness profile of the ECR IBD amorphous silicon coating on a silicon cantilever.*

The cantilever was fabricated by wet chemical etching from a (100) silicon wafer with the longest dimension of the cantilever aligned with the [110] crystal axis as shown in Figure 3-16 [134].

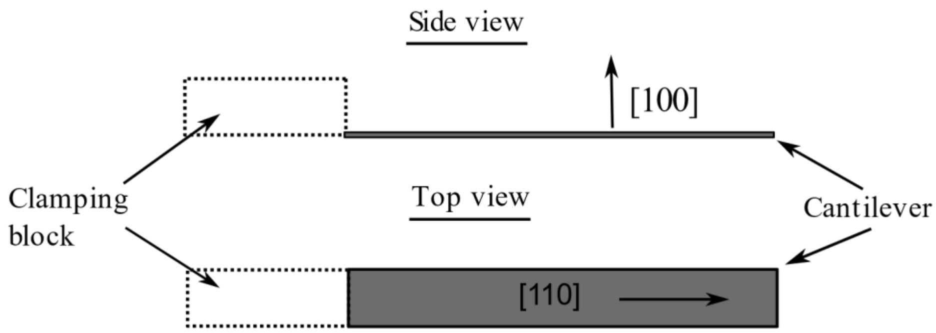

Figure 3-16: *Diagram showing the crystal orientation of the silicon cantilever.*

The thickness of the coated cantilever was estimated by using the measured bending mode frequencies of the coated cantilever and Equation (3.11). The average thickness of the coating was then subtracted from the average thickness of the coated cantilever to obtain the thickness of the cantilever (66×10^{-6} m). It should be noted that the first and second bending modes were not used to determine the thickness of this cantilever as these modes tend to significantly underestimate the thickness of a coated cantilever when using Equation

(3.11). Table 3-1, shows the calculated thicknesses of three different cantilevers based on using their first six bending modes. It can be observed that as the bending mode number increases, greater convergence of the cantilever thicknesses can be observed.

Bending mode number	**Cantilever 1 thickness (m)**	**Cantilever 2 thickness (m)**	**Cantilever 3 thickness (m)**
1	5.38×10^{-5}	3.95×10^{-5}	5.71×10^{-5}
2	6.97×10^{-5}	4.24×10^{-5}	6.07×10^{-5}
3	7.32×10^{-5}	4.31×10^{-5}	6.23×10^{-5}
4	7.38×10^{-5}	4.31×10^{-5}	6.23×10^{-5}
5	7.39×10^{-5}	4.32×10^{-5}	6.26×10^{-5}
6	7.39×10^{-5}	4.31×10^{-5}	6.26×10^{-5}

Table 3-1: Lists the calculated thicknesses of three different cantilevers using their mode frequencies.

It is hypothesised that the first two modes have this effect of underestimating the thickness of the cantilever because the cantilevers are not perfectly flat but instead follow a D shape as shown in Figure 3-17 [135].

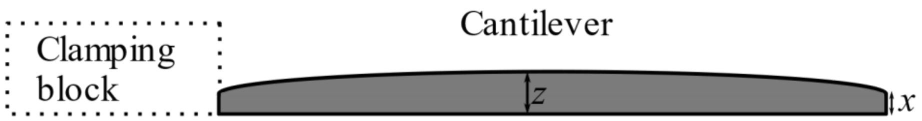

Figure 3-17: *Diagram illustrating the non-uniform thickness of the silicon cantilever. Not to scale.*

Since the first and second modes only bend near the clamping block, these modes effectively sample the thickness of the cantilever where it is thinnest (the bending of the first mode can be observed in Figure 3-5(a)). Higher order bending modes do not experience this problem as they bend (and thus sample the thickness) at multiple points along the length of the cantilever. The thickness calculated using higher order modes will therefore be more representative of the average thickness of the cantilever.

3.4.1.2 Uncoated, coated and coating losses

The losses of the coated cantilever were measured using the procedure outlined in Section 3.3.1. Since the cantilever was coated prior to this work for another project, it was not possible to measure the uncoated loss. However, it is possible to estimate the uncoated loss of a silicon cantilever, as its loss is dominated by thermoelastic loss which can be calculated by FEA. The procedure to calculate the thermoelastic loss of a silicon cantilever is described in Chapter 4. The accuracy of using this approach to estimate the uncoated loss of a silicon cantilever was tested. This involved measuring the uncoated loss of a silicon cantilever and comparing the values to the thermoelastic loss values calculated using FEA. This comparison of losses is shown in Figure 3-18. Very good agreement was obtained between the COMSOL estimates and the measured losses. Due to the frequency dependence of thermoelastic loss [136], the FEA losses were able to be linearly extrapolated to ensure the frequency of the measured and calculated losses exactly matched. The modal losses in the figure can be seen to follow two different trends corresponding to the bending modes (which follow a higher loss line) and torsional modes (which follow a lower loss line). This is due to the higher level of thermoelastic loss associated with the bulk motion of bending modes (see Chapter 4).

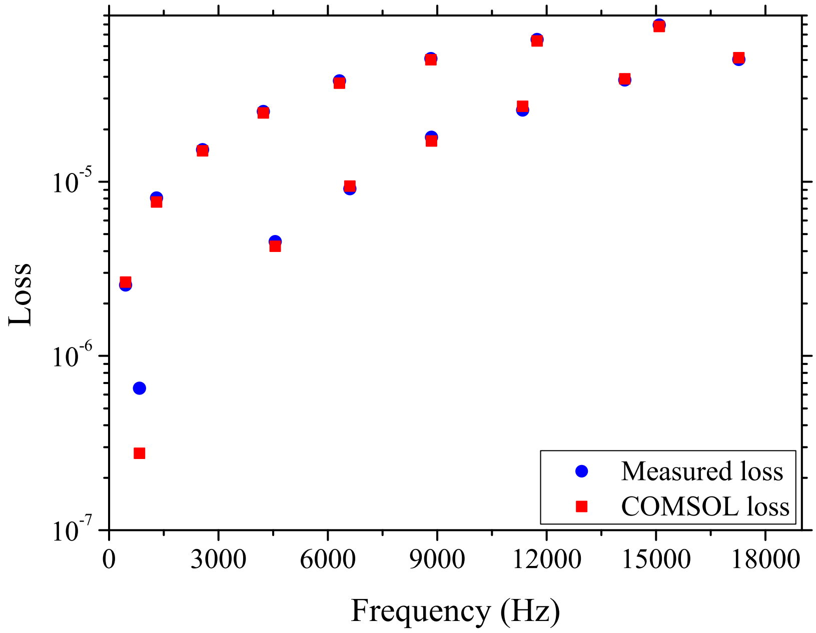

Figure 3-18: Measured loss of a cantilever against the calculated thermoelastic loss values obtained using COMSOL. The modal losses can be seen to follow two different trends corresponding to the bending modes (which follow a higher loss line) and the torsional modes (which follow a lower loss line).

For bending modes ≥ 450 Hz (second bending mode of the cantilever), COMSOL's thermoelastic loss values underestimated the measured bending mode losses by a maximum of 5 % and overestimated them by a maximum of 2 %. For torsional modes ≥ 4552 Hz (third torsional mode of the cantilever), it can be observed that COMSOL's thermoelastic loss values underestimated or overestimated the measured torsional mode losses by 6 %. The cantilever's first torsional mode at 885 Hz, shows an extremely large difference between measured and calculated loss of 58 %. The second torsional and first bending mode of the cantilever were unable to be measured. Similar level of agreement between measured and COMSOL losses were observed for two other bare silicon cantilevers (of similar thickness). Both of these cantilever COMSOL models overestimated the measured bending mode losses by a maximum of 2 % and underestimated them by a maximum of 10 % for frequencies equal to or greater than the cantilever's second bending mode. Similarly, both of these cantilever COMSOL models overestimated and underestimated the measured torsional mode losses by a maximum of 8 % for frequencies greater than or equal to the cantilever's third torsional mode.

Due to the similarity between the measured and COMSOL losses it was deemed suitable to use COMSOL's thermoelastic loss values as a substitute for the bare silicon cantilever losses. To improve similarity between the silicon cantilever's losses and COMSOL's, it was deemed that 4 % should be added to each of the calculated bending mode thermoelastic loss values as this was calculated to be the mean difference between the COMSOL and experimental values for all cantilevers measured. The mean difference for the torsional modes was 0 %. The error associated with the uncoated bending mode losses was assumed to be 6 % whilst the error associated with the uncoated torsional mode losses was assumed to be 8 %. Large variations in the difference between the measured and calculated uncoated losses of different cantilevers were observed for the 1st bending mode, 1st torsional mode and 2nd torsional mode. No repeatable offset or appropriate error bounds could be identified for these modes.

The measured coated and calculated uncoated loss values of the silicon cantilever (coated with amorphous silicon) is shown in Figure 3-19. The errors are not shown in the figure as they are very small on the y-axis scale.

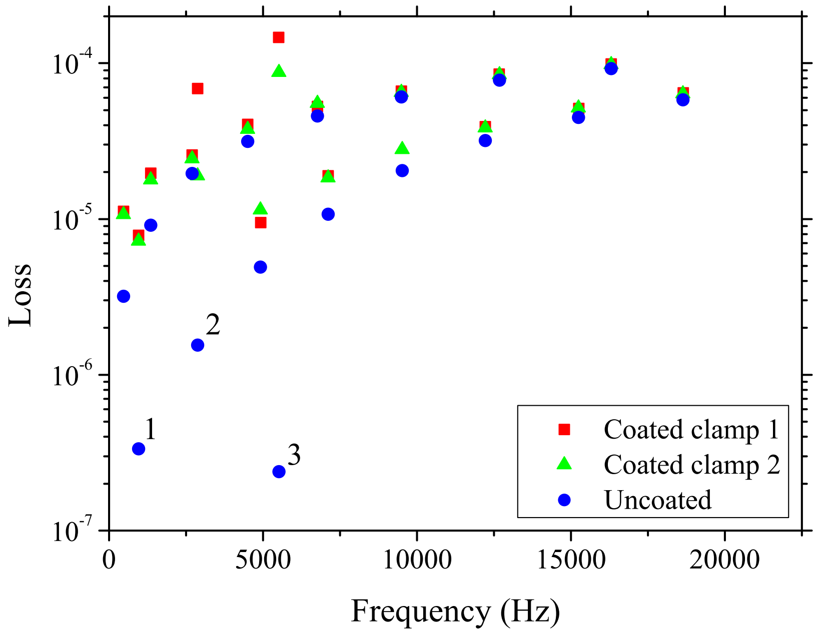

Figure 3-19: Measured loss of the silicon cantilever coated with amorphous silicon (two clamps) along with the calculated uncoated loss. The points labelled 1,2 and 3 will be referred to within the text.

The bending and torsional mode coating losses were determined from the lowest clamped coated cantilever losses and the uncoated losses using Equations (3.19) and (3.21), where the energy ratios were calculated using COMSOL. In the COMSOL model to calculate the energy ratios, it was assumed that the coating thickness did not vary along its length and had a thickness equal to the average of the DEKTAK measurements made. Three modes were excluded from the analysis. These were the 955 Hz and 2878 Hz modes (labelled 1 and 2 in Figure 3-19) as they were below the minimum torsional mode frequency at which reliable uncoated loss values could be estimated (as discussed above) while the 5507 Hz mode is a lateral mode (labelled 3 in Figure 3-19).

Figure 3-20 shows the mode-dependent coating losses. The errors were calculated using standard error propagation equations. The errors become progressively larger with frequency as the random error in the uncoated loss becomes greater with frequency. In order to compare the bending and torsional mode coating losses, the systematic uncertainties in the coating's Young's modulus and shear modulus must be accounted for (as they can shift the bending and torsional coating losses relative to one another, as shown in Equations (3.19) and (3.21)). The errors shown are therefore a combination of the random and systematic errors for each mode. Furthermore, these errors represent a one standard deviation uncertainty.

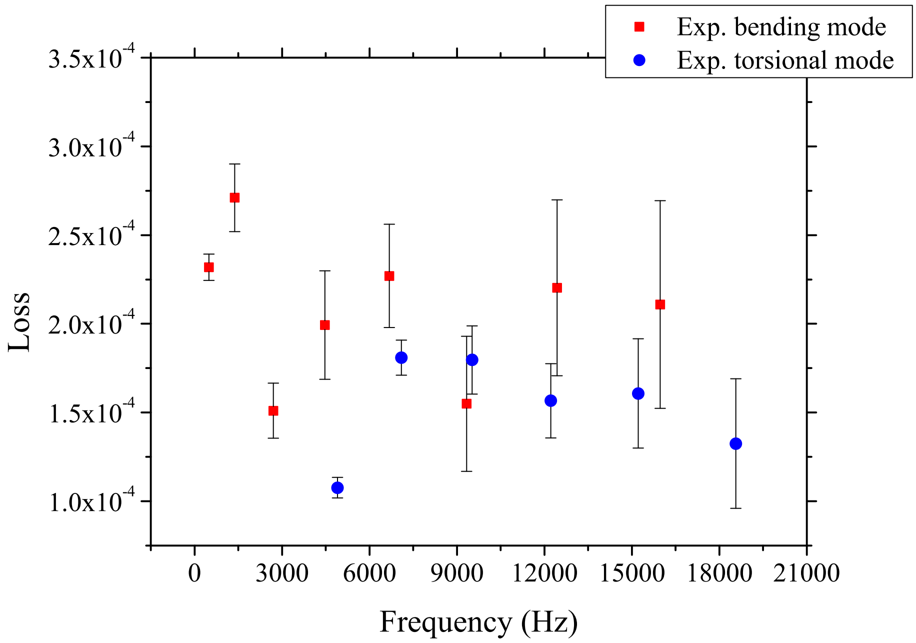

Figure 3-20: Amorphous silicon coating loss for both bending and torsional modes. The errors bars plotted are made from a combination of random and systematic errors.

It can be seen that the torsional mode coating loss tends to be lower than the bending mode coating loss. By separating each mode's coating energy into its bulk and shear components, it can be observed that bending modes store more energy in bulk motion than torsional modes as shown in Figure 3-21.

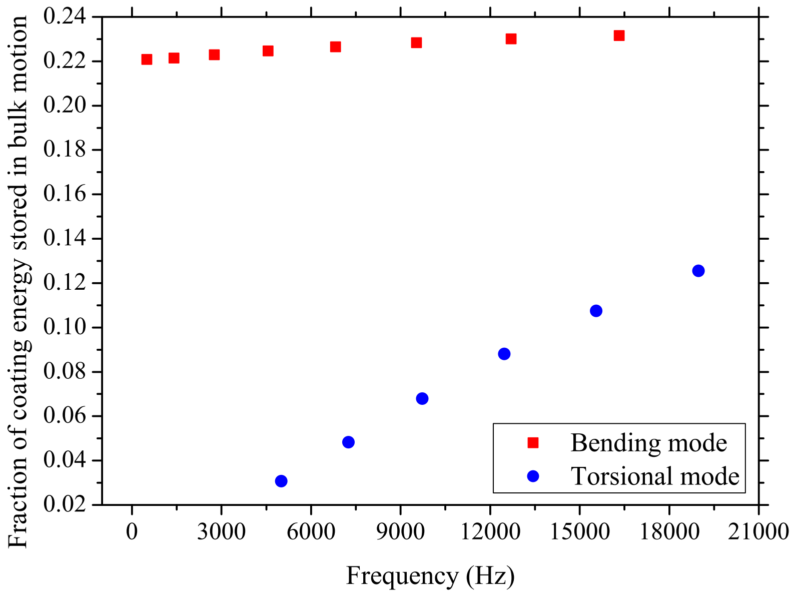

Figure 3-21: Fraction of coating energy stored in bulk motion for each of the coated cantilever's bending and torsional modes.

This observation of bending modes having a greater fraction of their energy stored in bulk motion than a torsional mode, along with the higher loss generally observed for bending modes than torsional modes (Figure 3-20), suggests that the bulk loss is likely to be greater than the shear loss in this coating.

3.4.1.3 Determining the bulk and shear losses

The two approaches described in Sections 3.3.4.1 and 3.3.4.2 were used to determine the bulk and shear losses of the amorphous silicon coating. The bulk, shear, coating and substrate energies used in this analysis were calculated from a COMSOL model that used the average coating thickness from the DEKTAK measurements. The simultaneous equation method yielded $(3.6 \pm 2.4) \times 10^{-4}$ bulk loss and $(1.4 \pm 0.7) \times 10^{-4}$ shear loss. The minimisation method calculated a bulk loss of $(3.4 \pm 0.8) \times 10^{-4}$ and shear loss of $(1.4 \pm 0.2) \times 10^{-4}$.

Figure 3-22 shows the *best fit* coating losses that were calculated using the bulk and shear losses determined by the minimisation method. Plotted against these values are the *experimental* coating losses. The experimental coating-loss values plotted in this figure are the same as those in Figure 3-20, however the errors are different. The errors bars shown in Figure 3-22 do not account for the systematic uncertainty in the coating's Young's modulus or shear modulus. They therefore represent only the random error in a mode's coating loss (as they incorporate only the random error associated with the uncoated loss and coated loss). The systematic errors were removed as it is misrepresentative to compare the fitted coating losses to the experimental coating losses when systematic errors are included as they act to shift all of the bending coating losses together and the torsional coating losses together. Including this uncertainty would therefore mean no fair comparison could be made between the trend in fitted and experimental losses. The random error in the mode's coating loss was calculated using standard error propagation techniques (i.e. addition in quadrature). From comparing the error bars in Figure 3-20 and 3-22 it can be observed that removing the systematic errors has a very small effect.

Figure 3-22 is useful in providing a check of the reasonableness of the bulk and shear loss values determined using the minimisation method. The more accurately the fitted values replicate the experimental values, the more trustworthy the calculated bulk and shear losses are. The trend followed by the bending modes and torsional modes in the fitted loss is the same as that for the experimental loss i.e. bending modes have higher coating loss than torsional modes. However, it is not possible for all of the fitted losses to simultaneously replicate the experimental losses. These differences between the fitted and experimental values could be explained by underestimating the range of uncoated losses particularly at low frequencies where this mismatch in experimental and fitted losses is greatest. To improve future analysis, the loss of the uncoated silicon cantilever should be measured prior to coating. This would enable a more trustworthy value of the uncoated loss to be used as well as help reduce the uncertainty in the experimental losses at high frequencies. Using the calculated bulk and shear losses, theoretical values of coated loss were calculated. These theoretical values are plotted against the measured values in Figure 3-23. The theoretical values can be observed to replicate the measured values well.

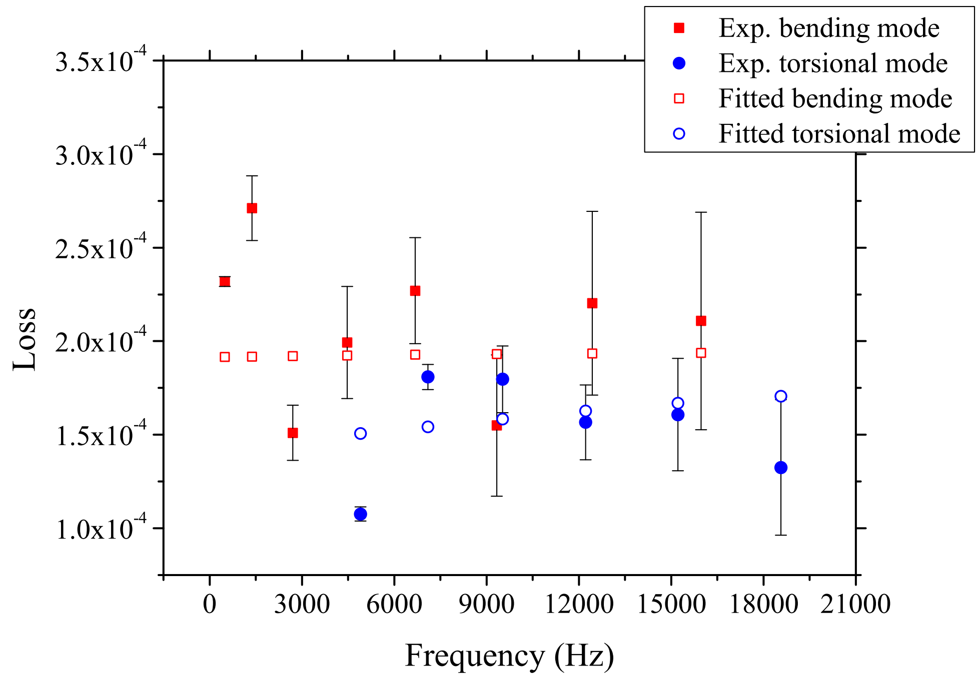

Figure 3-22: Experimental coating loss and fitted coating losses. The fitted coating losses were calculated using the minimisation method's bulk and shear loss values. The error bars shown are random errors.

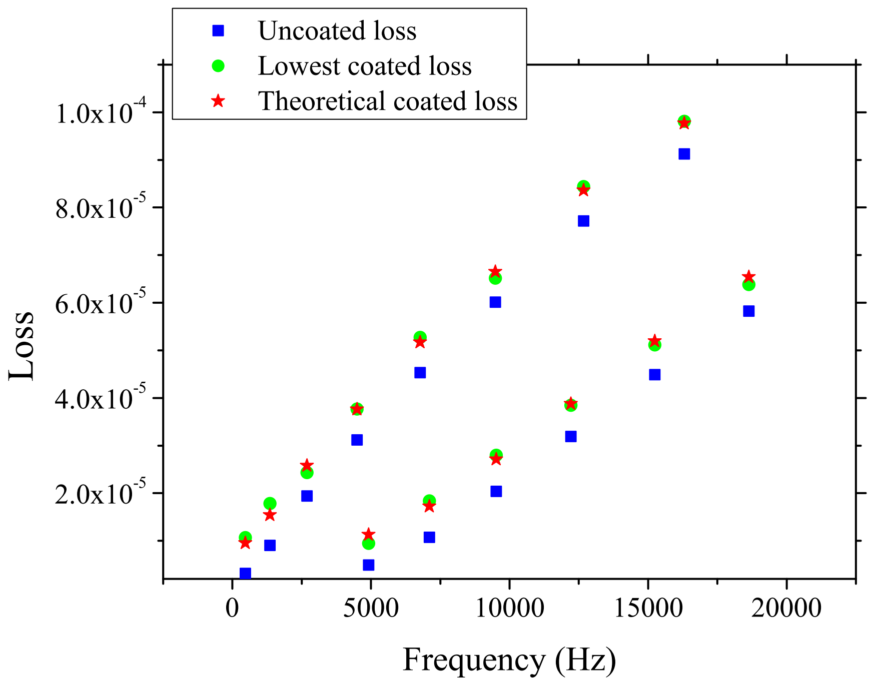

Figure 3-23: Coated and uncoated losses of a silicon cantilever coated with amorphous silicon. The coated losses were measured as well as calculated using the minimisation method's bulk and shear losses. The lowest, measured coated loss at each mode (shown in Figure 3-19) is shown. The uncoated loss was determined using COMSOL.

The errors of the bulk and shear losses (listed earlier in this section) are formed from both random and systematic errors and represent the one standard deviation range of the bulk and shear losses. However, these errors are not useful in evaluating whether the bulk and shear loss results of the two methods agree within error (as systematic errors shift the bulk and shear losses calculated using the different methods in the same way). To make this comparison, the systematic component of the errors must be ignored and purely random bulk and shear loss errors determined. These are calculated by repeating the simultaneous equation method and minimisation method when only the random errors in the coated and uncoated losses are accounted for (and no systematic error in the material properties are accounted for). The bulk and shear losses determined using the simultaneous equation method were $(3.6 \pm 2.3) \times 10^{-4}$ and $(1.4 \pm 0.7) \times 10^{-4}$ respectively while those calculated using the minimisation method were $(3.4 \pm 0.7) \times 10^{-4}$ and $(1.4 \pm 0.2) \times 10^{-4}$ respectively. While the bulk and shear losses from the two methods show agreement to within error, the error of the simultaneous equation method can be seen to be much larger than that of the minimisation method. As described in Section 3.3.4.2, the error of the simultaneous equation method is larger because it only uses two data points (one bending and one torsional mode) to estimate a bulk and shear loss at any one time which makes this method more prone to calculating outlying bulk and shear loss values. Due to the improvement in accuracy of the minimisation method, only this method will be used to solve for bulk and shear losses in the rest of this thesis.

The minimisation methods bulk and shear losses (with total error) are significantly different from one another. This is the first published bulk and shear loss analysis, known to the author, to reach this conclusion [137]. The implication of this result for gravitational wave detectors is extremely significant as it demonstrates the possibility of being able to lower thermal noise by carefully selecting a coating based upon it's bulk and shear losses (as illustrated in Figure 3-3). This marks a shift away from the current approach of identifying the optimum coating to use in a gravitational wave detector as being the one with lowest mode dependent coating loss.

A highly reflective mirror coating in a gravitational wave detector is formed by a stack of alternating layers of two materials. The reflectivity depends upon the difference in refractive index between the two materials, and on the number of pairs of layers which are used. Here, the thermal noise is calculated for a simplified coating composed entirely of amorphous silicon – i.e. assuming identical loss and material properties between the high and low index layers. The coating thermal noise was calculated using both Hong's method (Equations 3.6 to 3.8) and the parallel and perpendicular approach (Equation (2.26) when $\phi_{\text{substrate}} = 0$ and $\phi_{\parallel} = \phi_{\perp}$). Ideally it would have been best to calculate the thermal noise of a gravitational wave detector coating whilst accounting for its multi-layer structure using Hong's method and the parallel and perpendicular approach, however this is not possible as no estimate for the bulk and shear loss of a low index silica layer could be found.

The bulk and shear losses obtained from the minimisation method were used in Hong's equation. The mechanical loss used in the parallel and perpendicular thermal noise calculation was an average of the mode dependent coating losses ($\phi_{\parallel} = \phi_{\perp} =$ average $\phi_{coating}$). In these thermal noise calculations, a coating thickness of 2.6×10^{-6} m was assumed as this is the required thickness of an amorphous silicon/silica bilayer coating deposited upon a silicon test mass which would satisfy the reflectivity requirements of a gravitational wave detector at 1550 nm. It was assumed that the coating and test mass were held at a constant temperature of 120 K (LIGO Voyager temperature).

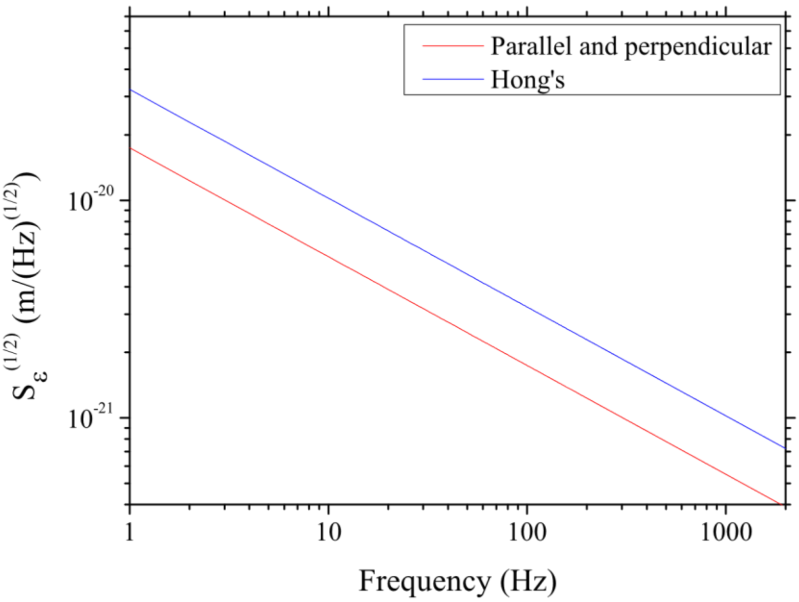

Figure 3-24: Thermal noise of a coating assumed to be composed entirely of amorphous silicon. The thermal noise was calculated using Hong's method as well as the parallel and perpendicular approach. Hong's method used the bulk and shear losses obtained from the minimisation method whilst the parallel and perpendicular approach used an average of the mode dependent coating losses. As convention dictates the square root of the thermal noise power spectral density is plotted.

Figure 3-24 shows the results of these calculations. The parallel and perpendicular method (which as mentioned earlier is the currently used method in the gravitational wave community to estimate a coating's thermal noise) estimates the coating's thermal noise to be 1.8× less than Hong's method. Whilst the absolute values in this figure are not directly relatable to a gravitational wave detector as it is not a multi-layer coating, this graph is useful in illustrating that Hong's approach to calculating thermal noise can result in a significantly different level of thermal noise compared to the traditional method. This difference has perhaps already been identified by Gras et al [138] who measured directly the thermal noise of the Advanced LIGO coating. The measured value was 22 % higher than that calculated using the parallel and perpendicular thermal noise equation. Abernathy et al [137] speculate that 82 % of this 22 % discrepancy is due to not using Hong's method to calculate thermal noise. Further direct coating thermal noise measurements are clearly of interest to fully verify that Hong's method gives the correct (or more accurate) results than the parallel and perpendicular approach.

The A+ future detector (discussed in Chapter 1) aims for a factor of 2 improvement in the Advanced LIGO detector's coating thermal noise. Using a coating thermal noise model (parallel and perpendicular approach) which could potentially be wrong by a factor 1.8 is therefore not particularly helpful.

3.4.1.4 Dependence of the bulk and shear losses on frequency

A recently published paper by Abernathy et al [137] calculates the bulk and shear losses of a 25 % titania doped tantala coating deposited on a silica disc. The method used to calculate the bulk and shear losses is not too dissimilar from the simultaneous equations method outlined in this thesis. However, one notable difference is that Abernathy calculates two values of bulk and shear loss: one set which is frequency independent and the other which is frequency dependent. The significance of the frequency dependent bulk and shear losses is not tested, and it is therefore not clear whether or not the losses are frequency dependent or not. Furthermore, their frequency dependent model results in unphysical negative bulk losses at high frequencies.

As a result of this uncertainty surrounding the frequency dependent nature of the bulk and shear loss, the frequency dependence of the amorphous silicon's bulk and shear losses was investigated. If a significant frequency dependent bulk and shear loss exists then the frequency dependent effect should be observed within the residuals of the fits in Figure 3-22. The residuals (r) were obtained by subtracting the experimental values from the best fit values. They were then regressed using the equation:

$$r_x = f_x \times \phi_{BB} \times \frac{U_{\mathrm{bulk},x}}{U_{c,x}} + f_x \times \phi_{SS} \times \frac{U_{\mathrm{shear},x}}{U_{c,x}}\,, \tag{3.26}$$

where f represents the frequency and x is used to denote the different mode numbers. The frequency dependent bulk and shear losses were obtained as being $\phi_{BB} = -1.55 \times 10^{-9}$ and $\phi_{SS} = -1.09 \times 10^{-9}$ respectively. Since both values are negative, this means that as the frequency increases the bulk and shear losses of the coating become smaller until they become negative, which is unphysical (as observed by Abernathy et al [137]). However, it is important to note the P values of these frequency dependent bulk and shear loss coefficients to identify whether these frequency dependent losses are significantly different from 0. P-values represent the probability of observing a coefficient's value given that a null hypothesis is true. For a null hypothesis of each coefficient (the frequency dependent bulk and shear coefficients) being equal to 0, P-values of 0.89 and 0.71 were obtained respectively. Within statistics, a null hypothesis is commonly rejected when the P-value is below 0.01, 0.05 or 0.1. The P-values determined from the regression are larger than 0.01, 0.05 and 0.1, and therefore the null hypothesis cannot be rejected (at a 1 %, 5 % or 10 % significance level). This result implies that these coefficients could likely be equal to 0. Due to the lack of significance shown for the amorphous silicon's bulk and shear frequency dependent losses as well as the untested significance of Abernathy et al's frequency dependent values [137] (described above), no compelling evidence points in the direction of a frequency dependent bulk and shear loss. For this reason, the coatings analysed in this thesis will only calculate frequency independent bulk and shear losses.

3.4.1.5 Effect of non-uniform coating thickness on the bulk and shear losses

The bulk and shear analysis presented in Sections 3.4.1.3 and 3.4.1.4 assumed a uniform coating thickness. As shown in Figure 3-15, this is known not to be the case. To assess the effects of this non-uniform thickness upon the results, a COMSOL model incorporating variations in coating thickness was constructed. In the COMSOL model, the thickness of the coating was approximated such that its thickness at the edges of the cantilever length and at the thickest point on the cantilever matched the measurements (which were the average of the two measurements at the same position along the cantilever length). In the COMSOL model, linear interpolation was then used to determine the coating thickness between these positions. Figure 3-25 compares the measured thickness profile with this model interpolation. The interpolated model replicates the measured values to within 7%.

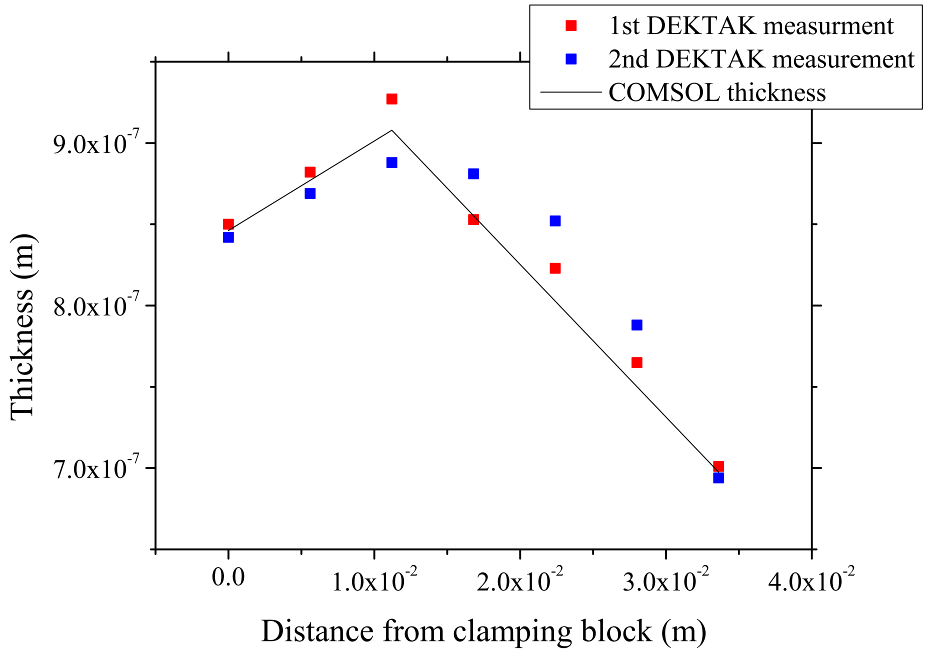

Figure 3-25: Coating thickness measured with DEKTAK against the approximated thickness used in the COMSOL model. The two sets of DEKTAK measurements taken along both sides of the cantilever length are plotted.

Using the minimisation method, the values of the amorphous silicon's bulk and shear losses were determined as being $(3.6 \pm 0.7) \times 10^{-4}$ and $(1.4 \pm 0.1) \times 10^{-4}$ respectively. The errors listed are the total error (systematic and random). Repeating the analysis, the purely random errors of the bulk and shear losses were determined to be $\pm 0.7 \times 10^{-4}$ and $\pm 0.1 \times 10^{-4}$ By comparing these random errors to those obtained assuming a uniformly thick coating ($(3.4 \pm 0.7) \times 10^{-4}$ bulk loss and $(1.4 \pm 0.2) \times 10^{-4}$ shear loss), it can be observed that accounting for thickness variations of this magnitude had no significant effect upon the bulk and shear losses calculated.

3.4.2 Bulk and shear loss of RLVIP

A bulk-shear analysis was carried out upon another amorphous silicon coating deposited by an alternative method called reactive low voltage ion plating (RLVIP) [139, 140].

3.4.2.1 Sample

A 1×10^{-6} m thick amorphous silicon coating was deposited upon a silica cantilever [140]. Using the same procedure as outlined in Section 3.4.1.1, the thickness of the cantilever was calculated to be 163×10^{-6} m. An approximate illustration of the silica cantilever is shown in Figure 3-4 (b).

3.4.2.2 Uncoated, coated and coating losses

Ownership of the cantilever was gained post-deposition. As a result, no uncoated loss values were measured. Silica cantilevers are made by hand welding a silica flexure to a silica clamping block using a 100 W CO_2 laser. The weld for each cantilever is unique, and these cantilevers can therefore display a relatively wide range of mechanical loss values. An empirical model for the loss of uncoated silica cantilevers was constructed from measurements of the loss of five uncoated cantilevers. Each cantilever was clamped a minimum of two times and the lowest measured loss for each of the cantilever's modes was plotted. Lines of maximum and minimum uncoated loss were then drawn on the graph such that they encapsulated all of the lowest experimentally measured uncoated losses. The result of this exercise is shown in Figure 3-26, where the two pink lines encapsulate all of the lowest measured bending mode losses whilst the orange lines do the same for the torsional modes. It should be noted that the pink lines do not go below 450 Hz in frequency and the orange lines below 1480 Hz as these were the lowest resonant modes measured for the uncoated cantilevers. Figure 3-26 also shows the loss of the coated cantilever as-deposited (purple diamonds: bending modes, hollow purple diamonds: torsional modes), after heat treatment for one hour at 300 °C (green squares: bending modes, hollow green squares: torsional modes) and then after heat treatment for 1 hour at 500 °C (blue circles: bending modes, hollow blue circles: torsional modes). Heat treatment was carried out in air. The cantilever was in the oven as the oven heated up and cooled down, with the heat treatment time corresponding to the length of time at the maximum temperature.

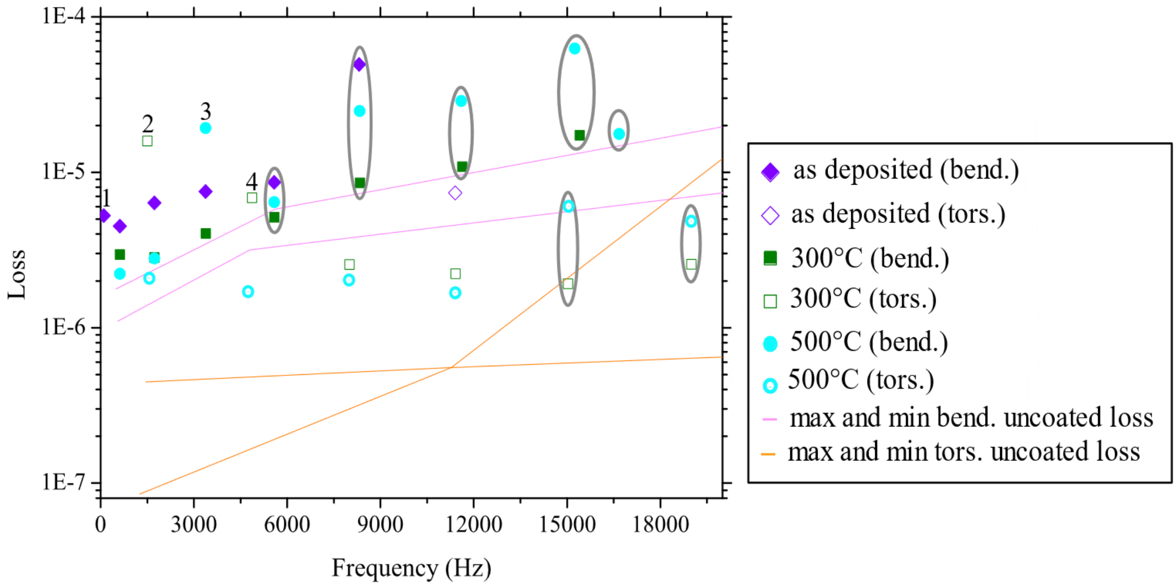

Figure 3-26: Measured loss of a silica cantilever coated with 1 um amorphous silica after deposition (purple diamonds: bending modes, hollow purple diamonds: torsional modes), after heat treatment for one hour at 300 °C (green squares: bending modes, hollow green squares: torsional modes) and after heat treatment for 1 hour at 500 °C (blue circles: bending modes, hollow blue circles: torsional modes). The lines indicate the maximum and minimum loss of the uncoated fused silica cantilevers, where the bending mode loss is between the two pink lines and the torsional mode loss between the two orange lines. The modes which are encapsulated by a grey circle and the modes marked with numbers 1-4 are discussed separately in Section 3.4.2.2

The loss of most modes of the coated cantilever reduces following heat-treatment. In some cases – possibly due to clamping loss effects – some modes do not follow this general trend. The bending and torsional mode coating losses were determined using Equations (3.19) and (3.21). The central value from the range of uncoated losses was used in these coating loss calculations, with the error in loss of the coating estimated using the maximum and minimum values of uncoated loss. The results are plotted in Figure 3-27, 3-28 and 3-29. It should be noted that not all modes were used to estimate the loss of the coating. The reasons for excluding some of these measurements are listed below:

- No coating losses were determined for bending modes at frequencies greater than 5500 Hz and for torsional modes greater than 11000 Hz as the spread of uncoated losses were too large to conclude anything significant. The excluded coated losses are encapsulated by grey circles in Figure 3-26.
- The loss value labelled 1 was not included as no modes in similar frequency were measured for the uncoated cantilevers that were used to estimate an upper and lower

bound of uncoated loss. If this point is used, an unquantifiable error would be associated with its uncoated loss and therefore the loss of the coating calculated.

- The coated losses labelled 2 and 4 were excluded as they break the trend of other modes where the torsional mode coated losses are less than the bending mode coated losses of similar frequency.
- The loss value labelled 3 was excluded as it broke a clear trend which all of the other measurements adhered to: heat treated loss being less than as deposited loss.

3.4.2.3 Bulk and shear losses of the coating

The minimisation method described in Section 3.3.4.2 was used to determine the bulk and shear losses of the coating in each of its three heat-treatment states. The results are presented in Table 3-2. For all heat-treatments of the coating, the bulk loss was found to be zero (this will be discussed later in more detail), while the shear loss was found to decrease with heat-treatment. The errors listed in the table are one standard deviation total errors (systematic and random errors). Figures 3-27, 3-28 and 3-29 compare the fitted and measured mode-dependent coating losses and show the quality of the fitting process. Two error bars are plotted for each mode's coating loss. The larger error bar is a combination of the coating loss's random error and systematic error (same as Figure 3-20). This uncertainty is useful in determining whether the bending and torsional mode coating losses are significantly different from one another. The smaller error bar is only the coating loss's random error (same as Figure 3-22). This error bar is useful for comparing the similarity between the fitted and experimental values. For most points, these error bars are very similar in size.

The best estimates of the bulk and shear loss values are able to fit the experimental mode-dependent coating-loss trends well as they consistently replicate the experimental trend of torsional coating loss being greater than bending-mode coating loss. However, discrepancies do exist between some of the measured and calculated loss values. This can most likely be explained by an underestimation of the spread in the uncoated losses or by excess loss associated with some modes of the coated sample (e.g. clamping loss effects at particular frequencies), which could have been improved by re-clamping the cantilevers more times. If more bending and torsional modes were able to be measured for the different heat treatments, coating losses which have extra loss associated with them would have been easier to identify as they would not follow the fitted trend. It is clear that all fits would have benefitted from more data points. Figure 3-30 is a 3D plot which helps show the improvement in the as-deposited mode dependent coating loss fit (Figure 3-27) as the theoretical bulk and shear losses are varied. It shows Figure 3-27's squared residual (the squared difference between a mode's fitted coating loss and its experimental value, summed over all modes) against random values of bulk and shear loss. The bulk and shear losses quoted in Table 3-2 can be seen to be the *best values* as they minimise the Figure 3-27's squared residual.

	Value	**Error**
As deposited		
Bulk loss	0	1×10^{-7}
Shear loss	1.53×10^{-4}	9×10^{-6}
300 °C		
Bulk loss	0	1×10^{-7}
Shear loss	4.9×10^{-5}	5×10^{-6}
500 °C		
Bulk loss	0	1×10^{-7}
Shear loss	3.8×10^{-5}	5×10^{-6}

Table 3-2: Table showing the calculated amorphous silicon's bulk and shear loss values when using the minimisation method. The table presents the losses for the as-deposited, 300 °C and 500 °C heat treatments. The table also lists the one standard deviation total error (systematic and random errors combined) of the bulk and shear losses.

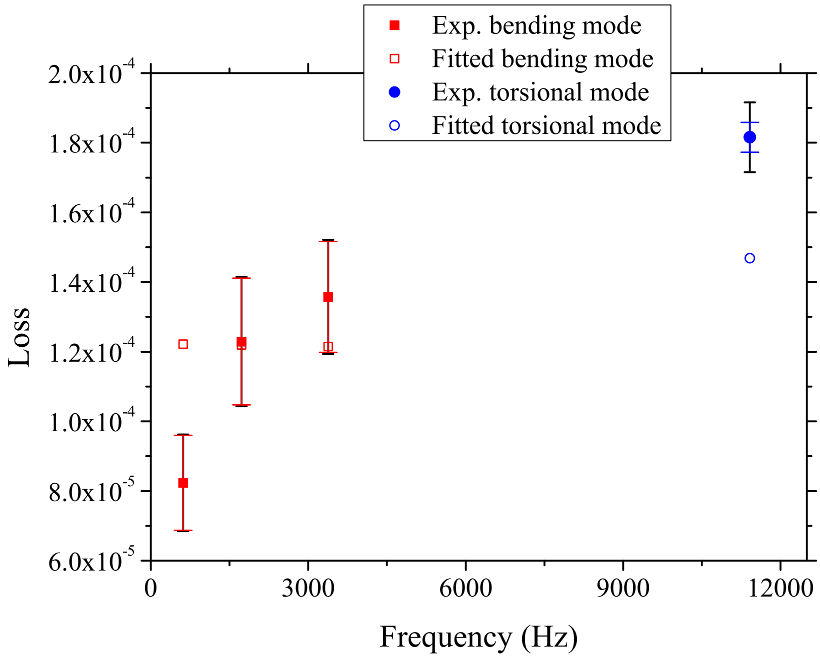

Figure 3-27: Experimental and fitted as-deposited amorphous silicon coating losses. The fitted coating losses were calculated using the bulk and shear losses obtained from the minimisation method. The black error bars represent the total error (systematic and random error) in the coating's loss. The blue and red error bars represent only the random uncertainty in the coating's loss.

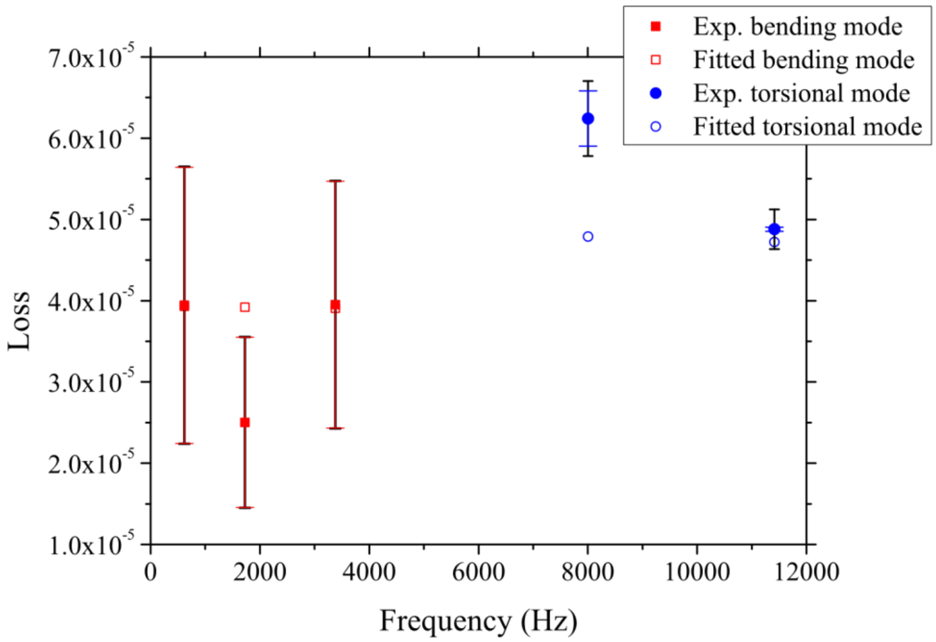

Figure 3-28: Experimental and fitted 300 °C heat treated amorphous silicon coating losses. The fitted coating losses were calculated using the bulk and shear losses obtained from the minimisation method. The black error bars represent the total error (systematic and random error) in coating loss. The blue and red error bars represent only the random uncertainty in coating loss.

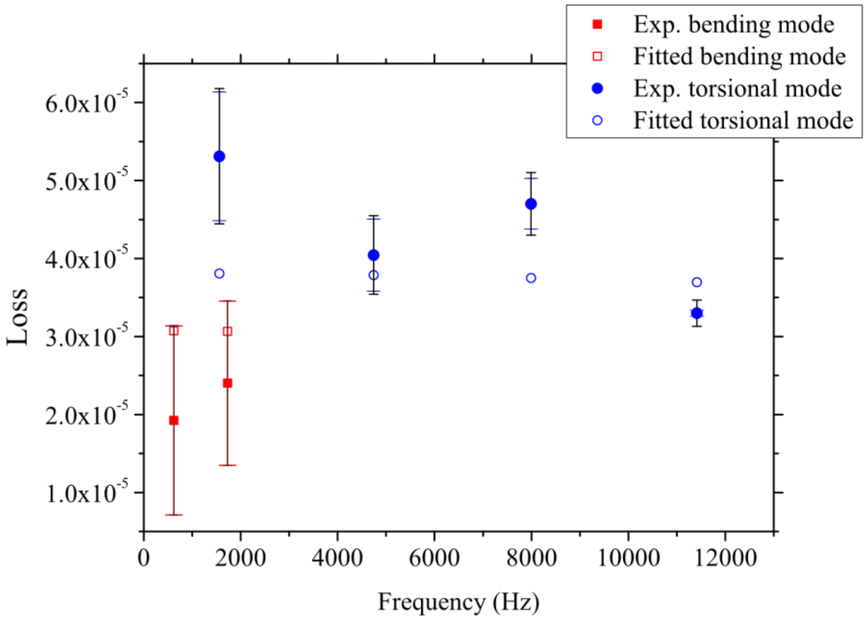

Figure 3-29: Experimental and fitted 500 °C heat treated amorphous silicon coating losses. The fitted coating losses were calculated using the bulk and shear losses obtained from the minimisation method. The black error bars represent the total error (systematic and random error) in coating loss. The blue and red error bars represent only the random uncertainty in coating loss.

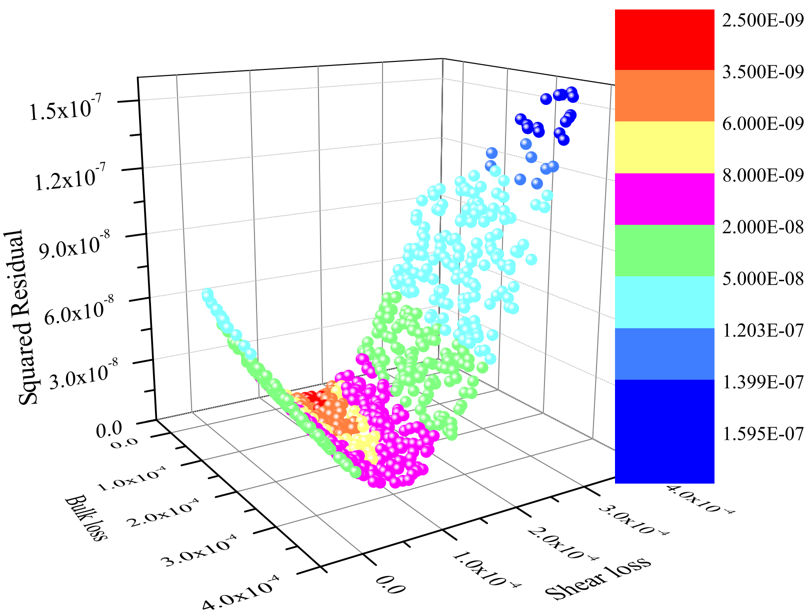

Figure 3-30: This figure illustrates the best fit bulk and shear losses for the as-deposited amorphous silicon coating. It shows the squared residual (the squared difference between a mode's experimental coating loss and fitted values, summed over all modes) for Figure 3-27 against random values of bulk and shear loss. Red points indicate fitted coating losses which replicate the experimental values well (in Figure 3-27) whilst blue points represent fitted values which represent the experimental values poorly.

Table 3-2 lists the total errors of the bulk and shear losses (systematic and random error). To determine whether the shear losses are significantly different from one another at different temperatures, the shear losses random errors should be calculated. However, since the shear losses are significantly different from each other at all heat treatments when using the total error and the total error is larger than just the random error, then the shear losses must also be significantly different from each other at all heat treatments when using random errors. It can therefore be concluded that the shear losses of the amorphous silicon coating are significantly different from each other for all heat treatments.

As noted above the bulk losses are all found to be equal to zero. Zero represents the minimum value which the bulk loss can take in the minimisation method. When this minimum value is marginally increased, the bulk loss is determined to be equal to this new value. Whilst this value of 0 bulk loss is perhaps unlikely and possibly represents a limitation of the minimisation method, the results strongly suggest that the shear loss is much larger than the bulk loss in the amorphous silicon material. The general trend of torsional modes having higher experimental coating loss than the bending modes is consistent with the shear loss dominating over the bulk loss in this coating, as torsional modes contain a greater proportion of their energy in shear motion. The error associated with the bulk loss was estimated from the "step size" specified in the VBA code used to estimate the bulk and shear losses of the coating (where the "step size" defines the minimum variation in the bulk and shear losses considered by the VBA fitting program). This "step size" was specified to be 1×10^{-7} in the program. Since the values of bulk loss for all of the 100 simulations performed by the VBA program were equal to 0, the error in bulk loss was estimated to be 1×10^{-7}. The shear loss error was defined as normal, equal to the standard deviation of the 100 calculated shear loss values.

Finding that the shear loss of a coating is greater than its bulk loss is an important result, as coating thermal noise is more sensitive to bulk loss than shear loss (see Section 3.2). As a result, this coating would be expected to have lower thermal noise compared to other coatings which have the same mode dependent coating loss but with a higher bulk loss than shear. This improvement in using a coating with low bulk loss would go unnoticed if the parallel and perpendicular method to calculating thermal noise is used. Figure 3-31 shows the thermal noise of a 2.6×10^{-6} m thick gravitational wave detector coating assumed to be composed entirely of as-deposited amorphous silicon. The thermal noise was calculated using the widely used parallel and perpendicular approach (with an average mode dependent coating loss) and Hong's method (with the calculated bulk and shear losses). A similar calculation was carried out for the ECR amorphous silicon coating (Section 3.4.1.3) It can be observed in Figure 3-31 that the parallel and perpendicular method overestimates Hong's thermal noise by a factor of 1.7.

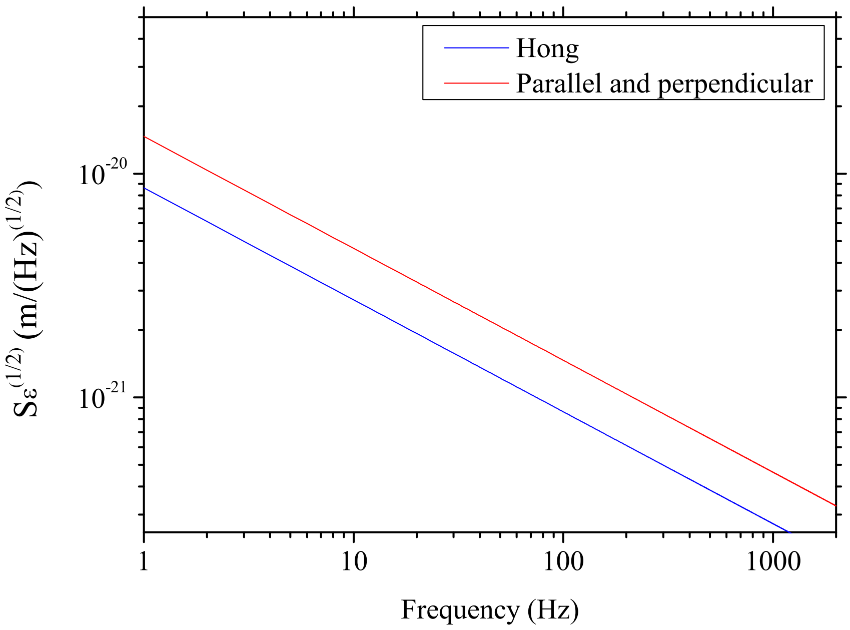

Figure 3-31: Thermal noise of a single layer amorphous silicon coating with thickness 2.6×10^{-6} *m (which is the required thickness of an amorphous silicon/silica bilayer coating that would satisfy the reflectivity requirements of a gravitational wave detector). The thermal noise was calculated using Hong's method as well as the parallel and perpendicular approach. As convention dictates the square root of the thermal noise power spectral density is plotted.*

Figure 3-32 shows the effect of heat-treatment upon the thermal noise of the amorphous silicon coating, calculated using Hong's method and using the same assumptions as for Figure 3-31. The observed reduction in mechanical loss with heat-treatment translates into a significant reduction in thermal noise. Whilst the thermal noise values shown in Figure 3-32 do not represent the noise of a detector coating (as a detector coating is made from two materials and the bulk and shear losses of a low index material which could accompany amorphous silicon is unknown), this figure is useful in highlighting the trend of a lower thermal noise for the gravitational wave detector's coating if higher heat treated amorphous silicon was used as one of the detector's alternating layers.

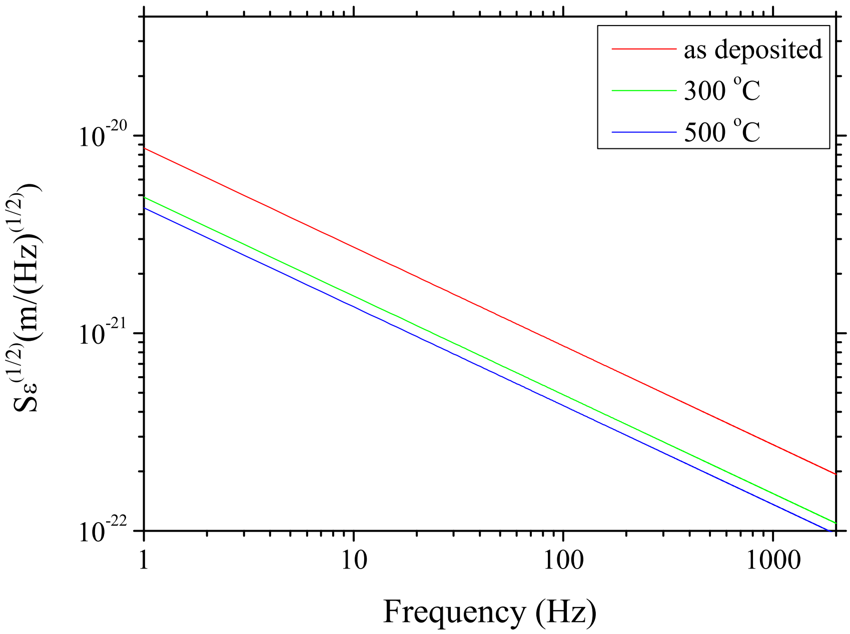

Figure 3-32: Thermal noise of a coating assumed to be composed entirely of as-deposited, 300 °C heat treated and 500 °C heat treated amorphous silicon. As convention dictates the square root of the thermal noise power spectral density is plotted.

Comparing the bulk and shear losses of the RLVIP amorphous silicon to the ECR IBD amorphous silicon it can be observed that the former was shear loss dominated whilst the latter was bulk loss dominated. It is speculated that this difference in loss could be due to the different deposition techniques creating dissimilar types of defects which couple differently to the loss mechanisms. Further investigations will need to be conducted to verify this.

3.5 Bulk and shear losses of a coating deposited onto two different substrate geometries

Section 3.4 showed that it is possible to determine the bulk and shear losses of a coating deposited upon a cantilever. However, a more powerful test of Hong's theory is to deposit a coating onto two different substrate geometries and determine whether - as predicted - the coating on each substrate exhibits the same bulk and shear losses.

To investigate this a silica cantilever and a silica disc were coated with a 500×10^{-9} m thick tantala coating using RLVIP deposition [139, 140]. The bulk and shear losses of the coating were determined for both substrates and compared.

3.5.1 Determining the bulk and shear losses of the tantala coating deposited onto a silica disc

3.5.1.1 Sample

Discs made of silica [118], silicon [141] and sapphire [142] are commonly used as substrates in mechanical loss experiments. The silica disc used in this experiment was measured to be 2.57×10^{-3} m in thickness and 38.07×10^{-3} m in radius. The most common resonant mode shapes of a disc with this geometry are often called butterfly and drum modes. Examples of these mode shapes are shown in Figure 3-33.

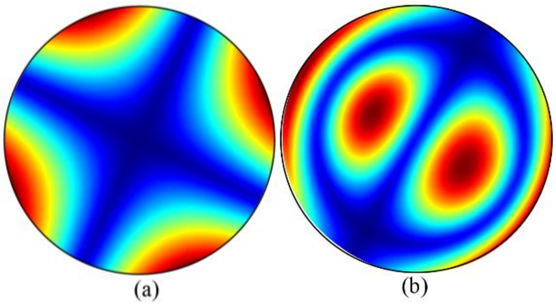

Figure 3-33: Images of a disc's (a) 1st order butterfly mode and (b) 2nd order drum mode. Red illustrates maximum deformation whilst blue represents minimum.

3.5.1.2 Method to measure the mechanical loss of a coated and uncoated disc

The mechanical loss was measured using a ring-down technique similar to that described in Section 3.3.1.2. The disc was held stationary inside a vacuum tank by two 50×10^{-6} m diameter tungsten wires. These wires support the disc at two points on it's circumference, 180 degrees apart, as illustrated in Figure 3-34. The amplitude of the disc's motion was measured using a SIOS SPS-120/500 laser interferometer [143]. The loss of each mode was measured at least 5 times and an average was taken. The disc was re-suspended 3 more times and the losses remeasured. The lowest loss measured at each mode for the different suspensions was assumed to be the best estimate of the loss of the disc.

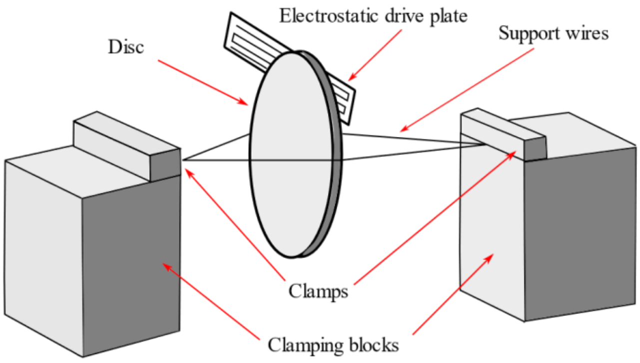

Figure 3-34: Setup up used to support a disc and excite it for ring down measurements.

3.5.1.3 Mode dependent coating loss

The mode-dependent coating loss was determined as follows:

$$\phi_{coating} = \frac{U_s}{U_c}(\phi_{coated} - \phi_{uncoated})\,. \tag{3.27}$$

This formula is a variation of Equations (3.19) and (3.21). The elastic strain energies stored in the coating and substrate were determined using a COMSOL model following the procedure used for the cantilevers described in Section 3.3.4.2. The disc was modelled with no fixed constraints (i.e. no suspension points). This is in line with others in the literature [141].

3.5.1.4 Uncoated, coated and coating losses

The uncoated disc was suspended 4 times and its modal losses measured each time. The results are shown in Figure 3-35 and discussed below. Error bars are approximately $\pm 1 \times 10^{-8}$ and are too small to see on the graph. It is interesting to note that the higher frequency modes have a greater spread of losses than the lower frequency modes. This is due to the higher frequency modes having quicker ring downs (than lower frequency modes) and therefore less points to fit to in determining the disc's mechanical loss.

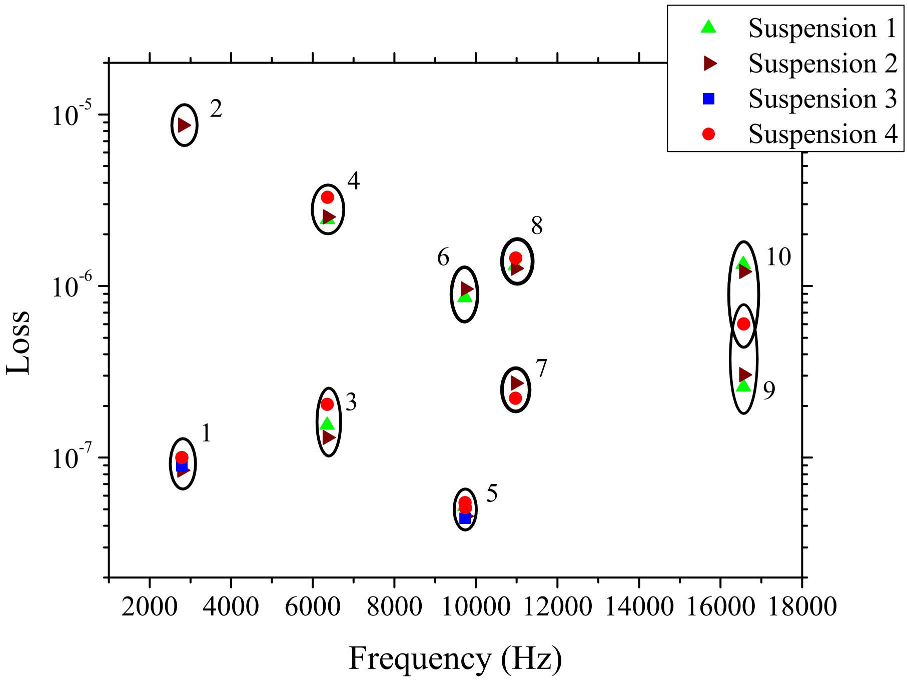

Figure 3-35: Uncoated loss of a 3" silica disc. The uncoated disc was suspended 4 times and its modal losses measured each time. These four sets of disc losses are plotted. The losses are grouped and labelled by mode number. Since the 4th suspension only measured one mode at 16500 Hz and the loss of this mode fell between the other modes measured with the different suspensions, it was not possible to identify which group it belonged to.

The losses for two suspensions of the coated disc are shown in Figure 3-36.

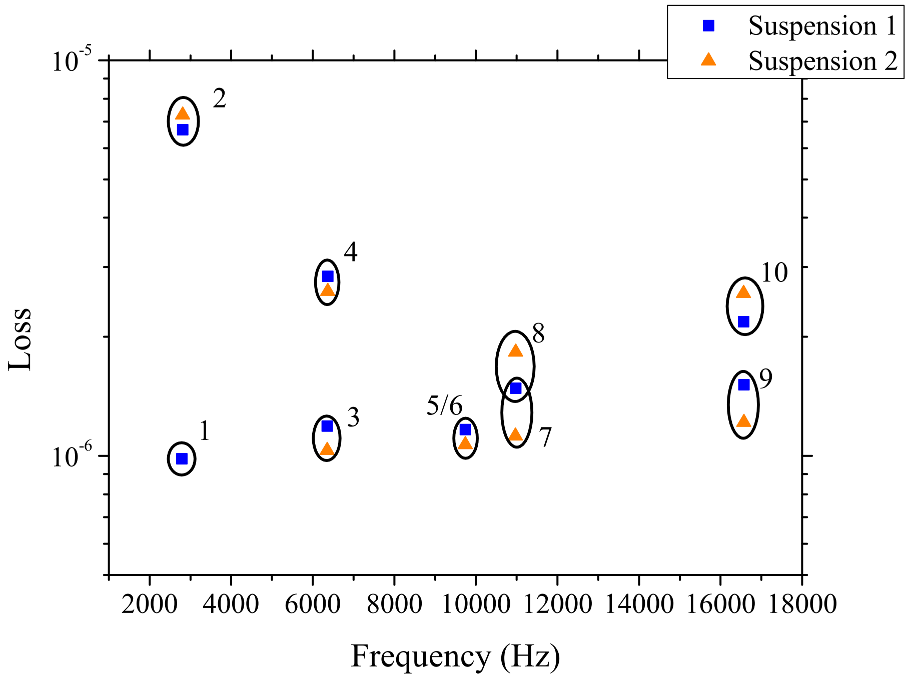

Figure 3-36: Losses of the silica disc coated with a 500×10^{-9} m thick layer of tantala. The coated disc was suspended 2 times and its modal losses measured each time. The losses are grouped and labelled by mode number. Since the 1st suspension only measured one mode at 11000 Hz and the loss of this mode fell between the other modes measured with the different suspension, it was not possible to identify which group it belonged to.

The tantala coating losses were determined using the uncoated disc losses (Figure 3-35), coated disc losses (Figure 3-35) and Equation (3.27). The losses of the coating are shown in Figure 3-37.

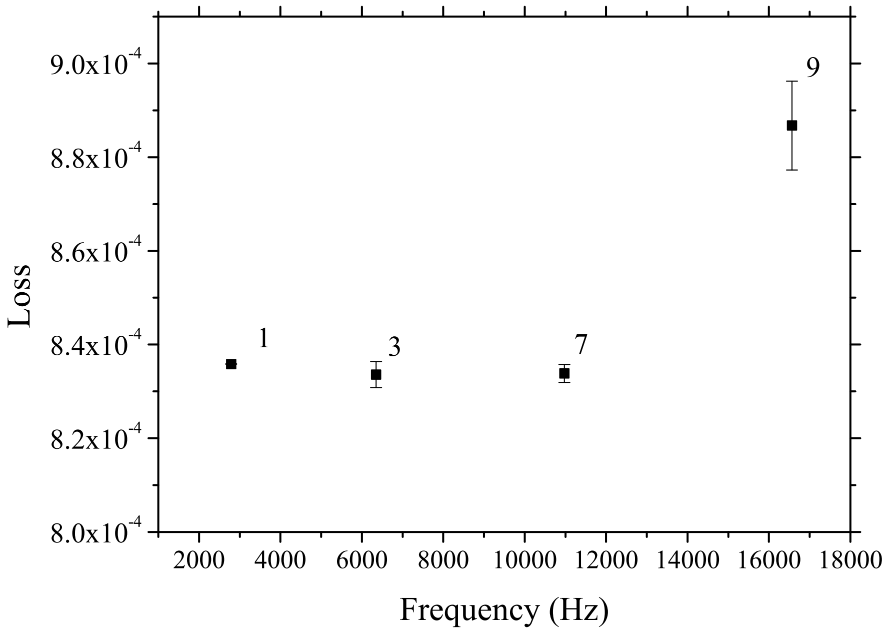

Figure 3-37: Losses of the tantala coating deposited upon the silica disc. Only some of the modes shown in Figures 3-35 and 3-36 were used to calculate the coating's loss. The reasons are explained in the text.

It should be noted that not all of the uncoated and coated losses were used to estimate coating loss. The reasons for excluding some of these measurements are listed below:

- Discs exhibit degenerate mode pairs [144] which resonate along different (orthogonal) axes with the same mode shape, as shown in Figure 3-38. A small frequency shift is observed between the pairs due to the imperfect symmetry of the discs. Theoretically, each resonant mode in a degenerate pair should have the same loss [144]. However due to the suspension technique used to support the disc, one mode in each pair will always exhibit higher loss due to the suspension wires adding extra damping into the system. The lowest loss mode is unaffected by the wires as the position where the wires touch the disc are defined as nodes. Modes 1 and 2, 3 and 4, 5 and 6, 7 and 8 and 9 and 10 are degenerate mode pairs with modes 2, 4, 8 and 10 being the higher loss members of their respective pairs, as shown in Figure 3-35 and 3-36. The higher loss mode from each pair was therefore excluded from the analysis.
- Modes 5 and 6 were excluded from calculating coating loss because only one of these modes was measured post coating deposition. Since the coated loss measured was

higher than both modes 5 and 6's uncoated losses it was not possible to identify which mode this coated loss belonged to.

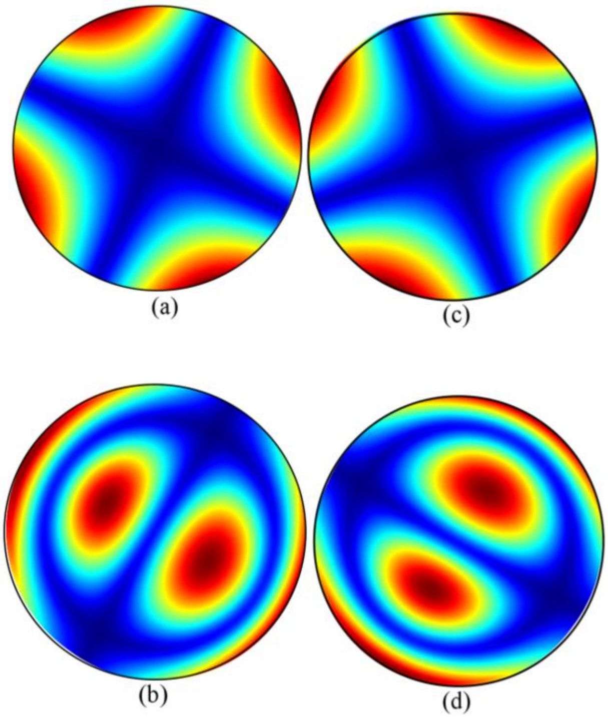

Figure 3-38: Disc's degenerate mode pairs: (a,c) butterfly modes and (b.d) drum modes.

3.5.1.5 Bulk and shear losses of the coating

The tantala coating's bulk and shear losses were calculated using the minimisation method to be $\phi_B = \left((1.28 \pm 0.22) \times 10^{-3}\right)$ and $\phi_S = \left((7.8 \pm 0.5) \times 10^{-4}\right)$. These values were determined to be significantly different from one another. The errors shown are one standard deviation total errors (combination of systematic and random). Figure 3-39 shows the measured and fitted values for the mode-dependent coating loss. While a similar magnitude of loss is observed between the fitted and experimental values, it can be seen that all but one of the fitted points do not match the experimental values within error. This highlights the same problem discussed previously of fitting to such a small number of data points. The errors plotted in Figure 3-39 represent the one standard deviation random coating loss error and they assume that the lowest coated and uncoated losses measured for the disc are in fact the discs intrinsic loss and no excess loss has occurred from the setup. Usually the lowest measured loss is accepted as the best value, as it's much easier to measure too high a loss value than too low a loss value. Repeat loss measurements are therefore useful as you need several measurements to know which one is lowest. Measuring the disc multiple times to identify this convergence in modal loss between suspensions is therefore essential since there is no method of quantifying this error. In Figure 3-35 and 3-36, it can be seen that the uncoated losses show better convergence in values than the coated losses. More suspensions of the coated disc would have been ideal, however the disc cracked after the second measurement making this impossible to do. This uncertainty in the coated losses coupled with the limited number of data points being fitted to could explain the differences in fitted and experimental values. For example, the shape of the fitted losses is dominated by mode 9. Figure 3-36 does not show absolute convergence for mode 9's coated losses. If it's coated loss were less, its coating loss would be less and therefore the fit could potentially be improved. One last point to note is that the four coating losses determined were all measured on butterfly modes. Therefore the modes only showed small variations in their bulk and shear energy ratios. Theoretically, it would be better to fit to modes which have large variations in bulk and shear energies as it makes the results less sensitive to errors associated with the experiment.

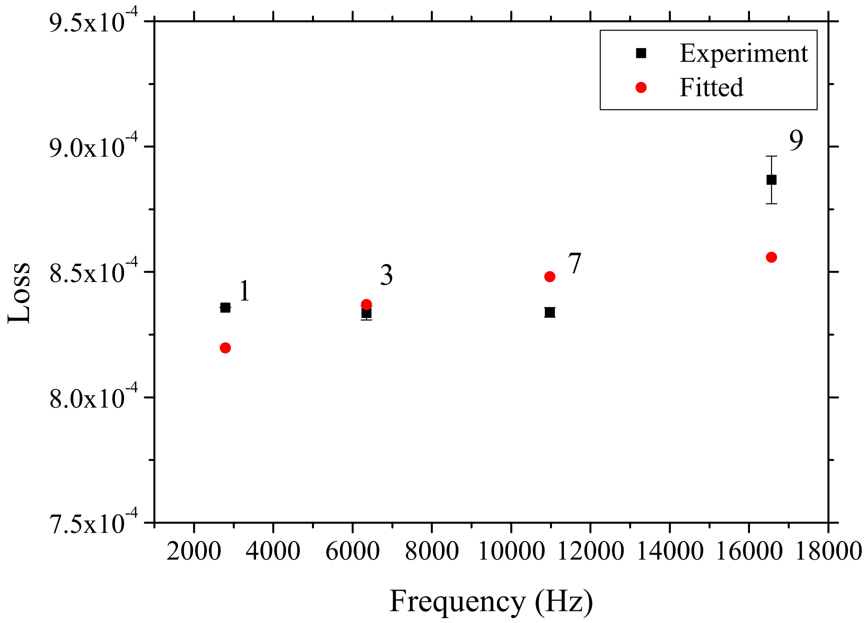

Figure 3-39: Experimental and fitted tantala coating losses. The fitted coating losses were calculated using the coating's bulk and shear losses that were determined using the minimisation method.

3.5.2 Determining the bulk and shear losses of the tantala coating deposited upon a cantilever

3.5.2.1 Sample

The same 500×10^{-9} m RLVIP tantala coating was deposited upon a silica cantilever of thickness 159×10^{-6} m [139]. The thickness of this substrate was determined using its uncoated bending mode frequencies.

3.5.2.2 Uncoated and coated losses

The uncoated cantilever was clamped twice and its modal losses measured after each clamping. As shown in Figure 3-40, there was good agreement between the losses measured for each clamp. Three modes were measured within the frequency range of 14300 – 14600 Hz. FEA modelling predicts only two modes to exist within this frequency range (a torsional mode and a higher frequency bending mode). It is suspected that this extra mode is caused by one or both of the modes being re-excited in some way and remeasured at a slightly different frequency (an observation that has been made in the past with other samples). It is also unknown as to whether this effect results in any coupling between the different modal losses similar in frequency. As a result of not being able to definitively identify which mode is bending and which is torsional, these modes have been labelled in Figure 3-40 as “not known”.

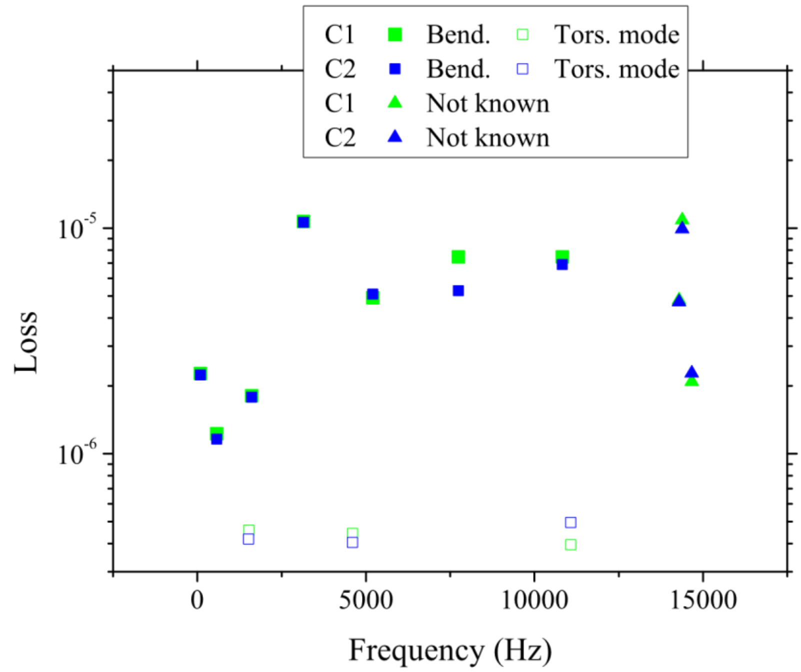

Figure 3-40: Uncoated losses of a silica cantilever. The losses were measured for two different clamps (clamp 1 (C1) and clamp 2 (C2)).

The coated silica cantilever was clamped eight times and its losses measured after each clamping. The lowest measured losses of the coated cantilever are shown in Figure 3-41. All but two of the coated cantilever's mode shapes could be easily identified as their frequencies shifted by less than 20 Hz from their uncoated frequencies. The 7804 Hz coated cantilever mode could not be definitively identified as being a bending or a torsional mode. This was because FEA modelling predicts both a bending and torsional mode to exist within 50 Hz of this frequency (and unlike for the uncoated cantilever mode at this frequency, its loss cannot be used for guidance as to whether this mode is a bending or torsional mode). It was also not possible to identify whether the measured mode of the coated cantilever at 14648 Hz was bending or torsional, as both theoretically exist at similar frequencies and only one mode was measured.

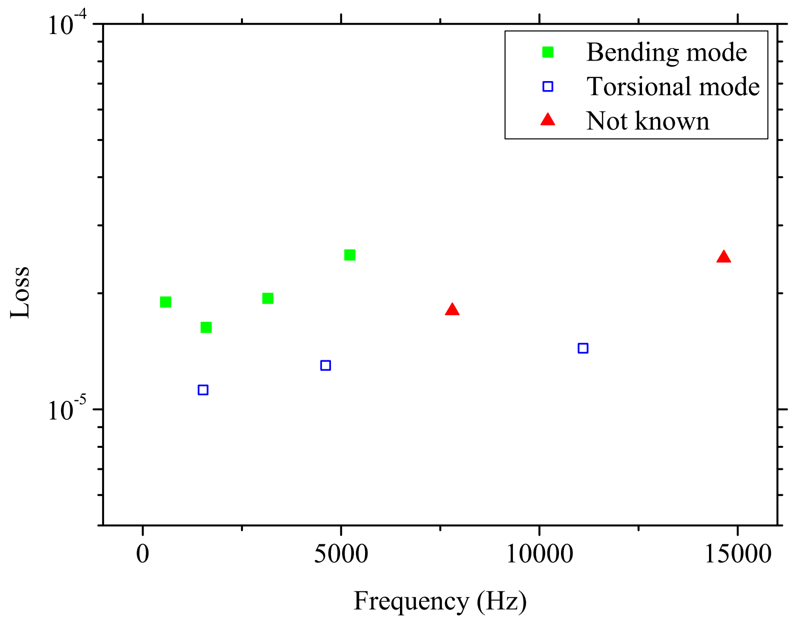

Figure 3-41: Lowest losses of the coated silica cantilever.

3.5.2.3 Bulk and shear losses of the coating

The mode-dependent coating losses were determined using the uncoated and coated losses (shown in Figures 3-40 and 3-41). Using these mode dependent coating losses, the minimisation method was then used to calculate the coating's bulk and shear losses. Figure 3-42 shows the measured mode dependent coating losses and the fitted values obtained from the minimisation method. Similarly to Figures 3-27, 3-28 and 3-29, two standard deviation error bars are plotted for each mode's coating loss. The larger error bar is a combination of the coating loss's random error and systematic error whilst the smaller error bar represents only the coating loss's random error. For most points, these error bars are very similar in size.

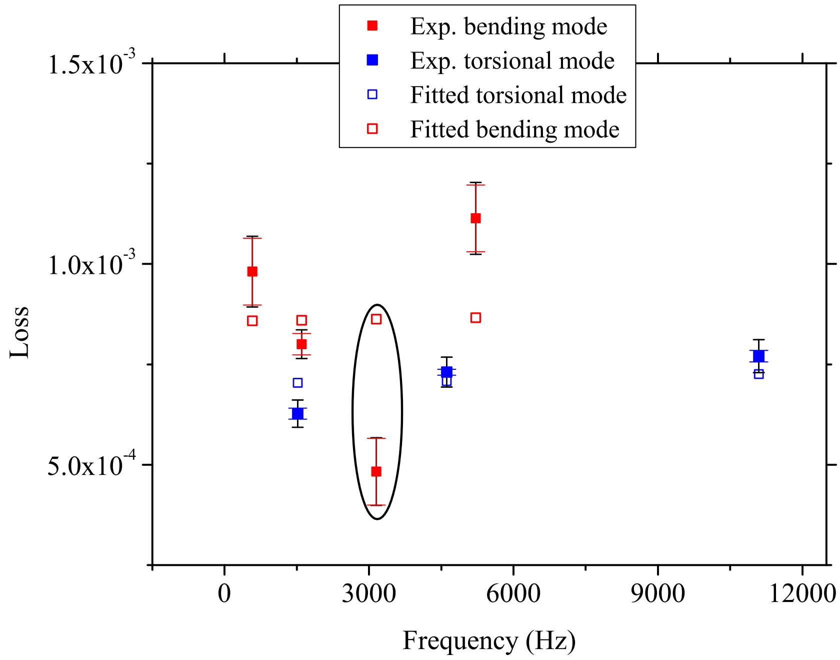

Figure 3-42: Measured and fitted coating loss values. The black error bars represent the one standard deviation total error (systematic and random error) in coating loss. The blue and red error bars represent only the one standard deviation random uncertainty in coating loss. The experimental and fitted losses of the 3100 Hz mode have been grouped together as they are explicitly referred to in the text.

The bulk and shear losses of the coating were determined to be $(1.4 \pm 0.2) \times 10^{-3}$ and $(7.0 \pm 0.5) \times 10^{-4}$ respectively, where the errors represent the one standard deviation total error (combination of systematic and random). However, it can be seen in Figure 3-42 that the quality of the fit used to determine these losses is suboptimum as a large difference exists between the experimental and fitted bending mode at approximately 3100 Hz (these modes are circled in the graph). This difference arises from the unusually high uncoated loss measured for this mode, which is likely to be an effect of excess loss. Evidence in support of this suggestion can be observed in Figure 3-26 which shows an upper and lower bound for the uncoated loss of 5 similar silica cantilevers. It can be observed that all of the measured uncoated silica losses (plotted in Figure 3-40) which were used to calculate coating loss (Figure 3-42) fall within these upper and lower bounds of uncoated loss (Figure 3-26), with the exception of the 3100 Hz mode. Using an uncoated loss for this mode which is equal to the average of the upper and lower bounds and repeating the minimisation method, substantially improves the quality of the fit. This improvement in fit is shown in Figure 3-43 which displays the experimental and fitted losses when this new uncoated loss is used for the 3100 Hz mode.

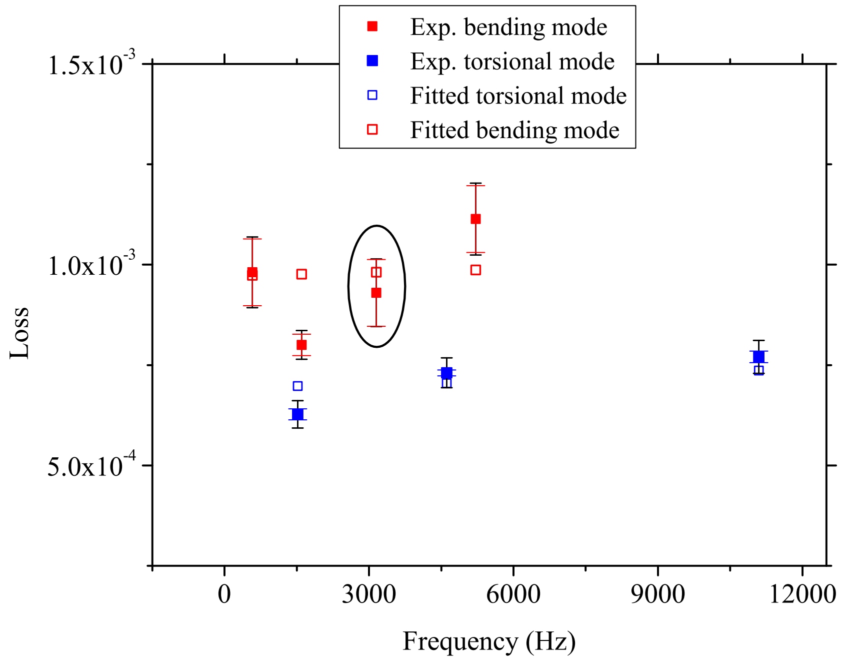

Figure 3-43: Measured and fitted coating loss values (using a new value of uncoated loss for the 3100 Hz mode). The black error bars represent the one standard deviation total error (systematic and random error) in coating loss. The blue and red error bars represent only the one standard deviation random uncertainty in coating loss. The experimental and fitted losses of the 3100 Hz mode have been grouped together to make it easier for the reader to identify them when reading the text.

The bulk and shear losses determined using this fit were $(1.90 \pm 0.30) \times 10^{-3}$ and $(6.7 \pm 0.5) \times 10^{-4}$ respectively, where the errors represent the one standard deviation total error (systematic and random). Due to the greater accuracy of this fit, there is more confidence in these loses being a closer representation of the true values of the sample.

3.5.3 Comparing the tantala coating's bulk and shear losses for the disc and cantilever substrates

The bulk and shear losses of the tantala coating deposited upon the cantilever were determined to be $(1.90 \pm 0.30) \times 10^{-3}$ and $(6.7 \pm 0.5) \times 10^{-4}$ respectively. Comparing these tantala coating bulk and shear losses to those calculated on the disc substrate ($(1.28 \pm 0.22) \times 10^{-3}$ bulk loss and $(7.8 \pm 0.5) \times 10^{-4}$ shear loss) it can be observed that there is good qualitative agreement between the results, with the bulk loss being significantly higher than the shear loss in both cases. However, the results from the two sample sets do not agree to within error. One possible explanation for this difference is that the uncoated and coated losses measured (of both the disc and cantilever) were not the true values of the sample and that with more re-clamping and more loss measuring, the uncoated and coated losses could change, thus affecting the coating losses and the bulk and shear losses. More possible causes for this discrepancy between the disc and cantilever's bulk and shear losses could be due to the limited number of modes measured on the disc and that all of the modes used in the disc fitting process being butterfly resonances (which makes the analysis more sensitive to experimental errors).

3.6 Validity of results

3.6.1 Frequency dependence of coating loss

Both the simultaneous equation method and the minimisation method rely upon the assumption that any frequency dependence in the mode dependent coating losses can be explained by the frequency dependent bulk/ shear energy ratio.

In the literature a frequency dependent coating loss has been calculated for three different coatings: aluminium oxide, tantala and silica which were all deposited upon silica discs of varying geometries [145]. The aluminium oxide coating showed a frequency dependence of approximately 1 % across 10 kHz, the tantala coating 4 % across 10 kHz and the silica coating 10 % across 10 kHz [145]. To obtain these results, the loss of the coatings were determined at different mode frequencies. The modal losses of each coating were then regressed only against frequency and did not account for the modes bulk and shear energies changing with frequency (for a cantilever this can be observed in Figure 3-20).

To investigate whether this frequency dependent coating loss could be caused by the bulk and shear energies of the coated disc's modes changing, a silica disc coated with silica was designed in COMSOL and the bulk and shear energy ratios of the coating calculated. It was observed for the disc's butterfly modes that the percentage of coating energy stored in bulk energy increased with frequency. Across 10 kHz, the percentage of coating energy stored in bulk energy increased by 12%. For a disc's drum modes, it was observed that the percentage of coating energy stored in bulk motion decreased with frequency. Across 10 kHz, the percentage of coating energy stored in bulk energy decreased by 18%. The frequency dependent coating loss identified in [145] can therefore be potentially explained by the modes of a coated disc exhibiting different ratios of bulk and shear energy at different frequencies.

It can therefore be concluded that the simultaneous equation method, the minimisation method and the bulk and shear losses determined using these methods remain valid.

3.6.2 Stress and curvature

Stress and curvature are known to develop in coated substrates during (intrinsic stress) and after the coating deposition process (thermal stress) [146]. In this section, the effect of stress and curvature upon the coating/substrate energy ratio is investigated to identify whether these parameters can significantly affect the energy ratio (and therefore the calculated bulk and shear losses).

3.6.2.1 Coated cantilevers

To observe the effect of stress and curvature upon the energies stored in the coating and substrate, an FEA model replicating the stress pattern and curvature caused by a coating was constructed. A silicon cantilever coated with amorphous silicon was used for this study. The coating was 1×10^{-6} m thick and the length, width and thickness of the cantilever were 34×10^{-3} m, 5×10^{-3} m and 63.5×10^{-6} m respectively. The mode frequencies, vertical displacement (Figure 3-44) and direction of stress within the coating and the cantilever were all measured experimentally.

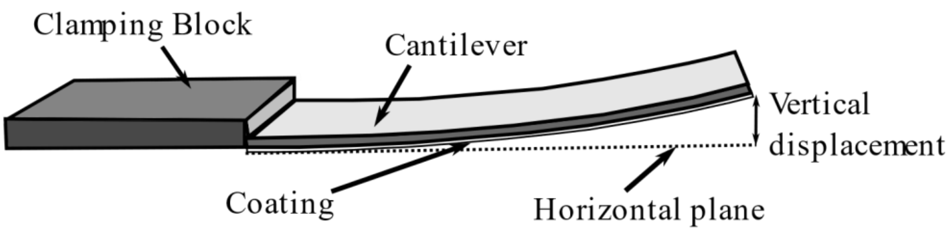

Figure 3-44: Vertical displacement of a cantilever as a result of depositing a coating on its surface.

Multiple approaches in ANSYS were tested to stress and curve the coated cantilever. The ANSYS model which most accurately replicated the measured mode frequencies, vertical displacement and direction of stress, applied a thermal stress to the cantilever and coating within a *static structural analysis*. The thermal stress was applied by defining the coating and substrate thermal expansion coefficients (3×10^{-9} K^{-1} and 4.6×10^{-5} K^{-1} respectively) along with a thermal temperature condition which was used to replicate the cooling effect from deposition temperature (100 °C) to room temperature (22 °C). When the ANSYS model was ran, the coated cantilever would cool from an initial temperature of 100 °C to 22 °C which would result in a thermal stress being developed in the coated cantilever (as a result in the mismatch of the coating and substrate's thermal expansion coefficients). The coefficients of thermal expansion (CoTE) of the substrate and coating were chosen such that they enabled the FEA model to: bend and stress the coated cantilever in the same directions as measured in the lab (Figure 3-45); introduce a similar magnitude of vertical displacement (the cantilever in the lab was measured to have $(3 \pm 0.5) \times 10^{-3}$ m vertical displacement whilst the model predicted 3.23×10^{-3} m) and to reproduce comparable modal frequencies and shapes to those measured (Table 3-3).

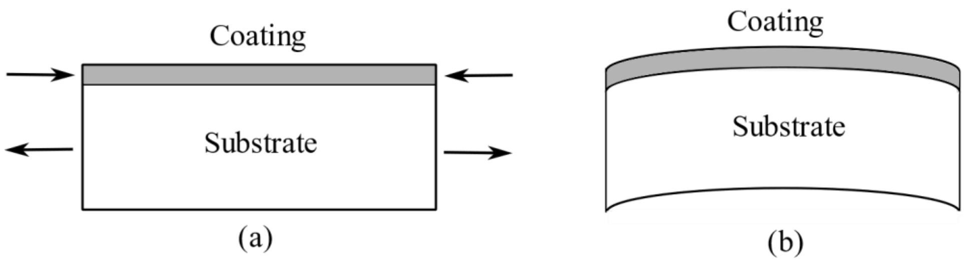

Figure 3-45: *(a) direction of stress measured in the laboratory for the amorphous silicon coating and silicon substrate (compressive stress in the coating and tensile in the substrate). (b) deformation caused as a result of the stress in the coating and substrate.*

Exp. freq. (Hz)	Unstressed ANSYS freq. (Hz)	Stressed ANSYS freq. (Hz)	Diff. between exp. and unstressed freq. (%)	Diff. between exp. and stressed freq. (%)
	76 (b)	77 (b)	N/A	N/A
465 (b)	479 (b)	480 (b)	-3.0	-3.3
715 (t)	921 (t)	725 (t)	-28.8	-1.4
1338 (b)	1344 (b)	1352 (b)	-0.4	-1.1
2159 (t)	2800 (t)	2066 (t)	-29.7	4.3
2635 (b)	2633 (b)	2652 (b)	0.1	-0.6
4147 (t)	4782 (t)	4166 (t)	-15.3	-0.5
4362 (b)	4352 (b)	4379 (b)	0.2	-0.4
	5828 (l)	6237 (l)	N/A	N/A
6516 (b)	6502 (b)	6529 (b)	0.2	-0.2
7567 (t)	6927 (t)	8147 (t)	8.5	-7.7
9095 (b)	9082 (b)	9102 (b)	0.1	-0.1
9174 (t)	9289 (t)	9170 (t)	-1.3	0
11438 (t)	11916 (t)	11412 (t)	-4.2	0.2
12096 (b)	12092 (b)	12100 (b)	0	0
	14845 (t)	14261 (t)	N/A	N//A
15519 (b)	15532 (b)	15524 (b)	-0.1	0
17344 (t)	18108 (t)	17477 (t)	-4.4	-0.8
	19400 (b)	19370 (b)	N/A	N/A
	21370 (t)	21063 (t)	N/A	N/A

Table 3-3: Experimental and computational mode frequencies for a coated silicon cantilever. The table includes computational mode frequencies for two ANSYS models. One which used the thermal condition to pre-stress a modal analysis and the other which did not. The letters b, t and l describe the type of mode as being either bending, torsional or lateral respectively. The cantilevers substrate dimensions were: length 34×10^{-3} *m, width* 5×10^{-3} *m and thickness* 63.4×10^{-6} *m. The coating was made of amorphous silicon and had a thickness of* 1×10^{-6} *m. Some modes were unable to be measured experimentally. These are represented in the table by a blank cell.*

Figure 3-45 shows the direction of the stress and curvature of the coated silicon cantilever. Table 3-3 compares the measured mode frequencies to those obtained from the FEA model. The stressed and curved FEA model replicates the measured mode frequencies more closely than the unstressed model. The addition of stress and curvature to the FEA model has a small effect upon bending mode frequencies (less than 2% change) but a large effect upon torsional mode frequencies (up to 35%). Due to the similarity in results between the stressed ANSYS model and the experimental values, it seems reasonable to assume that the stress and curvature in this FEA model approximates that present in the coated cantilever as well.

The magnitude of the stress and curvature of the coated cantilever can be increased by varying the CoTE values. However, it should be noted that the relative magnitudes of the coating and substrate CoTE should be kept the same (substrate CoTE > coating CoTE) as this ensures the stress pattern remains unchanged (uniformly compressed coating and tensile substrate). By running multiple FEA models with varied thermal expansion coefficients and extracting the coating and substrate elastic strain energies for each mode, the effect of stress and curvature upon the coated cantilever's energy ratios can be determined. Figure 3-46 (a) shows the effect upon bending mode energy ratios whilst (b) shows the effect upon torsional and lateral modes. It should be noted that compressive stresses of between 0 and 300×10^6 Pa are commonly measured in coatings that are used in mechanical loss experiments [54].

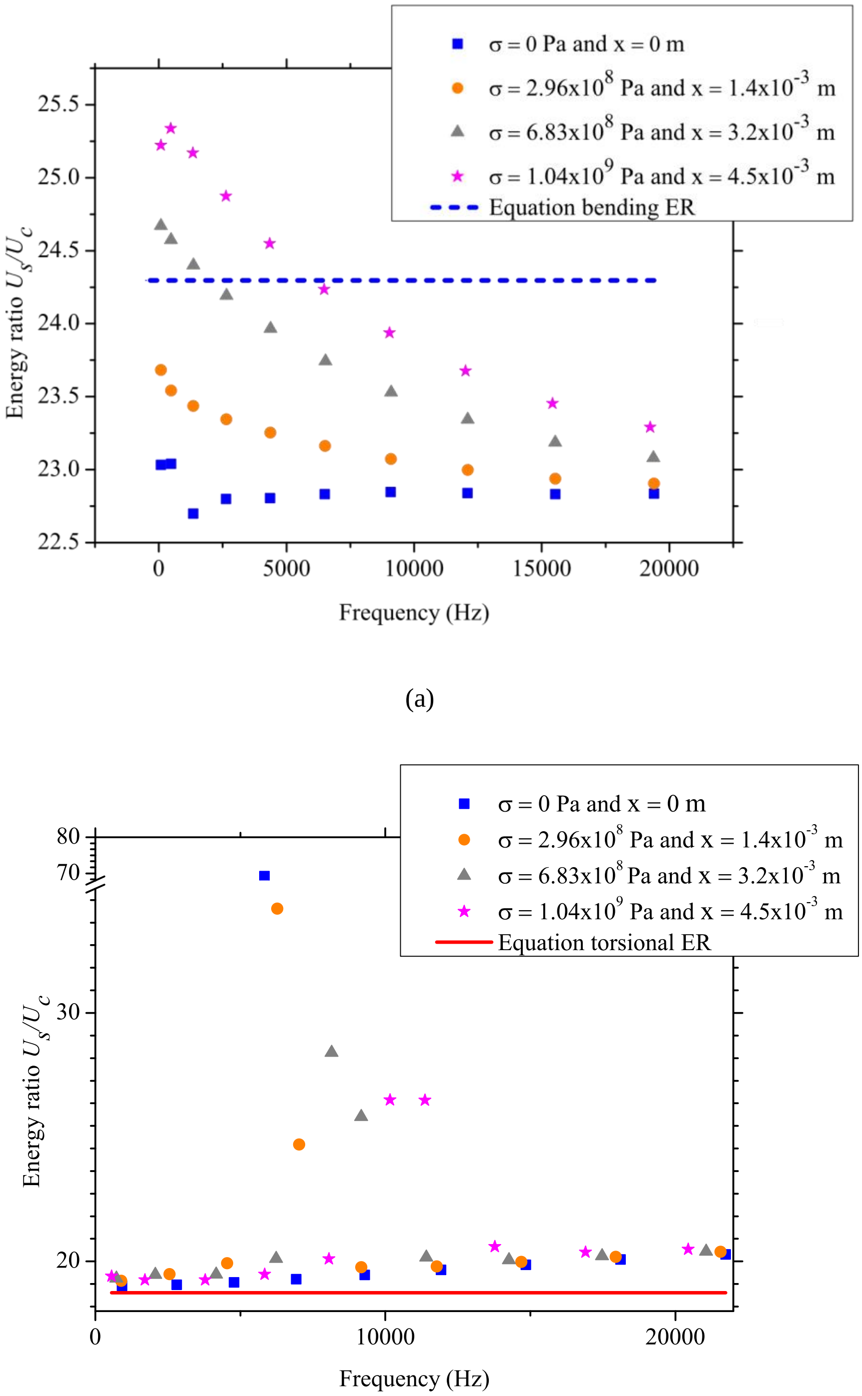

(a)

(b)

Figure 3-46: Effect of stress and curvature upon FEA modelled energy ratios (ER) for a coated silicon cantilever. (a) shows the effect upon bending modes whilst (b) shows the effect upon torsional and lateral modes. The energy ratios predicted by Equations (3.18) and (3.20) are also shown. The elastic energy ratios of four differently stressed FEA models are plotted. The compressive stress of the coating (σ) and vertical displacement (x) of the cantilevers are quoted in the legend. These stresses were measured acting along the length of the coating. The cantilever's substrate dimensions were: length 34×10^{-3} *m, width* 5×10^{-3} *m and thickness* 63.4×10^{-6} *m. The coating was made of amorphous silicon and had a thickness of* 1×10^{-6} *m.*

The bending mode energy ratio can be seen to increase as the stress and curvature increases. However, as the frequency increases, the sensitivity of the bending mode energy ratio to stress and curvature decreases. In contrast, the stresses and curvatures investigated have little effect upon the majority of the torsional mode energy ratios. The large outlier in Figure 3-46 (b) with an energy ratio of approximately 70 in the unstressed and non-curved model is a lateral mode and should not be expected to conform to the torsional mode trend. However, it can also be observed from Figure 3-46 (b) that the torsional modes near the lateral mode in frequency are more sensitive to the effects of stress and curvature. This could be due to coupling of these torsional modes with lateral motion. It can be concluded from Figure 3-46 that both stress and curvature have a significant effect upon a coated cantilever's substrate/coating energy ratio and that this effect becomes larger at higher values of stress and curvature. Additional FEA models of coated cantilevers with different geometries and material properties corroborated these results. Table 3-4 lists the other models investigated.

Body	**Material**	**Length** ($\times$ $\mathbf{10^{-3}}$ **m)**	**Width** ($\times$ $\mathbf{10^{-3}}$ **m)**	**Thickness** ($\times$ $\mathbf{10^{-6}}$ **m)**
Substrate	cSi	34	5	66.1
Coating	aSi	34	5	8.24
Substrate	cSi	34	5	74.9
Coating	aSi	34	5	1
Substrate	Silica	42	5	100
Coating	Tantala	42	5	1

Table 3-4: List of the ANSYS models that were investigated to identify the effect of stress and curvature upon energy ratios. Each substrate and coating pair was an ANSYS model. The table states the different materials of the coatings and substrates used in each of the ANSYS models as well as their geometries. Stresses of similar magnitude to those listed in Figure 3-45 were applied to these models. cSi and aSi refer to crystalline silicon and amorphous silicon materials.

The coated cantilevers investigated in Sections 3.4 and 3.5 were not particularly curved as they all exhibited $\leq$ 1 mm vertical deflection. The stresses of these coatings were unmeasured. The most curved coated cantilever investigated (in this chapter) was the silicon cantilever coated with ECR IBS amorphous silicon (vertical deflection ≈1 mm). Repeating the minimisation method using stressed energy ratios, the best fit bulk and shear loss values were determined to be 3.5×10^{-4} and 1.5×10^{-4} respectively. These values fall within the random error of the bulk and shear losses determined earlier in Section 3.4.1.2 (bulk loss of $(3.4 \pm 0.7) \times 10^{-4}$ and shear loss of $(1.4 \pm 0.2) \times 10^{-4}$). It can therefore be concluded that it is unlikely that curvature and stress significantly affect the bulk and shear losses presented in this chapter. However, the effect of curvature and stress on the energy ratios of a coated cantilever should be considered when the coated cantilever is more highly curved and stressed.

3.6.2.2 Coated disc

When a disc is coated, there is no obvious measurable bending of the coated substrate and no large shift in resonant mode frequencies observed, as shown in Table 3-5.

Tantala coating on silica disc		**Silica coating on silica disc**	
Uncoated freq. (Hz)	**Coated freq. (Hz)**	**Uncoated freq. (Hz)**	**Coated freq. (Hz)**
2789	2790	2808	2779
2813	2813	2832	2834
6351	6355		4240
6361	6365	6395	6395
9734	9742	6405	6403
9738			9802
10969	10973		9806
10980	10975	11039	11038
16559	16564	11050	11039
16564	16573	16661	16662
		16670	16669
		16856	16845
		16848	16851

Table 3-5: Measured resonant mode frequencies for two uncoated silica discs. Table 3-5 also lists the mode frequencies of the discs when they have been coated with tantala and silica.

An ANSYS model of a disc which includes stress must therefore replicate these experimental observations. The method used to incorporate stress in the ANSYS model of a cantilever can not be used for a disc. This is because a fixed constraint is not used to determine a disc's resonant modes (whilst it was used for the cantilever to model its clamped condition) and a fixed constraint is a requirement of the *static structural* interface which employs thermal properties to stress the cantilever. A new method therefore had to be developed.

To stress the coated disc, a modal interface was used along with an 'inistate' command which applied a uniform radial stress to either the coating, substrate or both. The stress was varied until the modelled resonant mode frequencies matched the measurements. No deformation of the disc was observed in this analysis. The energy ratios for each mode were calculated and the mesh checked for convergence. The energy ratios for differently stressed ANSYS disc models for a 500×10^{-9} m tantala coating deposited upon a silica disc (with dimensions 2.57×10^{-3} m thickness and 38.07×10^{-3} m radius) are shown in Figure 3-47 where it can be seen that the different stress patterns have an insignificant effect upon the substrate/coating ratio.

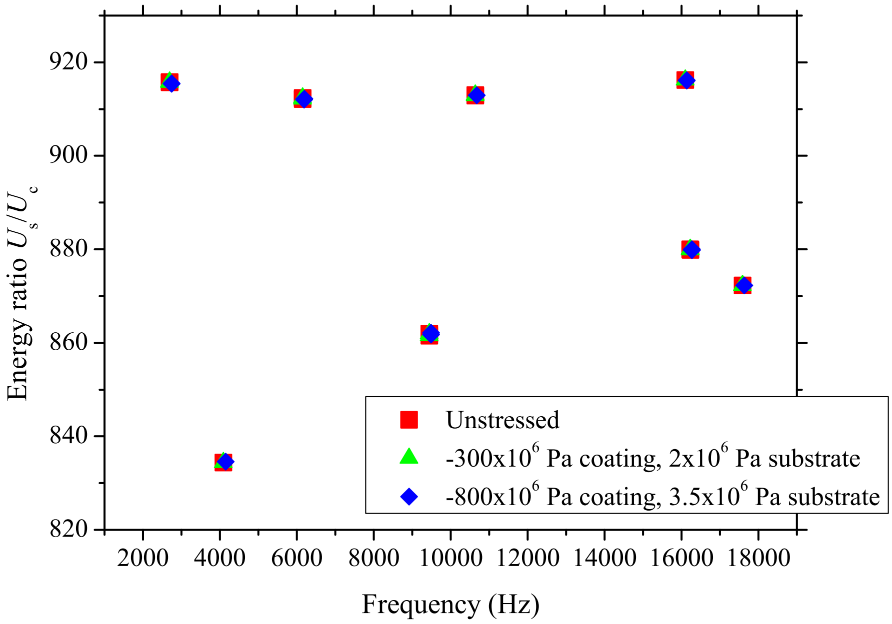

Figure 3-47: Ratio of energy stored in the substrate with respect to that stored in the coating for a 500×10^{-9} *m tantala coating deposited upon a silica disc (with dimensions:* 2.57×10^{-3} *m thickness and* 38.07×10^{-3} *m radius).*

Since stress has no effect upon the substrate/coating energy ratio, stress therefore has no effect upon the losses of a coating deposited upon a disc which remains uncurved after deposition.

In Section 3.6.2.1, coated cantilever energy ratios were shown to be affected by curvature and stress. In this section, a coated disc's energy ratios were shown to be unaffected by stress. It is therefore hypothesised that curvature causes coated cantilever energy ratios to change in value.

3.6.3 Effect of stress and curvature upon bulk and shear energy ratios

It was unknown how to define Equations (3.2) and (3.3) within ANSYS to determine bulk and shear energies.

Attempts were therefore made at investigating the combined effect of stress and curvature upon a cantilever's bulk and shear energy ratios using COMSOL. However, due to a problem with the software, confirmed by COMSOL [147], it was not possible to investigate the effects of stress upon a cantilever's bulk and shear energy ratio as it led to unphysical results (in some cases negative energy ratios). However, it was possible to identify the relationship between a cantilever's vertical displacement (Figure 3-44) and its bulk and shear energy ratios.

Coated silicon cantilevers of varying vertical displacements were designed in a *solid mechanics interface* and an *eigenvalue analysis* solved (which determined the mode frequencies of the coated cantilever). Using Equations (3.2) and (3.3) the bulk and shear energies of the coating were calculated using COMSOL. The silicon cantilever was 60×10^{-6} m thick, 5×10^{-3} m wide and 34×10^{-3} m long. The coating was made from amorphous silicon and was 1×10^{-6} m thick. The results are presented in Figure 3-48.

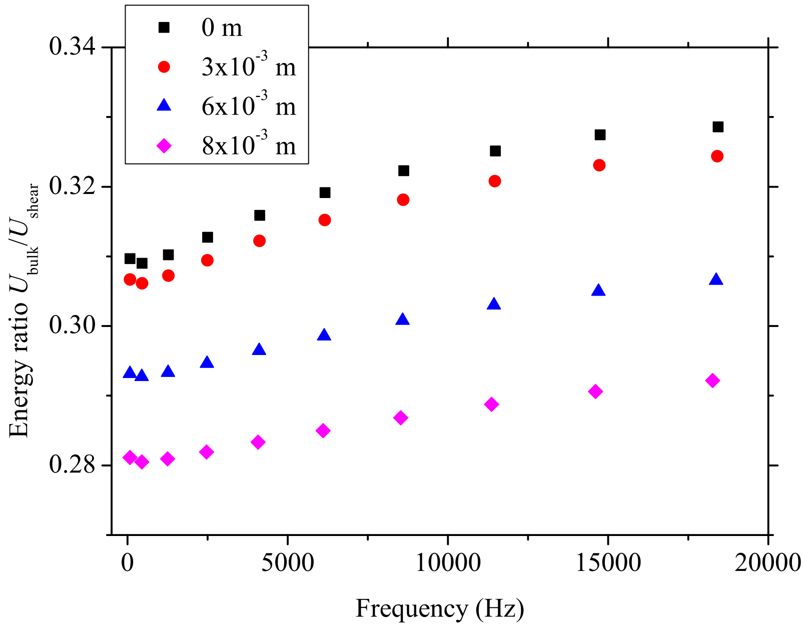

(a)

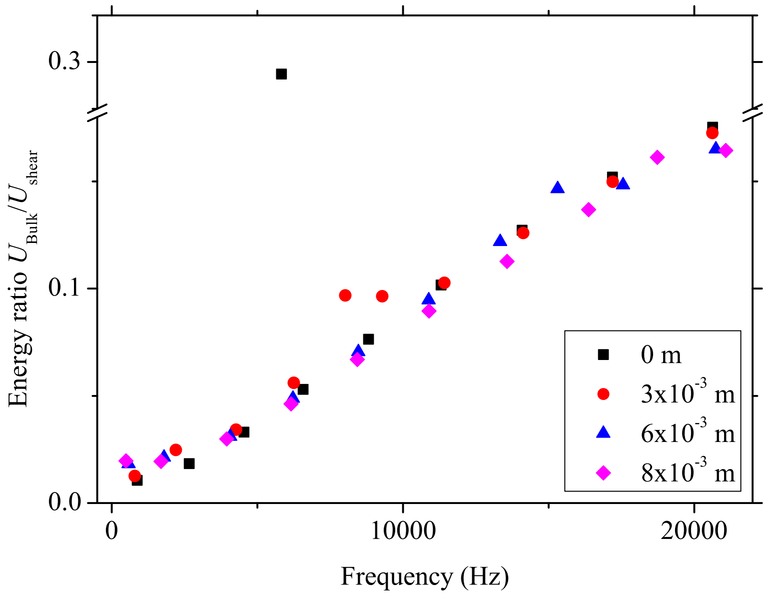

(b)

Figure 3-48: Bulk/shear energy ratio for varyingly curved coated cantilevers. The coating was 1×10^{-6} *m thick and made from amorphous silicon, whilst the cantilever was* 5×10^{-3} *m wide,* 34×10^{-3} *m long,* 60×10^{-6} *m thick and made from silicon. (a) Bulk/shear energy ratio for the bending modes whilst (b) shows the torsional and lateral modes. The vertical displacements of the coated cantilevers are shown in the legends.*

In Figure 3-48 (a) it can be observed that as the vertical displacement of the cantilever increases, less of the coating's elastic strain energy for a bending mode is stored in bulk motion compared to shear motion. This effect can be seen to increase as frequency increases.

It is difficult to identify any clear trend in Figure 3-48 (b). The pure lateral mode in the uncurved cantilever (approximately 5000 Hz with a bulk/ shear ratio of 0.3) acts more like a torsional mode as curvature increases (as its bulk/shear energy ratio begins to more closely approximate the torsional mode values). Furthermore, this lateral mode affects the pure torsional modes energy ratios' near it in frequency distorting any trend further (similar effect was observed in Figure 3-46 (b)). Finally, this picture is further complicated by the curvature affecting torsional mode frequency. However, it can be said that vertical displacement can be seen to affect the bulk and shear energy ratios for both bending and torsional modes.

In order to maximise the accuracy in calculating bulk and shear losses, the effect of a cantilever's vertical displacement upon bulk and shear energies should be accounted for.

The effect of a cantilever's vertical displacement upon the coating's bulk and shear energies is not expected to have any significant effect upon the bulk and shear losses listed in this chapter as the vertical displacement of the cantilevers was small (< 1 mm). By accounting for the effect of vertical displacement and recalculating the bulk and shear losses of the amorphous silicon coated silicon cantilever (most vertically displaced cantilever investigated in this chapter), the losses of the coating were still determined to be within error of the original losses (which did not account for vertical displacement effects).

It should be noted that in reality when a cantilever is coated at an elevated temperature and is allowed to cool to room temperature, the coated cantilever curves both along its length (called vertical displacement and shown in Figure 3-44) and its width (called width displacement and shown in Figure 4-33). Due to the dimensions of the cantilever, the vertical displacement is much larger than the width displacement. In Section 3.6.2.1, ANSYS was used to curve a coated cantilever before the stressed and curved coating/substrate energy ratios were calculated. This method used in ANSYS to curve and stress the coated cantilever, curved the coated cantilever both along its length and width. It should be emphasised that the COMSOL bulk/shear energy ratios shown in Figure 3-48 only investigated the effect of a coated cantilever's vertical displacement upon the bulk/shear energy ratio. It was not possible to determine the effect of width displacement (either individually or in conjunction with vertical displacement) upon a coating's bulk/shear energy ratio due to meshing problems in the *static structural* COMSOL model (inverted mesh elements within the coating caused models not to solve and a failure of solutions to converge i.e. mesh convergence issues). Whilst the cause of this FEA problem is unknown, it is suspected that it could be due to a combination of the small coating dimensions and its width deformations. The effect of width displacement upon bulk and shear energies is an area of further work that should be investigated in the future.

No investigation was conducted into the effect of curvature upon a coated disc's bulk and shear energy ratios, as coated discs do not exhibit any noticeable curvature when they are coated in the lab (see Section 3.6.2.2).

3.6.4 Effect of Fejer coating thermoelastic loss

Fejer coating thermoelastic loss (Fejer loss) [109] describes a loss mechanism of a coated substrate which is caused by the mismatch in the substrate's and coating's material properties. Unlike coating thermoelastic loss (Section 2.4.1) which only accounts for random temperature fluctuations in the coating, Fejer coating thermoelastic loss investigates random temperature fluctuations in the coating and substrate. It is calculated using [109]:

$$\phi_{\text{Fejer}}(\omega) = \frac{2Y_f \alpha_f^2 T}{C_f(1-v_f)}\left(1 - \frac{\alpha_s}{\alpha_f}\left(\frac{Y_s(1-v_f)C_f}{Y_f(1-v_s)C_s}\right)\right)^2 g(\omega)\,, \qquad (3.28)$$

where Y is the Young's modulus, T is the temperature, α is the thermal expansion coefficient, C is the specific heat capacity, v is the Poisson ratio, the subscripts s and f represent the substrate and film respectively and $g(\omega)$ is defined as:

$$g(\omega) = Im\left(-\frac{1}{\sqrt{i\omega\tau_f}}\left(\frac{\sinh\left(\sqrt{i\omega\tau_f}\right)}{\cosh\left(\sqrt{i\omega\tau_f}\right) + \left(\frac{K_f C_f}{K_s C_s}\right)^{\frac{1}{2}}\sinh\left(\sqrt{i\omega\tau_f}\right)}\right)\right), \qquad (3.29)$$

where ω is the angular frequency, τ_f is defined as:

$$\tau_f = \frac{t^2 C_f}{K_f}\,, \qquad (3.30)$$

t is the coating thickness and K is the thermal conductivity. It can be observed in Equation (3.28) that the greater the difference in the coating and substrate material properties, the greater the Fejer loss. Strictly speaking, the effect of Fejer loss should always be considered when the mode dependent mechanical loss is calculated [148]:

$$\phi_{coating} = \frac{U_s}{U_c}(\phi_{coated} - \phi_{uncoated}) - \frac{U_{\text{bulk}}}{U_c}\phi_{\text{Fejer}} \ . \tag{3.31}$$

However, as a result of the Fejer loss being a small effect (particularly at low frequencies), it is often not considered when calculating the loss of a coating [137] and is therefore not expected to effect the bulk and shear losses determined in this chapter.

To test the significance of Fejer loss upon the bulk and shear losses previously presented in this chapter, the bulk and shear losses of the as-deposited amorphous silicon coating and tantala coating were re-calculated accounting for this effect. The bulk and shear losses of the as-deposited amorphous silicon coating when Fejer loss was accounted for was 0 ± 0 and $(1.45 \pm 0.09) \times 10^{-4}$, which is within error of the losses when Fejer loss was ignored 0 ± 0 (bulk loss) and $(1.53 \pm 0.09) \times 10^{-4}$ (shear loss). The bulk and shear losses of the tantala coating when Fejer loss was accounted for was $(1.98 \pm 0.30) \times 10^{-3}$ and $(6.6 \pm 0.5) \times 10^{-4}$, which is within error of the losses when Fejer loss was ignored $(1.90 \pm 0.30) \times 10^{-3}$ (bulk loss) and $(6.7 \pm 0.5) \times 10^{-4}$ (shear loss).

Since neither the Fejer loss associated with the amorphous silicon coating (deposited onto a silica cantilever) or a tantala coating (deposited onto a silica cantilever) had a significant effect upon the bulk and shear losses previously determined (which did not account for Fejer loss), it is not expected that the bulk and shear losses of the amorphous silicon coating (deposited onto a silicon cantilever) will show any significant effect either (as this coating and cantilever have similar material properties and therefore will have a relatively small Fejer loss).

3.7 Conclusion

A method of calculating the bulk and shear mechanical losses of a coating has been developed and tested for various coatings applied to different types of substrate. These are the fundamental losses associated with an amorphous coating material and knowledge of them allows more accurate thermal noise calculations, using a model developed by Hong, to be carried out.

For amorphous silicon deposited by an ECR IBD process, the bulk loss was found to be significantly higher than the shear loss. This is believed to be the first analysis showing that the bulk and shear loss of a coating can indeed be different from each other. In contrast, for amorphous silicon deposited by the RLVIP process, the shear loss dominated the loss of the coating, with the bulk loss being found to be zero. The shear loss of this coating reduced significantly with heat-treatment at both 300 and 500 °C.

The method was also applied to a tantala coating deposited onto a silica cantilever and a silica disk, and the calculated bulk and shear losses were compared. For both samples, the bulk loss was significantly higher than the shear loss. However, the values of bulk and shear loss obtained from the two samples did not agree to within error. This was possibly due to too few modes being measured and to inaccurate loss measurements due to substrate loss effects.

The effect stress and curvature has upon a coated cantilever's substrate/coating energy ratios was investigated. It was observed that the stress and curvature of the coated cantilever significantly affects its modal energy ratios. As the stresses and curvatures of the coated cantilever increases, the effect upon its modal energy ratios increases. The effect stress has upon a coated disc's substrate/coating energy ratios was also investigated. It was observed that stress had no affect upon the coated disc's modal energy ratios. From these results it was hypothesised that the energy ratios of a coated substrate change as a result of curvature alone. It was concluded from this work that in order to determine the most accurate values of a coating's bulk and shear losses, the effect of curvature and stress upon a coated substrate's energy ratios should be considered.

A more accurate method (Hong's method) of determining a coating's thermal noise can be used when the coating's bulk and shear losses are known. Hong's method of calculating thermal noise determines significantly different values from the more widely used method called the parallel and perpendicular approach (a difference of 1.7× and 1.8× for the IBS ECR and RLVIP amorphous silicon coatings respectively). To ensure the thermal noise of future detectors is most accurately calculated, Hong's method should be used to calculate detector thermal noise.

To conclude, Hong's theory of bulk and shear loss enables the gravitational wave community to identify materials which will contribute least to the thermal noise of a detector.

4 The effect of stress and curvature upon thermoelastic loss

4.1 Introduction

Coating thermal noise limits the performance of second-generation gravitational wave detectors at their most sensitive operating frequencies. To minimise the magnitude of a detector's coating thermal noise, detector coatings are made from materials which exhibit low mechanical loss. To identify these low loss materials, mechanical loss experiments (as described in Chapter 3) are used. However, recent research and observations call into question the accuracy of coating losses which have been determined using loss experiments involving silicon cantilever substrates. If some of these coating losses have been determined incorrectly, then it is possible that the magnitude of a gravitational wave detector's coating thermal noise is unnecessarily high and could be reduced by using a previously tested coating.

Silicon cantilevers are frequently used as a substrate in mechanical loss experiments [121, 122]. Coatings are frequently deposited onto cantilevers at elevated temperatures to optimise the coating properties. When the coated cantilevers cool, they become stressed and curved due to thermal expansion. At room temperature the mechanical loss of a silicon cantilever is approximately equal to its thermoelastic loss [136]. If the stress or curvature of a coated cantilever significantly affects the substrate's thermoelastic loss, then the originally measured/calculated uncoated loss of the cantilever, no longer applies as the loss of the substrate in the coated sample. Coating losses determined under these circumstances (using the method outlined in Section 3.3) would therefore be incorrect.

Studies conducted by Kumar and Haque [149], Pan et al [150] and observations made by the author suggest that the stress and/or curvature of a coated silicon cantilever significantly affects the substrate's thermoelastic loss thus invalidating some previously determined coating losses.

Kumar and Haque [149] derived a theoretical model which describes how the thermoelastic loss of a beam can change due to residual stress. Their work showed that as the initial tensile stress of the beam is increased, it's thermoelastic loss decreases. Whilst Kumar and Haque's model shows a clear relationship between stress and thermoelastic loss, the stress pattern they investigated is very simple (axial stress). Furthermore, the beam they investigate does not exhibit any curvature and as a result their model does not represent the pattern of stress or curvature a substrate would experience when a coating has been deposited upon it and allowed to cool.

Pan et al [150] investigated the effect of a coating's stress upon its mechanical loss. They deposited varyingly stressed silicon nitride and amorphous silicon coatings upon silicon cantilevers. They observed that the coated losses measured at bending modes decreased with increased coating stress and increased curvature to such an extent that the coated cantilever losses were less than the uncoated cantilever losses. For torsional modes, they did not see this inversion of coated and uncoated losses. Pan et al speculated from their results that the coated substrate's stress pattern was having an effect upon the substrate's thermoelastic loss, however they were not able to prove this.

A similar observation to Pan et al was made by this author. A highly stressed amorphous silicon coating was deposited upon silicon cantilevers. The coated cantilevers were extremely curved as a result. It was observed that some of the measured bending mode coated losses were lower than the substrate's theoretically calculated values of bending mode uncoated loss (which can be approximately determined using Liftshitz and Roukes theoretical model). These results are shown in Figures 4-1 and 4-2 which compare the measured losses of two coated silicon cantilevers of different substrate thickness with their theoretically calculated bending mode uncoated loss (Liftshitz and Roukes thermoelastic loss)

From Kumar and Haque's study, it is clear that stress can affect thermoelastic loss. However, it is not clear from their work whether the more complicated stress pattern present in a coated cantilever significantly affects the substrate's thermoelastic loss (and therefore is responsible for the lowering of the coated cantilever's loss, as observed in Figures 4-1 and 4-2). It is possible that curvature of the coated cantilever could also affect the coated losses.

In this chapter the effect of a coated cantilever's stress and curvature upon the cantilever's/substrate's thermoelastic loss will be investigated. This chapter introduces thermoelasticity and its most well-established theories before outlining the mathematics of how FEA was able to calculate thermoelastic loss. FEA models which are shown to be able to replicate the theory and experimental data will then be used to investigate the effects of a coated cantilever's stress and curvature upon the cantilever's/substrate's thermoelastic loss. Experiments will then be used to try and replicate the conclusions drawn from the FEA models.

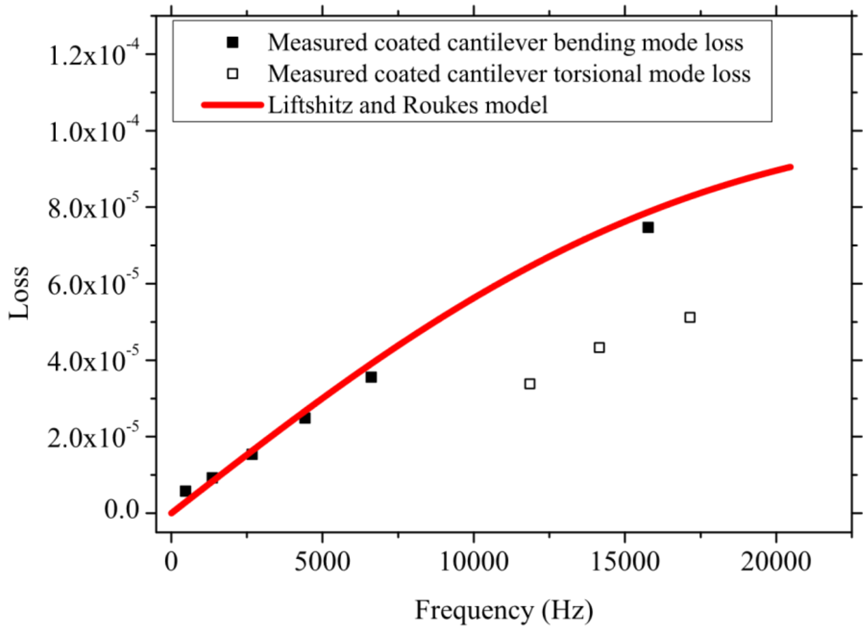

Figure 4-1: Measured losses of a silicon cantilever coated with an amorphous silicon coating. The thermoelastic loss of the uncoated silicon cantilever's bending modes is also shown. This was calculated using Liftshitz and Roukes theoretical model (see Section 4.2.2). No theoretical model exists within the literature to determine the thermoelastic loss of a cantilever's torsional modes. The cantilever investigated was 64×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34×10^{-3} *m long. The coating was* 1×10^{-6} *m thick.*

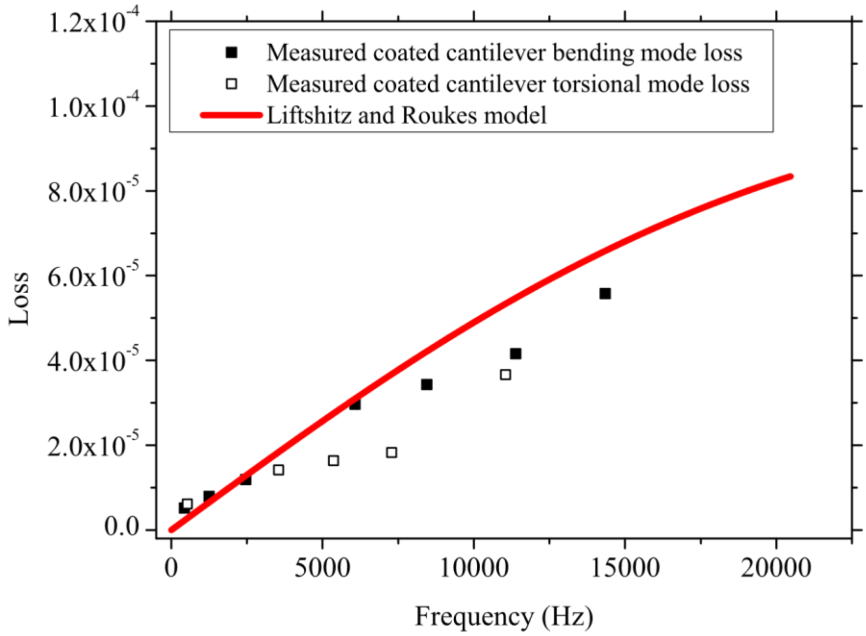

Figure 4-2: Measured losses of a silicon cantilever coated with an amorphous silicon coating. The thermoelastic loss of the uncoated silicon cantilever's bending modes is also shown. This was calculated using Liftshitz and Roukes theoretical model (see Section 4.2.2). No theoretical model exists within the literature to determine the thermoelastic loss of a cantilever's torsional modes. The cantilever investigated was 59×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34×10^{-3} *m long. The coating was* 1×10^{-6} *m thick.*

4.2 Theory of thermoelasticity

Thermoelastic loss is an anelastic relaxation process which describes an irreversible heat flow across a body caused by the coupling of the body's temperature and strain fields via its non-zero thermal expansion coefficient. There are two popular analytical models which are used to approximate the thermoelastic loss of a cantilever's bending modes in the literature. These were derived by Zener [151, 152] and Liftshitz and Roukes [136].

4.2.1 Zener model

Zener investigated the thermoelastic loss of a beam, hypothesising that when an undeformed beam is taken out of its equilibrium and bent into an arc (as shown in Figure 4-3), half of the beam will compress and heat up whilst the rest will expand and cool down. This difference in temperature results in heat flowing across the beam.

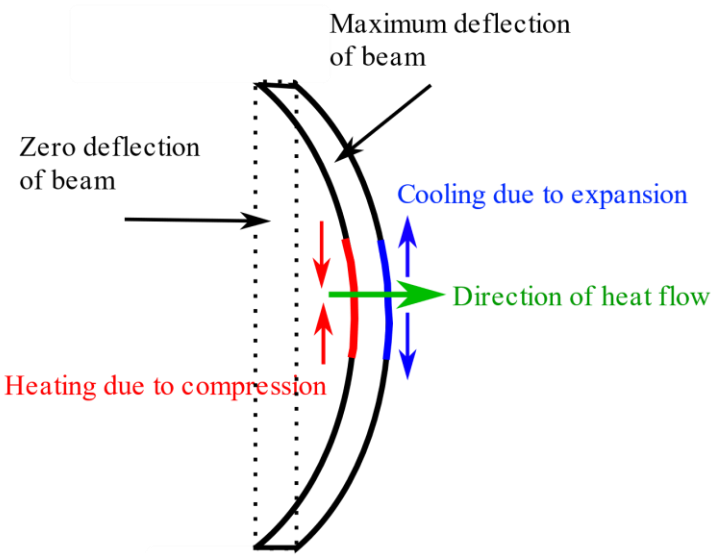

Figure 4-3: A schematic diagram of a bending beam. The initial geometry of the beam is illustrated using dashed lines. When the beam deflects, one side of the beam experiences compression leading to an increase in temperature whilst the other side of the beam experiences expansion resulting in a decrease in temperature. The difference in temperature results in heat flowing across the beam.

When the beam is forced to vibrate and not remain statically bent, the heat flow oscillates from one direction to the other. This coupling between the strain and temperature fields provides an energy dissipation mechanism which allows the beam to relax back to its equilibrium. Zener derived the thermoelastic loss of the beam to be equal to the ratio of the imaginary $I\big(Y_0(\omega)\big)$and real $R\big(Y_0(\omega)\big)$ parts of the complex frequency dependent elastic modulus at constant stress:

$$\phi(\omega) = \frac{I\big(Y_0(\omega)\big)}{R\big(Y_0(w)\big)} = \frac{Y\alpha^2 T}{\rho C}\left(\frac{\omega\tau}{1+\omega^2\tau^2}\right), \tag{4.1}$$

where Y is the Young's modulus, α is the thermal expansion, T is the temperature, ρ is the density, C is the specific heat capacity and ω is the angular frequency. τ is the relaxation time:

$$\tau = \frac{\rho C t^2}{\pi^2 K}, \tag{4.2}$$

where t is the thickness of the beam and K is the thermal conductivity. The relaxation time is associated with the time taken for the heat to cross the beam. According to Zener's expression (Equation 4.1), the thermoelastic loss of a beam as a function of frequency follows a Loretzian distribution whose peak value is dependent upon the beam's relaxation rate (inverse of relaxation time). When a beam oscillates at a frequency much greater than its relaxation rate, the beam has no time to relax, the vibrations are therefore adiabatic and the loss therefore tends to a minimum value. The loss similarly tends to its other minimum value when the frequency of oscillation is much less than the relaxation rate as the body dissipates very little energy and remains effectively in equilibrium. The loss peak of the Lorentzian occurs when the relaxation rate is equal to its angular frequency.

4.2.2 Liftshitz and Roukes model

To obtain the simple closed form expression of thermoelastic loss (Equation (4.1)), Zener limited his converging infinite series solution to its first term. Liftshitz and Roukes believed this approximation was unnecessary and instead derived an exact analytical formula for the thermoelastic loss of a beam using a mathematically more rigorous approach.

Liftshitz and Roukes calculated the thermoelastic loss of a beam by coupling the beams equation of motion and its heat equation together by noting that the strain (u) of a bent beam is due to the stress (σ) acting along the bent beam as well as an expansion caused by the beam changing in temperature (ΔT):

$$u = \frac{\sigma}{Y} + \alpha \Delta T \,. \tag{4.3}$$

Solving the coupled differential equations (equation of motion and heat equation) under the condition of harmonic vibrations (beam is oscillating), the normal mode angular eigenfrequencies of the beam were determined. These eigenfrequencies are complex. The real part of the angular eigenfrequency is the angular mode frequency of the beam whilst the imaginary component describes the magnitude of the damping of the mode. Taking the ratio of the imaginary and real components enables the thermoelastic loss of a beam to be found:

$$\phi(\omega) = \frac{Y\alpha^2 T}{\rho C}\left(\frac{6}{\epsilon^2} - \frac{6}{\epsilon^3}\frac{(\sinh\epsilon + \sin\epsilon)}{\cosh\epsilon + \cos\epsilon}\right), \tag{4.4}$$

where:

$$\epsilon = t\sqrt{\frac{\pi f}{\chi}}\,, \tag{4.5}$$

and χ is the thermal diffusivity:

$$\chi = \frac{K}{\rho C} \ . \tag{4.6}$$

4.2.3 Comparison of theoretical models

The two thermoelastic loss models (Zener model and Liftshiz and Roukes model) show good agreement at frequencies close to the thermoelastic loss peak. At frequencies well below the peak a small difference of up to 2 % can occur, while at frequencies much higher than the peak there can be a difference of up to 15 %.

It should be emphasised that both theories calculate the thermoelastic loss for the bending modes of a beam, they do not calculate the thermoelastic loss of a beam's torsional modes. No analytical formula in the literature exists to calculate the thermoelastic loss of the torsional modes of a beam.

4.3 Calculating thermoelasticity using FEA

A range of FEA methods were investigated to calculate the thermoelastic loss of a cantilever. Different FEA software and methods were used to identify which method performed optimally. A brief summary of how these approaches calculated thermoelastic loss is described below.

4.3.1 ANSYS

A *harmonic project* was used to determine the thermoelastic loss of a cantilever using ANSYS. In the project, the geometry of the body was specified, the thermal and mechanical boundary conditions stipulated, and the material properties defined. To calculate the thermoelastic loss, a nominal, oscillatory force was applied to the cantilever. This force would cause the cantilever to oscillate with a particular shape and at a user-defined frequency. To calculate the cantilever's thermoelastic loss due to bending, the force would be applied to the cantilever's length/width face (as shown in Figure 4-4). To calculate the thermoelastic loss due to twisting, the force would be applied to two edges of the cantilever's length/width face in opposite directions from one another (as shown in Figure 4-4). An appropriate mesh was also designed for the body and the mathematics of the analysis defined by specifying a coupled field element.

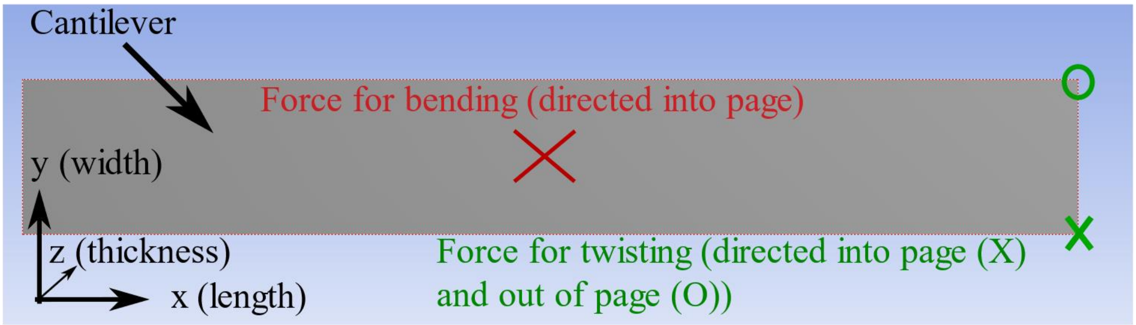

Figure 4-4: Direction and position of forces applied to the cantilever to induce bending and torsional motion.

For ANSYS to calculate thermoelastic loss, it must solve a combination of structural and thermal analyses in such a way that the inputs of one analysis depends upon the results of the other. To do this, a *coupled field analysis* is required, which can be activated by specifying a coupled field element type using a command line. For a three-dimensional cantilever, the coupled field element called Solid 226 was used [153]. Once this element had been defined in ANSYS, ANSYS automatically coupled the cantilever's equation of motion to its heat equation (as described for Liftshitz and Roukes analytical model in Section 4.2.2).

A code written by ANSYS [153] was used to calculate the thermoelastic loss of a cantilever at incremental frequencies within a specified range (by stipulating an incremental value as well as an upper and lower value of the frequency range). An oscillating force was applied to the cantilever at the minimum frequency specified in the code, causing the cantilever to vibrate. ANSYS then solved for the unknowns of the coupled differential equations at this oscillation frequency. The stress and strain solutions from these equations were then used to determine the time averaged total strain energy (U_t) for each element [153]:

$$U_t = \frac{1}{4}\int \{\sigma\}^T\{\varepsilon\}^* dv \quad , \tag{4.7}$$

where $\{\sigma\}$ is the 3×3 stress matrix, $\{\varepsilon\}^*$ is the 3×3 complex conjugate of the strain and the integral is performed over the elements volume dv. The time averaged total strain energy has both a real and imaginary component. The real part represents the average stored strain energy whilst the imaginary component describes the energy lost [153]. The total imaginary strain energy and real strain energy are found by summing the values for each element, allowing the thermoelastic loss to be calculated at the oscillation frequency [153]:

$$\text{Thermoelastic loss} = \frac{\sum_{\text{elements}} Im(U_t)}{\sum_{\text{elements}} Re(U_t)} \ . \tag{4.8}$$

The code written by ANSYS, then repeats these calculations for each specified frequency to produce a curve of thermoelastic loss versus frequency.

4.3.2 COMSOL method 1

COMSOL method 1 replicated the same approach described using ANSYS (in Section 4.3.1), except it used COMSOL FEA software.

In COMSOL, a *thermoelasticity* project [154] was selected along with a *frequency domain* study type. The geometry, material properties, boundary conditions and mesh were specified. A *boundary load* was also applied to the top surface of the cantilever or edges (to replicate the effects of the force applied in ANSYS). The thermoelastic loss of the cantilever was then calculated over a specified frequency range. COMSOL automatically solved the cantilever's coupled equation of motion and heat equation at the specified frequencies. Using volume integrals, the real and imaginary strain energies were determined and Equation (4.8) was used to calculate thermoelastic loss.

4.3.3 COMSOL method 2

This method was similar to the first COMSOL method (and ANSYS method) but utilized an *eigenfrequency* study type which meant that the thermoelastic loss was only calculated at the resonant modes of the cantilever (and not at incremental values between a lower and upper value of frequency). No pressure was applied to the cantilever in this model, as COMSOL automatically calculated the modal frequencies where the thermoelastic loss was to be determined. The thermoelastic loss was again calculated by dividing the imaginary strain energy by its real component as shown in Equation (4.8).

4.4 Building a working FEA model

Before the effect of stress and curvature (caused by coating deposition) on a substrate's thermoelastic loss can be investigated, an FEA model which is able to replicate the Zener model, Liftshitz and Roukes model and experimental data must first be built. Showing that the FEA model can replicate these simple theories as well as experimental data will give greater confidence in the results of the more complicated FEA modelling which is required to investigate the effects of stress and curvature upon thermoelastic loss.

4.4.1 Thermoelastic loss due to bending motion

An FEA model can be shown to accurately determine thermoelastic loss when its values can replicate those obtained using analytical models and/or experimentally.

The mechanical loss of an uncoated silicon cantilever 5×10^{-3} m wide, 77×10^{-6} m thick and 34×10^{-3} m long was measured using the setup described in Section 3.3. Using Equations (4.1) and (4.4), the thermoelastic loss of a cantilever with these dimensions was calculated using Zener's model and Liftshitz and Roukes's model. Similarly, using the ANSYS and COMSOL method 1 approaches, the thermoelastic loss due to bending of a cantilever with these dimensions was calculated. In the FEA models, one of the cantilever's thickness $\times$ width faces was constrained to have zero displacement to simulate the clamp, all of the faces were specified to be thermally insulated (same condition as the analytical models) and a force was applied to one of the cantilever's length $\times$ width faces (as shown in Figure 4-4). The material properties used in the analytical calculations and FEA models are listed in Table 4-1. The measured, FEA and analytical loss values are presented in Figure 4-5.

Material properties (units)	**Values**
Youngs modulus (Pa)	166×10^{9}
Poisson ratio	0.27
Density (kg/m^3)	2330
Thermal expansion coefficient (1/K)	2.6×10^{-6}
Specific heat capacity (J/kg.K)	705
Thermal conductivity (W/m.K)	140

Table 4-1: List of the silicon material properties which were used in both the analytical and FEA models.

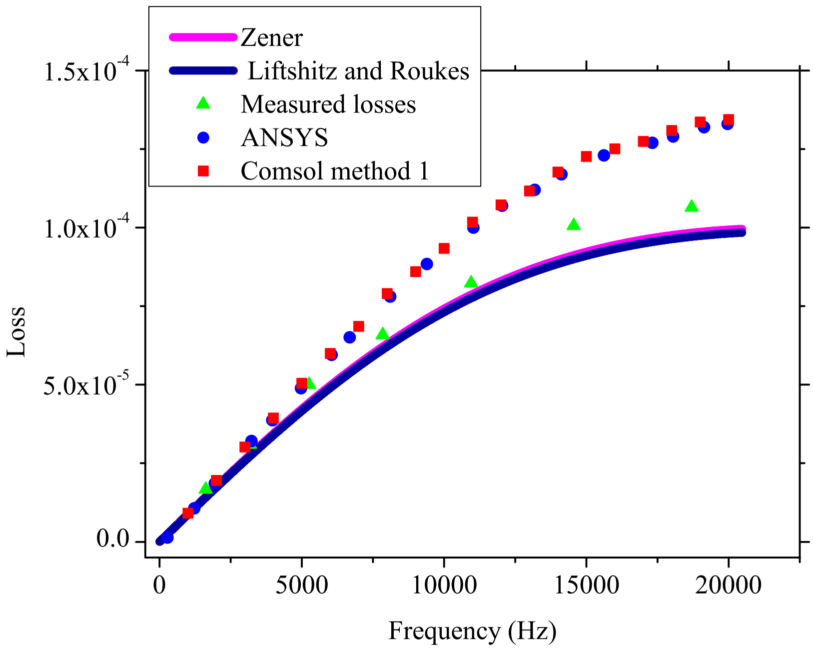

Figure 4-5: FEA and analytically determined thermoelastic loss values of a 77×10^{-6} m thick, 5×10^{-3} m wide and 34×10^{-3}m long silicon cantilever. The green triangles show the measured uncoated losses of a silicon cantilever with the same dimensions. The measured losses were obtained at the cantilever's mode frequencies. The FEA values were not determined at mode frequencies but at incremental intervals between an upper and lower frequency bound (as described in Section 4.3.1 and 4.3.2).

It can be seen that the FEA models fail to replicate the analytical or measured values well. Furthermore, it can be observed that a difference also exists between the analytical model values and measured losses. The next subsection will explore these differences in values and describe how a more accurate FEA model was built.

4.4.1.1 Difference in FEA and analytical thermoelastic models

The difference in loss values between the FEA and analytical models shown in Figure 4-5 is due to the different model assumptions made in the calculations.

The FEA model accounts for the stresses of the cantilever acting in all directions and allows for the cantilever to exhibit oscillations which vary spatially across its width. These width dependent variations were observed by splitting the cantilever into 25 equally sized strips along its length. The thermoelastic loss of each strip was found to vary at different positions across the width of the cantilever. Figure 4-6 shows the outline of some strips of the cantilever (labelled 1, 4, 9 and 13) whilst Figure 4-7 shows the thermoelastic loss of each of these strips.

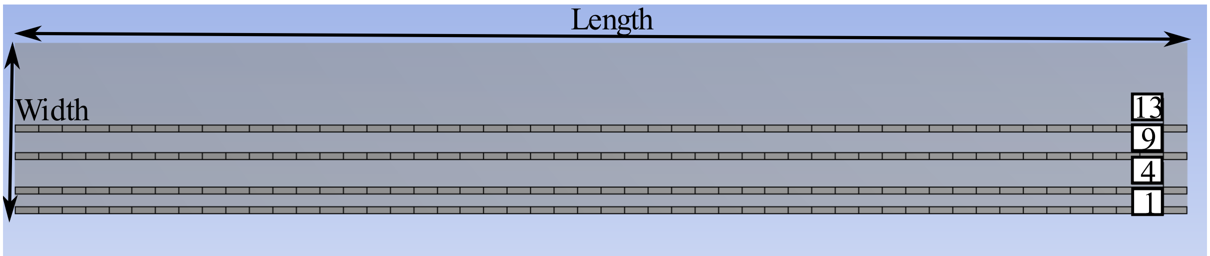

Figure 4-6: Outline of 4 of the 25 equally sized strips which the cantilever was split into. The 25 strips of the cantilever are labelled 1-25. Strips 1 and 25 are at the edges of the cantilever and strip 13 is in the middle of the cantilever.

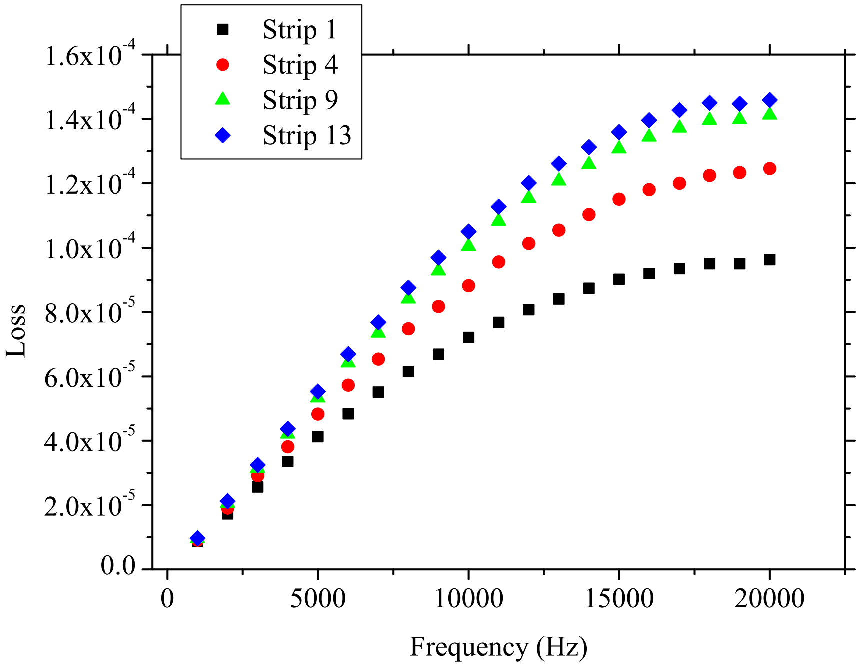

Figure 4-7: Thermoelastic loss of the four strips shown in Figure 4-6. The thermoelastic loss of the cantilever strips is symmetrical across the middle strip (strip 13).

The analytical models (Zener's model and the Liftshitz and Roukes model) do not account for oscillations which vary spatially across the width of the cantilever as they assume a Poisson ratio equal to 0 in their model. When the Poisson ratio of the FEA model is made equal to 0, no width dependent oscillations occur in the FEA model and the FEA's thermoelastic loss can be seen to be equal to Zener's model and Liftshitz and Roukes's model. Figure 4-8 shows that no width dependent oscillation occurs when the Poisson ratio of the FEA model is equal to 0 and Figure 4-9 shows that as the Poisson ratio tends to 0, the values of the analytical and FEA models converge.

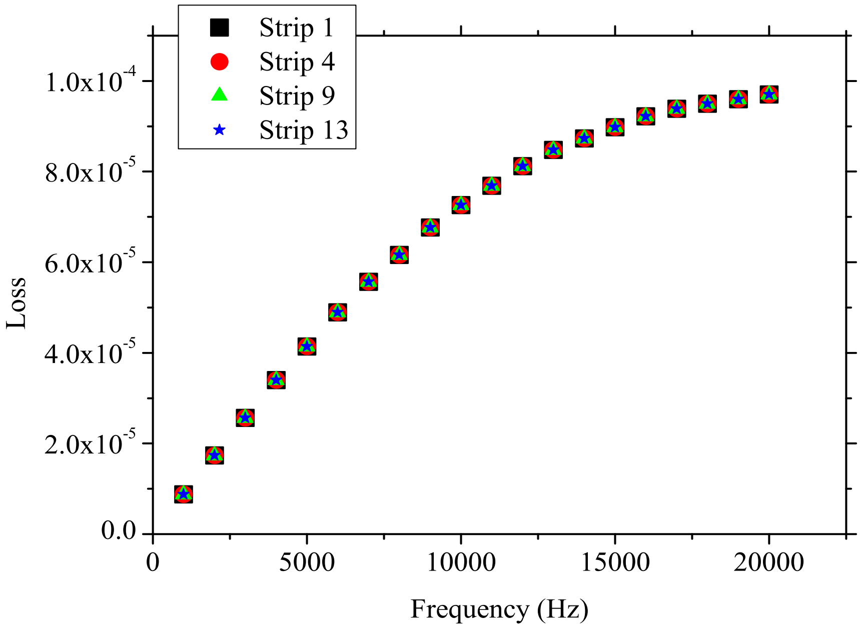

Figure 4-8: Thermoelastic loss of the four strips shown in Figure 4-6, when it is assumed that the cantilever has a Poisson ratio of 0.

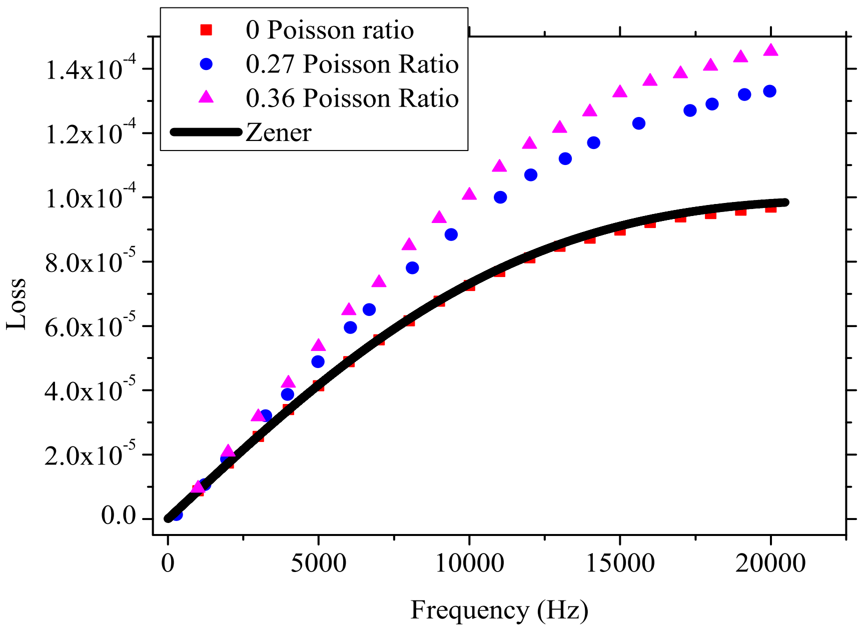

Figure 4-9: Thermoelastic loss of three geometrically identical cantilevers each with a different Poisson ratio. Zener's values of thermoelastic loss are also shown in the graph.

Having identified the difference and superiority of the FEA model compared to the analytical models (as it accounts for the more realistic cantilever conditions of width dependent variations), the question still remains as to why the analytical model is able to replicate the measured data better than the FEA model (as shown in Figure 4-5). The answer relates to the importance of modelling the silicon's anisotropic elastic matrix.

Isotropic material properties were originally used in the FEA model as both analytical models assume an isotropic beam with uniform Youngs modulus. However, in reality, silicon is anisotropic and therefore for the most accurate FEA thermoelastic loss values, the cantilever should be modelled as such. The crystal orientation of the silicon cantilevers used in the mechanical loss experiments presented here is shown in Figure 4-10.

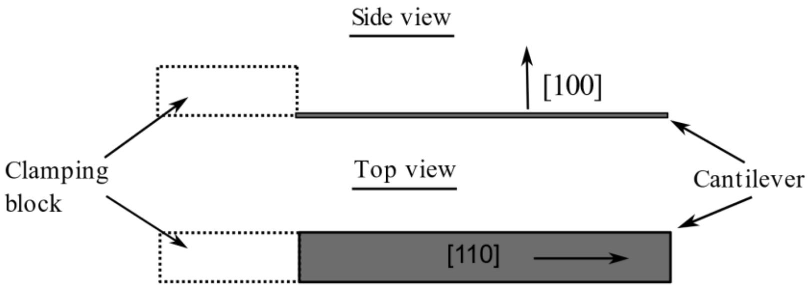

Figure 4-10: Crystal orientation of the silicon cantilever.

To define an anisotropic material in a FEA model, a 6×6 matrix known as a stiffness matrix (C) is required which is used to relate the stress of a three-dimensional anisotropic body to its strain and therefore contains all information regarding the bodies directionally dependent Young's modulus, shear modulus and Poisson ratio. When the geometry of a silicon cantilever is specified in the FEA program such that it's length, width and thickness are in the x, y and z directions, the following 6×6 stiffness matrix is required to correctly specify the cantilever's material properties as that represented in Figure 4-10:

$$C = \begin{vmatrix} 194.5 & 35.7 & 64.1 & 0 & 0 & 0 \\ 35.7 & 194.5 & 64.1 & 0 & 0 & 0 \\ 64.1 & 64.1 & 165.7 & 0 & 0 & 0 \\ 0 & 0 & 0 & 79.6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 79.6 & 0 \\ 0 & 0 & 0 & 0 & 0 & 50.9 \end{vmatrix} \tag{4.9}$$

where the matrix is ordered in standard form (x, y, z, yz, zx and xy) and the values have units of $\times\ 10^9$ Pa [155]. Figure 4-11 shows the FEA thermoelastic loss of an anisotropic silicon cantilever compared to the analytical and experimental values. It can be observed that both models now replicate one another well and the FEA model replicates the measured values more closely than the analytical model due to it's more realistic assumptions (material anisotropy).

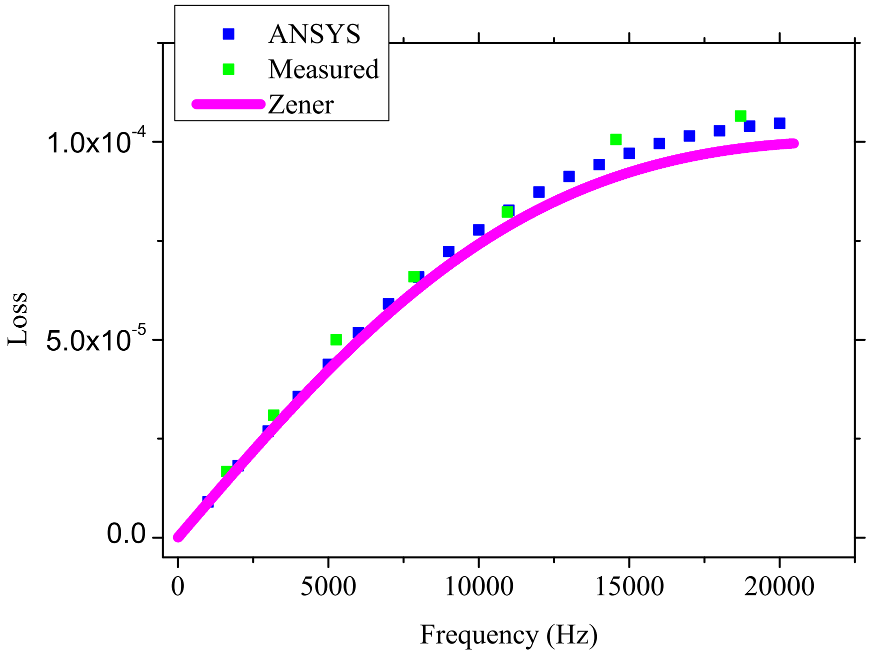

Figure 4-11: FEA thermoelastic loss of an anisotropic silicon cantilever compared to Zener's analytical values as well as the experimental values.

Moving from an isotropic to an anisotropic FEA model gives better agreement with the analytical model as the relevant Poisson ratio, which couples length and width deformations together, is significantly smaller in the anisotropic case (0.064 compared to 0.27). Thus the anisotropic case agrees more closely with the analytical model, which assumes a Poisson ratio of 0. Whilst the Poisson ratios which couple cantilever deformations in other directions are still much greater than 0 (0.36 and 0.27), it is only the Poisson ratio which couples the cantilever's length and width deformations that affects its thermoelastic loss. The sensitivity of the thermoelastic loss of the cantilever upon the length/width Poisson ratio can be observed in Figure 4-12.

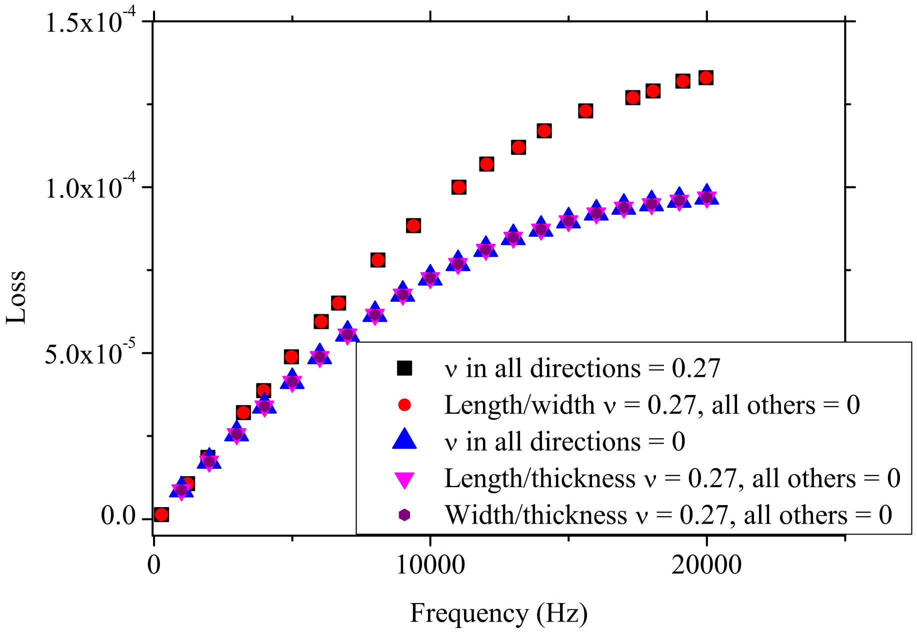

Figure 4-12: FEA thermoelastic loss of cantilevers with different Poisson ratios (ν). It can be observed that the Poisson ratio which couples the length and width deformations together has an effect upon the cantilever's thermoelastic loss whilst the Poisson ratios in the other directions do not.

The small difference in the length/width Poisson ratio from the FEA and analytical model is responsible for the greater similarity shown between the measured and FEA values as opposed to the measured and Zener model's values (shown in Figure 4-11). This can be observed in Figure 4-9 where the FEA model, which assumes a uniform Poisson ratio of 0, is shown to be equal to Zener's values. By accounting for a non-zero Poisson ratio and the width dependent variations which occur as a result of this assumption, the FEA model is able to replicate the experimental results more accurately than the theoretical model.

In summary, an FEA model which replicates both the analytical and experimental values of thermoelastic loss due to bending was built and any discrepancies between the model, analytical and experimental values was explained.

4.4.2 Thermoelastic loss due to torsional motion

There are no analytical models in the literature which determine the thermoelastic loss of a beam's torsional modes. To determine the accuracy of the FEA models in being able to calculate the thermoelastic loss of a beam's torsional modes, the FEA values must be compared to experimental results.

Figure 4-13 shows the measured torsional mode losses of an uncoated silicon cantilever, with dimensions 5×10^{-3} m wide, 77×10^{-6} m thick and 34×10^{-3} m long. Using the ANSYS and COMSOL method 1 approaches, the thermoelastic loss of the cantilever due to torsional motion was calculated (at incremental values between an upper and lower frequency bound). In the FEA models one of the cantilever's thickness × width faces was clamped, all of the faces were specified to be thermally insulated and a force was applied to two edges of the cantilever's length × width faces (as shown in Figure 4-4). The silicon cantilever was modelled as having anisotropic mechanical properties (Equation (4.9)) whilst its thermal properties are listed in Table 4-1. The FEA model's thermoelastic losses are shown in Figure 4-13 together with the measured losses.

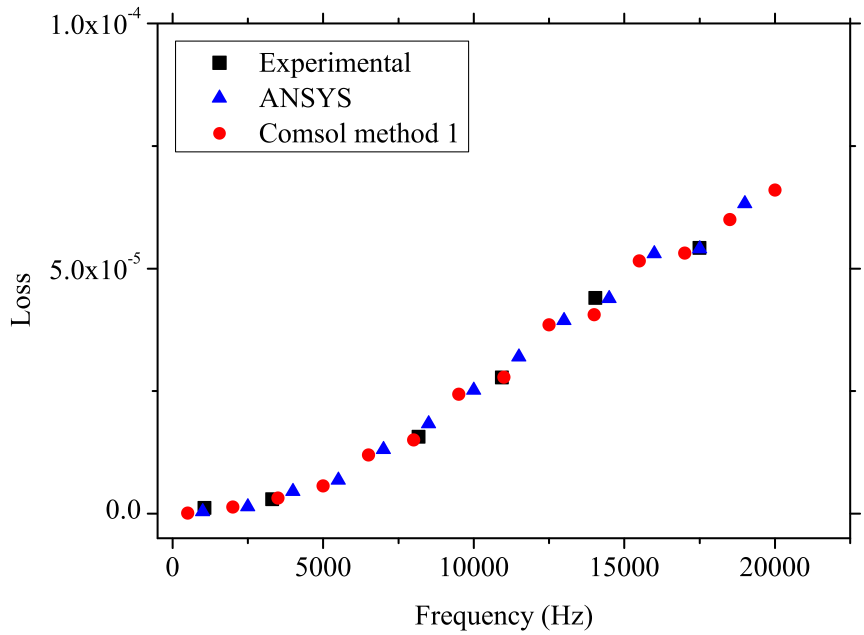

Figure 4-13: Experimental and FEA thermoelastic losses of a cantilever: 5×10^{-3} m wide, 77×10^{-6} m thick and 34×10^{-3} m long. The mechanical losses of the cantilever were measured at its torsional mode frequencies. The thermoelastic losses were calculated using COMSOL method 1 and ANSYS. The FEA thermoelastic losses were determined at incremental values between an upper and lower frequency bound.

Figure 4-13 shows that the two FEA models determine approximately identical losses to one another and replicate the experimental losses well (within 5%). It can therefore be concluded that an FEA method has been developed which can approximate experimental values. Figure 4-13 also shows that the FEA thermoelastic losses experience a small, almost sinusoidal effect (small bumps in losses with frequency). This is not thought to be a real effect. It is hypothesised that the sinusoidal effect is caused by some boundary condition of the model. A similar sinusoidal effect was observed within FEA bending thermoelastic losses and later removed by a change in clamping constraint. However, changing the clamping constraint (or any other constraint) could not be observed to have any effect upon the sinusoidal effect within the torsional thermoelastic losses.

4.4.3 Limitations of FEA models for bending and torsional motion

Whilst both ANSYS and COMSOL method 1 are able to determine non-resonant thermoelastic losses of a cantilever, neither method is able to determine the thermoelastic loss at exactly a bending or torsional resonant mode. Figure 4-14 and Figure 4-15 show the result of using the ANSYS method to determine the thermoelastic loss at a bending and torsional resonant mode. The same effect was observed when using COMSOL method 1.

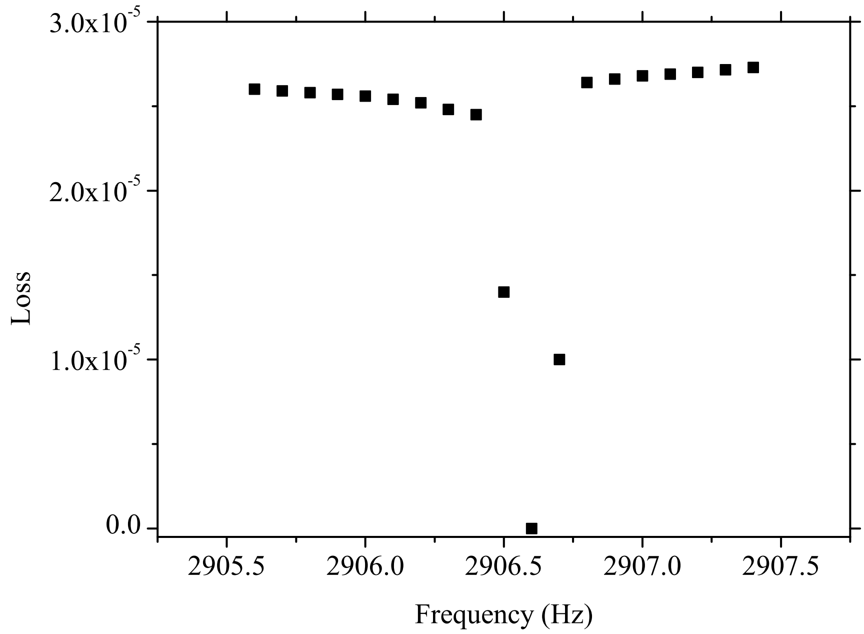

Figure 4-14: FEA thermoelastic loss of a silicon cantilever that is 72.5×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34.4×10^{-3} *m long. A bending mode exists 2906.6 Hz.*

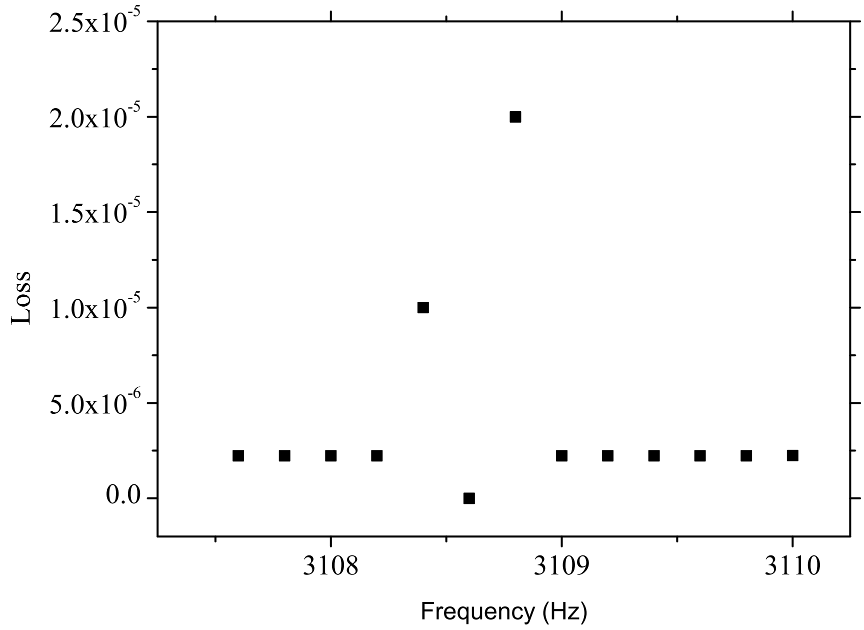

Figure 4-15: FEA thermoelastic loss of a silicon cantilever that is 72.5×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34.4×10^{-3} *m long. A torsional mode exists 3108.6 Hz.*

This observation of the model not working at resonant frequencies was confirmed by ANSYS support who summarised the problem as being a result of nonlinear instabilities occurring at resonant frequencies. Essentially, ANSYS support's explanation was that the ANSYS model calculates thermoelastic loss by forcing the cantilever to bend/twist at a particular frequency with a hypothetical amplitude and then calculates the cantilever's real and imaginary strain energies which are used to obtain the thermoelastic loss. At a non-resonant frequency, this approach works fine. However, at a resonant frequency the applied force causes the amplitude of the cantilever to tend to infinity which results in the cantilever's real strain energy (stored energy) tending to infinity and therefore the thermoelastic loss tending to 0 (Equation 4.8). Thermoelastic losses calculated slightly off resonant frequency however can be used to approximate the resonant loss. COMSOL method 1 suffers from the same problem as ANSYS.

COMSOL method 2 was built to directly determine the thermoelastic loss of a cantilever at both its bending and torsional resonant mode frequencies. Figure 4-16 shows the FEA thermoelastic loss against the analytically determined and experimentally measured values. Good agreement can be observed between the measured and COMSOL method 2's values. For bending modes, the COMSOL values underestimate and overestimate the measured losses by a maximum of 10 % and 2 % respectively. For torsional modes (excluding the first mode which shows a 50 % difference due to convergence issues), the COMSOL values underestimate and overestimate by a maximum of 8 %. From Section 3.4.1.2, it is speculated by modelling the losses of three other silicon cantilevers, that on average, COMSOL method 2 underestimates measured bending mode losses (2nd bending mode and higher) by 4 % whilst for torsional modes (3rd torsional mode and higher) there is no systematic offset. Figure 4-16 confirms these results as on average COMSOL underestimated the bending and torsional mode losses by 4.3 %, and 0.1 % respectively. Good agreement can be observed between the analytical model and COMSOL method 2's values. The difference between the analytical values and COMSOL method 2's values was earlier identified (as being a result of the FEA model having a non-zero Poisson ratio, Figure 4-11).

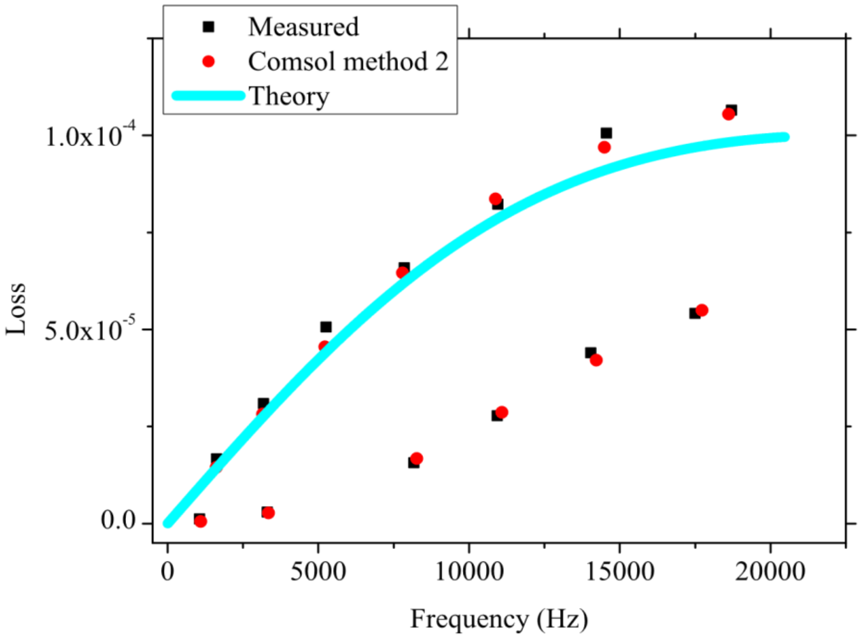

Figure 4-16: FEA thermoelastic loss of a silicon cantilever that is 77×10^{-6} m thick, 5×10^{-3}m wide and 34×10^{-3} m long. The FEA thermoelastic loss was calculated using COMSOL method 2 and the values were determined at mode frequencies. Shown alongside these values are the measured losses of a bare cantilever, with the same dimensions, as well as the analytical values determined using Zener's equation (theory). The measured losses equal to or greater than Zener's values are bending modes whilst those which are lower are torsional modes.

Being able to determine the thermoelastic loss of both bending and torsional modes is essential if the effect of stress and curvature upon both mode types is to be investigated.

4.5 Effect of stress on thermoelastic loss

4.5.1 Building and verifying a stressed FEA model

Before the effect of stress and curvature on the thermoelastic loss of a coated cantilever can be investigated, it must first be shown that the FEA models developed are able to accurately replicate the effect of simple stress patterns upon thermoelastic loss. Both Kumar and Haque [149], and Lepage [156] investigated the effect of simple stress patterns upon thermoelastic loss.

Kumar and Haque derived an equation which describes how the thermoelastic loss of a beam, clamped at both ends, changes with tensile stress (where the beam's length is much greater than its width and thickness, the beam is clamped along its width × thickness faces and tensile stress is applied axially along the length of the beam). The derivation is an extension of the Liftshitz and Roukes theory of thermoelastic loss (Section 4.2.2), in which the initial stress of the beam, before deformation, is taken into account. Kumar and Haque modified Liftshitz and Roukes equation for the strain of a bent beam (Equation (4.3)) to include the magnitude of the beam's initial stress (σ_0), in addition to the strain induced by bending ($\frac{\sigma}{Y}$) and thermal expansion due to the temperature change of the deformed beam ($\alpha \Delta T$):

$$u = \frac{\sigma}{Y} + \alpha \Delta T - \frac{\sigma_0}{Y} \ . \tag{4.10}$$

Equation (4.10) shows that as the initial tensile stress increases, the beam becomes more difficult to bend resulting in less deformation, less strain and less thermoelastic loss:

$$\phi(\omega) = \frac{\left(\frac{Y\alpha^2 T}{C}\right)\left(\frac{6}{\epsilon^2} - \frac{6}{\epsilon^3}\frac{(\sinh\epsilon + \sin\epsilon)}{\cosh\epsilon + \cos\epsilon}\right)}{1 + 0.97\left(\frac{\sigma_0}{a_1\pi^2 Y}\right)\left(\frac{l}{d}\right)^2} \ , \tag{4.11}$$

where l represents the length of the beam, d represents the thickness of the beam and a_1 is a boundary condition factor which was determined by knowing that:

$$0.97\left(\frac{\sigma_0}{a_1\pi^2 Y}\right)\left(\frac{l}{d}\right)^2 = \frac{\left(\left(\frac{f_{\mathrm{stress}}}{f_{\mathrm{unstressed}}}\right)^2 - 1\right)}{0.97} \ , \tag{4.12}$$

where f_{stress} represents one of the stressed beam's modal frequencies and $f_{\text{unstressed}}$ represents the same mode's frequency when the beam is unstressed.

Lepage developed a finite element method to determine the thermoelastic loss of a clamped-clamped beam which was stressed axially along its length. A mathematical model of the thermopiezoeletric properties of a body is formulated, in which the body's thermal, electric and mechanical fields are coupled together in matrix form [156]. By setting a bodies electrical effects to zero, the thermoelastic loss of a body can be calculated by determining the real and imaginary parts of its eigenfrequency (see Section 4.2.2). This was shown to replicate Zener's results. Lepage then used this model to investigate the effects of stress upon thermoelastic loss. The initial stress of the beam was incorporated into the finite element method by amending the beam's stiffness matrix [156] (which relates the magnitude of deformation a beam will experience when a force is applied to it). The beam's stiffness matrix depends upon the beam's geometry, stress conditions and material properties. Lepage observed that as the beam's tensile axial force increased, the stiffness of the beam increased, resulting in less strain of the beam and ultimately less thermoelastic loss.

Using both the theory of Lepage and the theory of Kumar and Haque, the thermoelastic loss of a clamped-clamped beam's first resonant mode was calculated as its initial stress was varied. The beam was 90×10^{-6} m long, 4.5×10^{-6} m wide and 4.5×10^{-6} m thick. The results are shown in Figure 4-17. COMSOL method 2 was used to replicate these results. This method was chosen as it is the only FEA method able to calculate the loss at resonant modes and it therefore allows direct comparison with the other models. To incorporate an axial stress within COMSOL, the *initial stress and strain* functionality was used. This addition to the model enabled analytical frequencies of a pre-tensioned clamped-clamped wire to be replicated in COMSOL, confirming the validity of the approach. The thermoelastic loss values calculated using COMSOL method 2 are shown in Figure 4-17.

Figure 4-17 shows that as the beam becomes stiffer and more difficult to bend, the thermoelastic loss decreases in the Kumar and Haque, Lepage and COMSOL models. For the range of stresses studied, these models agree to within 10 %. The difference between Kumar and Haque's loss values and the other two models is due to different model assumptions being used. The Lepage and COMSOL models account for stresses and strains acting in all three directions of the beam (i.e. they account for width dependent variations in thermoelastic loss) whilst Kumar and Haque's model only considers stress acting along its length. This can be observed by rerunning the COMSOL model with a Poisson ratio equal to 0, which removes any width dependent variations in thermoelastic loss (as explained in Section 4.4.1.1). The values obtained are shown in Figure 4-17 and are labelled 'Comsol ν = 0'. It can be observed that by removing the width dependent variations in thermoelastic loss, the COMSOL model is able to replicate Kumar and Haque's analytical values almost exactly (within 1 %). Since Lepage's and COMSOL's values fall within 2 % of one another and the (less than 10 %) variation between COMSOL's and Kumar and Haque's values can be explained, it can be concluded that a COMSOL model has been built which is able to successfully replicate axially stressed thermoelastic loss models (from in the literature).

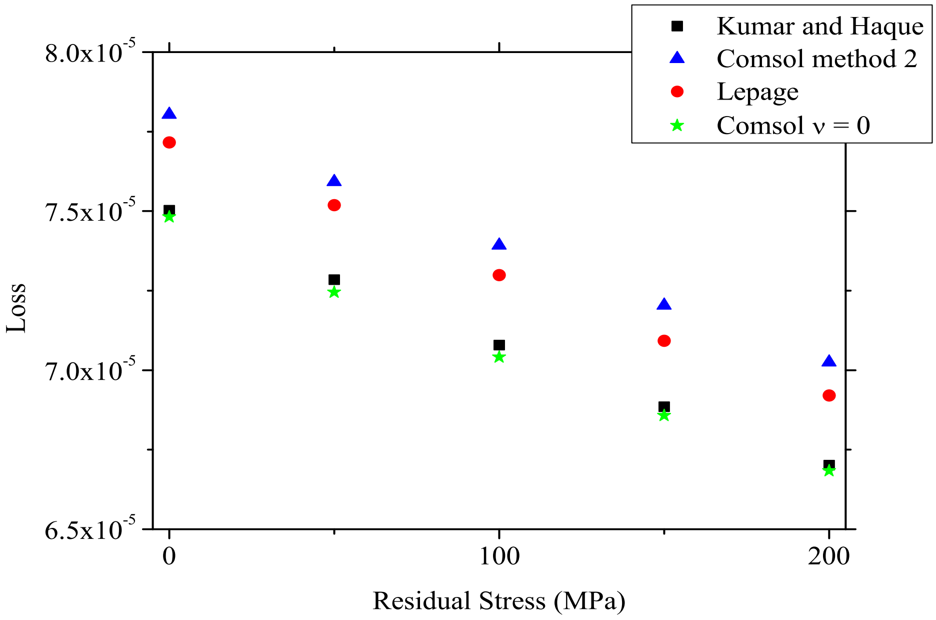

Figure 4-17: Varyingly stressed, silicon beam's first resonant mode's thermoelastic loss. The silicon beam was 90×10^{-6} *m long,* 4.5×10^{-6} *m wide and* 4.5×10^{-6} *m thick. The thermoelastic loss was calculated using COMSOL method 2, Lepage's approach and Kumar and Haque's equation. The thermoelastic loss of the beam was also calculated using the COMSOL method 2 approach when the material was assumed to have a Poisson ratio equal to 0.*

It is interesting and important to note that the thermoelastic loss of the beam decreases as the stress increases due to both:

- a change in mode frequency and
- the stress intrinsically changes the thermoelastic dissipation of the beam (imaginary part of the eigenfrequency, as described in Section 4.2.2).

This can be observed in Figures 4-18 and 4-19. Figure 4-18 shows the thermoelastic loss of the beam's first resonant mode for different axial stresses (i.e. same points as the red circles in Figure 4-17). Shown against these values are the thermoelastic losses of the unstressed beam at the same frequencies as the stressed resonant modes (off resonant values). It can be observed that the decrease in thermoelastic loss caused by increasingly stressing the beam is greater than the expected decrease in thermoelastic loss caused by only the resonant mode frequency changing in value. Figure 4-19 explains this larger than expected decrease in loss as being a result of the stress reducing the dissipation of the beam. The dissipation values were determined using COMSOL method 2 and Lepage's method to further highlight the similarity between both three-dimensional models. To conclude, a COMSOL model has been built which is able to successfully replicate axially stressed thermoelastic loss models (discussed in the literature).

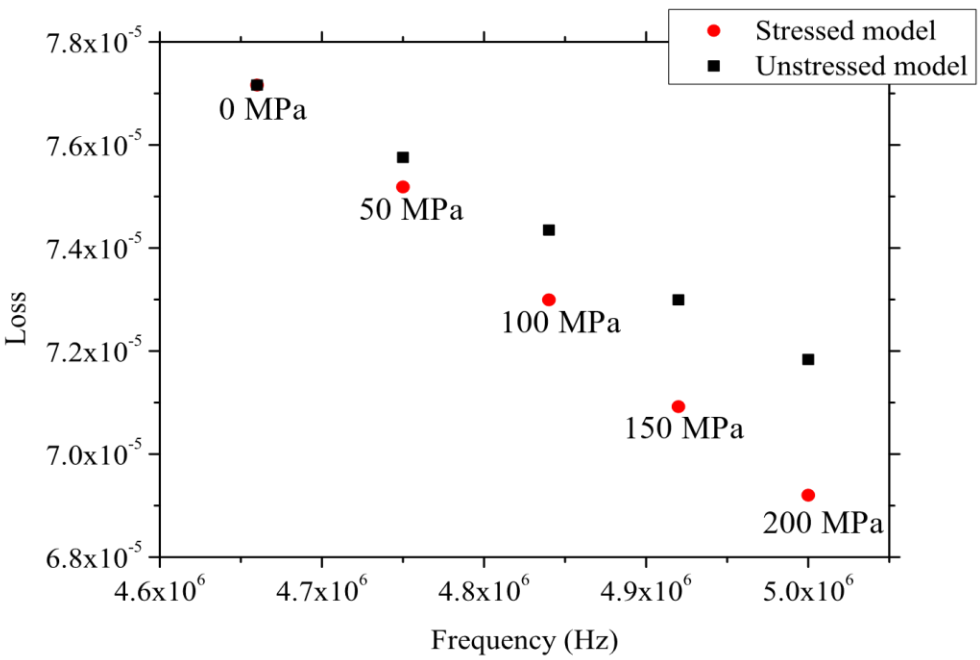

Figure 4-18: Silicon beam's first resonant mode's thermoelastic loss under different stress conditions (red points). Off-resonance thermoelastic losses of the unstressed silicon beam are also shown. It can be observed that the decrease in thermoelastic loss due to stress can not be fully explained by the beam resonating at a higher frequency. The silicon beam modelled was 90×10^{-6} *m long,* 4.5×10^{-6} *m wide and* 4.5×10^{-6} *m thick.*

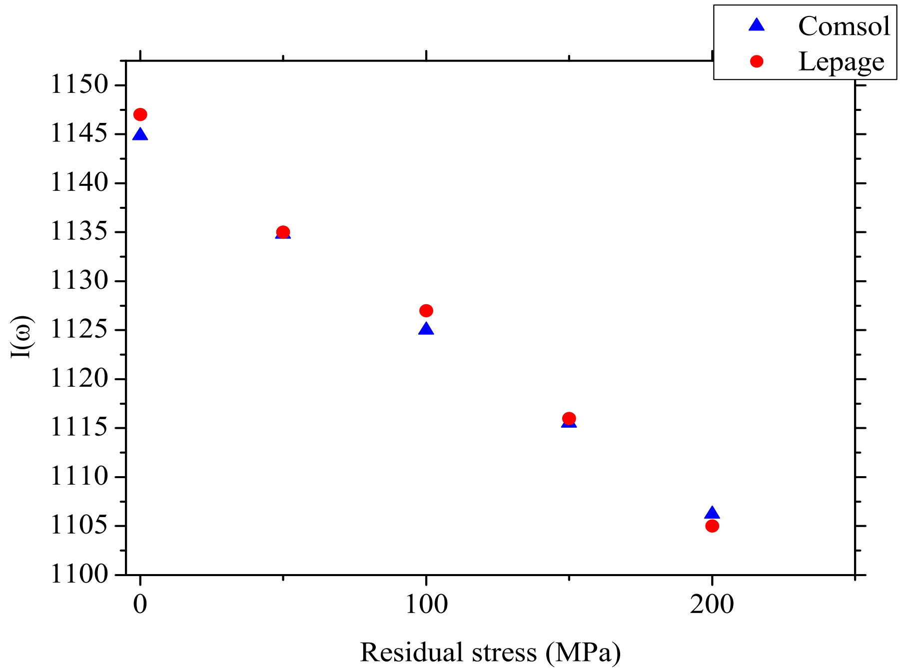

Figure 4-19: Dissipation (imaginary component of angular eigenfrequency $I(\omega)$) for a varyingly stressed silicon beam's first resonant mode. The silicon beam was 90×10^{-6} m long, 4.5×10^{-6} m wide and 4.5×10^{-6} m thick. The dissipation was calculated using COMSOL method 2 and Lepage's approach. Both models agree very well with one another.

4.5.2 Effect of different directions of stress upon a cantilever's thermoelastic loss

Using the COMSOL model outlined in Section 4.5.1, the effects of simple axial stresses upon the thermoelastic loss of a cantilever was investigated. The cantilever's geometry was 34×10^{-3} m long, 5×10^{-3} m wide and 72.5×10^{-6} m thick. One of the cantilever's width $\times$ thickness faces was clamped using a fixed constraint. The effects of various compressive and tensile stresses acting in different directions was then investigated for both torsional and bending modes. In this analysis, the first bending and torsional mode frequencies were excluded due to lack of convergence in the FEA model.

Initially, the effects of stress along the length of the cantilever were investigated. As shown in Figure 4-20, the thermoelastic loss reduces as the tensile stress acting along the length increases (where positive stress represents tensile stress whilst negative stress represents compressive stress). This is a result of the stress stiffening the cantilever, resulting in less bending deformation and therefore less damping due to heat flow. Another effect of the increasing cantilever stiffness is an increase in the mode frequencies, also seen in Figure 4-

20. The opposite effect can be observed for a compressively stressed cantilever, where its frequency can be observed to decrease and its loss to increase. It should also be noted that the change in frequency and loss is greater for bending modes than torsional modes. The most likely explanation for this is that by changing the stress along the length of the cantilever, the bending mode frequencies and bending mode thermoelastic loss will be directly affected as their values are dependent upon the stiffness of the cantilever along its length (Youngs modulus). This is not the case for torsional mode frequencies (and perhaps torsional mode thermoelastic loss) as they are dependent upon the material's stiffness in its length and width plane (shear modulus). This explanation is consistent with the effect of different stresses acting across the width of the cantilever (see Figure 4-21). Changing the stress along the width has no effect on the bending modes but does affect the torsional modes frequency and loss. It can once again be observed in Figure 4-21 that tensile stress leads to a reduction in thermoelastic loss whilst compressive stress leads to an increase.

Figure 4-22 shows the thermoelastic loss of the silicon cantilever when varying magnitudes of stress is applied through its thickness. Stress in this direction has a similar effect as stress along the length of the cantilever, with the exception that the torsional modes frequencies and losses are more greatly affected. It is not yet understood why stress acting through the thickness has an effect upon the loss and frequencies of the cantilever.

It should be noted that in Figures 4-20, 4-21 and 4-22, the lowest in frequency modal losses of the most compressively stressed cantilever was not shown. This is because as the cantilever became more compressively stressed, the frequency of its lowest resonant modes begin to tend to 0. As the frequency tends to 0, instabilities occurred within the software which led to purely imaginary losses being determined.

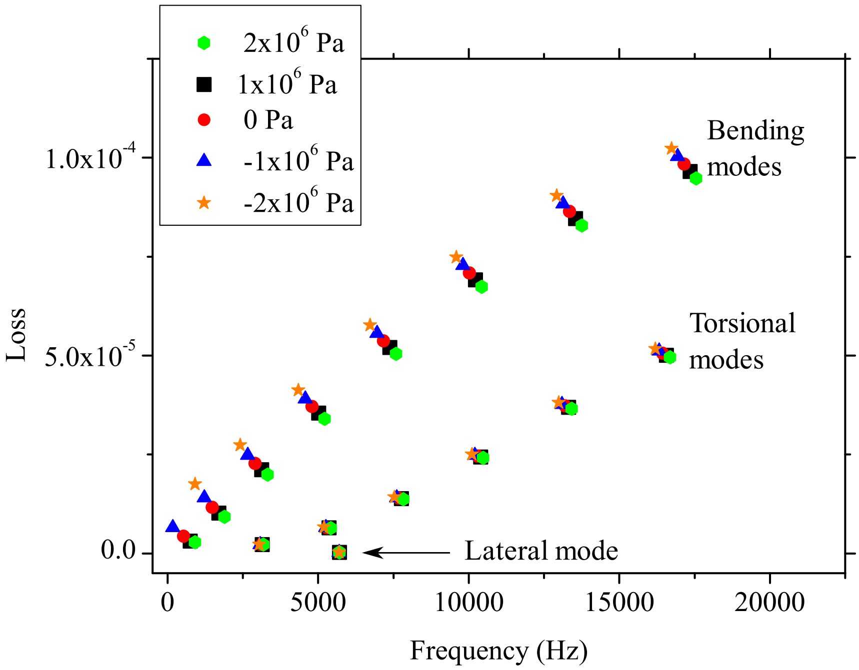

Figure 4-20: Varyingly stressed, silicon cantilever's thermoelastic loss. The stress was applied along the length of the cantilever. The silicon cantilever was 34×10^{-3} *m long,* 5×10^{-3} *m wide and* 72.5×10^{-6} *m thick. Positive stress represents tensile stress whilst negative stress represents compressive stress.*

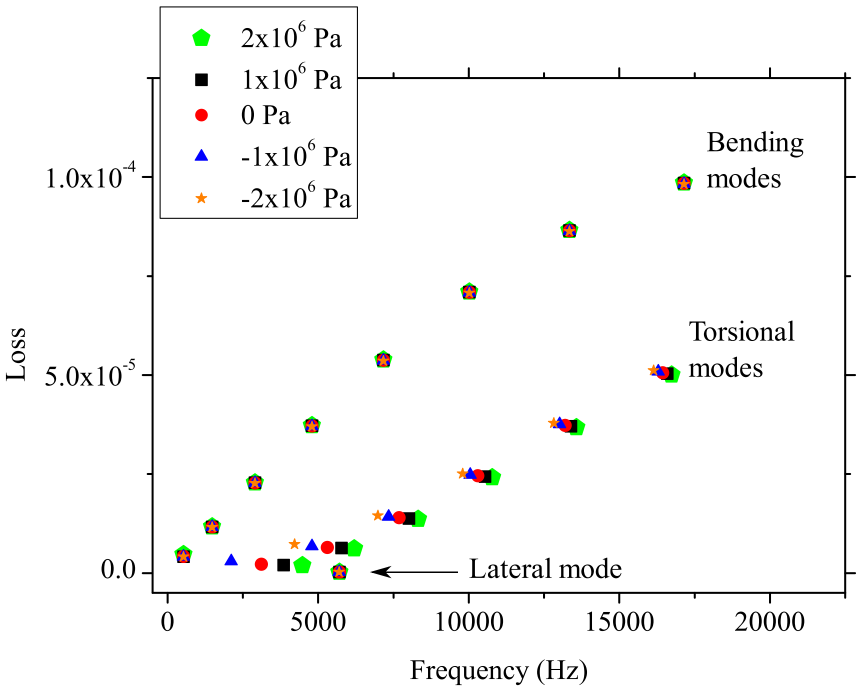

Figure 4-21: Thermoelastic loss for a silicon cantilever with stress of varying magnitude applied along the cantilever's width. The silicon cantilever was 34×10^{-3} *m long,* 5×10^{-3} *m wide and* 72.5×10^{-6} *m thick. Positive stress represents tensile stress whilst negative stress represents compressive stress.*

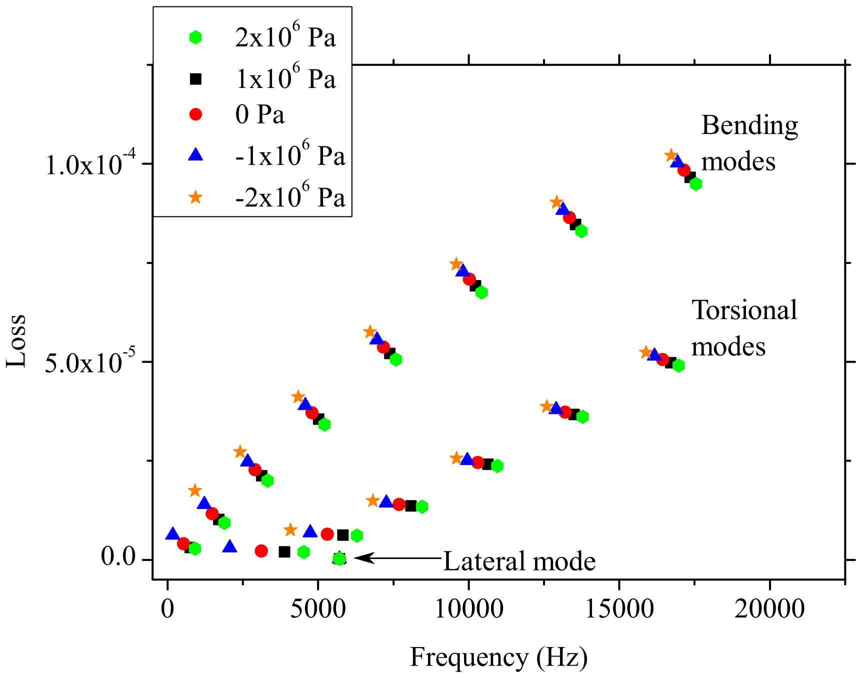

Figure 4-22: Thermoelastic loss for a silicon cantilever with stress of varying magnitude applied along the cantilever's thickness. The silicon cantilever was 34×10^{-3} *m long,* 5×10^{-3} *m wide and* 72.5×10^{-6} *m thick. Positive stress represents tensile stress whilst negative stress represents compressive stress.*

This section presents how different magnitudes and different directions of simple axial stresses affect the thermoelastic loss of a cantilever used in mechanical loss experiments. This chapter will now focus on investigating the more complicated effect of how a cantilever's thermoelastic loss could be affected by the curvature and stress induced by depositing a coating on its surface.

4.5.3 Change in a cantilever's thermoelastic loss when a coating has been deposited upon it

Coatings are often deposited upon cantilevers at high temperatures [157]. However, when the coated cantilever cools to room temperature after deposition, large stresses and curvatures are induced in the coated cantilever. As discussed in the introduction, it is speculated (from the work carried out by others and observations made by the author) that these stresses and curvatures could affect the original thermoelastic loss of the cantilever.

It would have been ideal to have stressed and curved a coated cantilever as described in Section 3.6.2.1 and then fed the stressed and deformed geometry of the coated cantilever into one of the FEA approaches (detailed in Section 4.3) which was able to calculate thermoelastic loss. The perfect scenario would have been then to use the FEA software to separate the thermoelastic loss of the coated cantilever into its cantilever and coating components. This would have enabled the effect of stress and curvature upon the cantilever substrate to be identified. However, this was not possible to do using either ANSYS or COMSOL for two reasons. Firstly, neither COMSOL or ANSYS permitted the solutions from a separate study (which would have been used to **stress and deform** the coated cantilever) to be directly exported into a model able to calculate thermoelastic loss. Secondly, it was not possible to separate a coated cantilever's FEA thermoelastic loss into it's cantilever and coating contributions, as heat flows from the coating into the cantilever and vice versa, affecting their individual loss. As a result, new methods were developed to investigate how a substrate's thermoelastic loss is affected by the stress and curvature induced by coating deposition. This also meant that only one of the two effects (stress or curvature) upon thermoelastic loss was examined at one time.

4.5.4 Stress and thermoelastic loss

When a coated cantilever bends (at a bending mode) and twists (at a torsional mode), both the coating and cantilever bend and twist together. The amount of bending and twisting the coated cantilever experiences is dependent on both the stress in the coating and the stress in the cantilever. As a result, both the stress in the coating and the stress in the cantilever must be incorporated within a coated cantilever model when its thermoelastic loss is calculated. Unfortunately, in a computer model of this type it is not possible to determine only the thermoelastic loss of the cantilever as heat will flow from the coating into the cantilever and vice versa, thus affecting the cantilever's value of thermoelastic loss. However, it should still be clear from comparing the unstressed and stressed coated cantilever thermoelastic loss values whether the addition of stress increases / decreases or does not affect the thermoelastic loss of the cantilever. The reason for this is that the thermoelastic loss of the coated cantilever is dominated by the loss of the cantilever.

The coated cantilever modelled in COMSOL was the same as that used in Figure 4-1, as it was observed experimentally that something (perhaps stress and/ or curvature) affected the coated cantilever's thermoelastic loss, resulting in the measured coated losses being less than the calculated uncoated loss. This cantilever was 64×10^{-6} m thick, 5×10^{-3} m wide and 34×10^{-3} m long. The coating was 1×10^{-6} m thick and was made from amorphous silicon. Using COMSOL method 2, the thermoelastic loss of the coated cantilever assuming no stress and no curvature was calculated and is shown in Figure 4-25. To determine the thermoelastic loss of the stressed coated cantilever, the stresses of both the coating and cantilever had to first be determined. These stresses were obtained using the same ANSYS method described in Section 3.6.2.1 to determine the stressed energy ratios i.e. the coated cantilever was cooled from its deposition temperature to room temperature and taking advantage of the mismatch in the coating's and cantilever's thermal expansion coefficients, the coated cantilever was stressed and curved until it replicated the experimentally measured vertical deflection, direction of stress and mode frequencies. The stresses at each mesh element of the coated cantilever were determined and polynomials were used to fit the normal (x, y and z) and shear (xy, xz and yz) stresses of the cantilever. Examples of the quality of fits are shown in Figure 4-23 and 4-24. The coated cantilever was designed in ANSYS such that its length, width and thickness were along the x, y and z axes respectively.

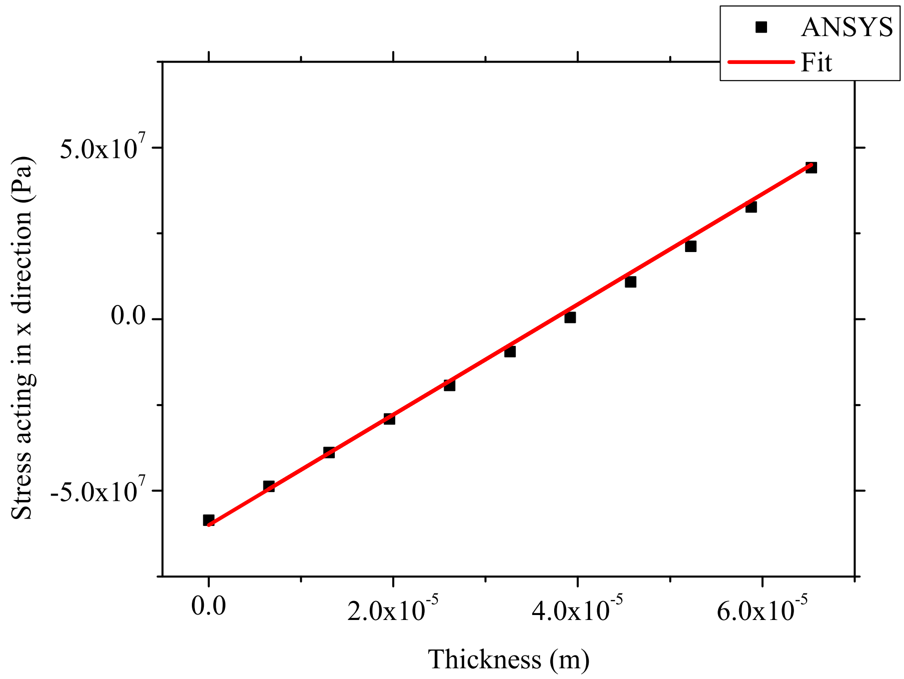

Figure 4-23: ANSYS and fitted stress values through the thickness of the cantilever. The stress was directed along the length of the coated cantilever (x axis). The values were taken 1.5×10^{-3} *m along the coated cantilever's length away from the clamping block and on one of the width edges. Positive stress represents tensile stress whilst negative stress represents compressive stress.*

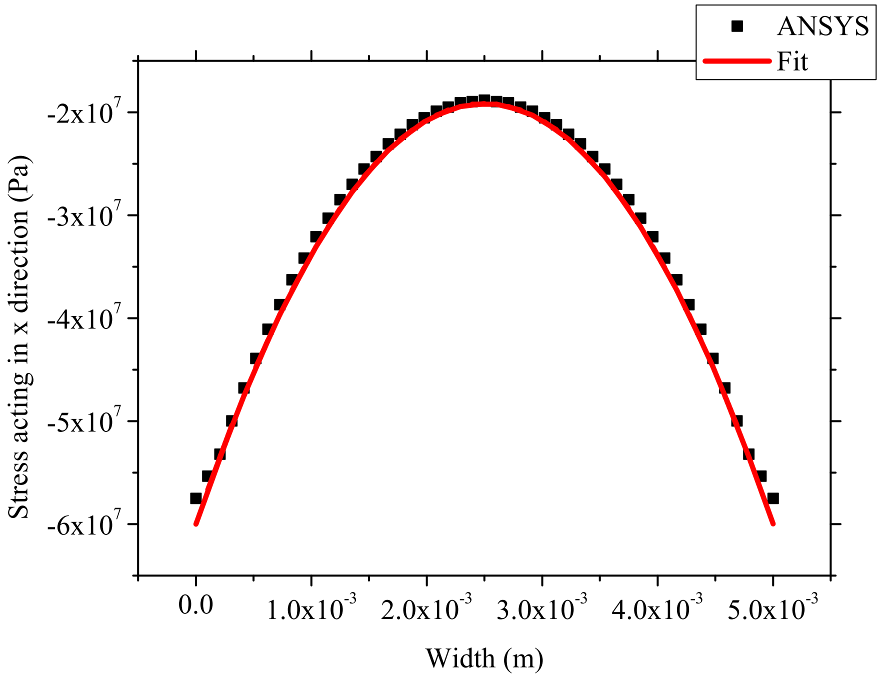

Figure 4-24: ANSYS and fitted stress values across the width of the coated cantilever. The stress was directed along the length of the coated cantilever (x axis). The values were taken 1×10^{-3} *m along the coated cantilever's length (away from the clamping block) and at the bottom of the coated cantilever (opposite side from coating). Positive stress represents tensile stress whilst negative stress represents compressive stress.*

These fitted polynomials were then used to stress the coated cantilever in COMSOL and the thermoelastic loss of the body was calculated. The thermoelastic losses of the stressed and unstressed coated cantilever are shown in Figure 4-25.

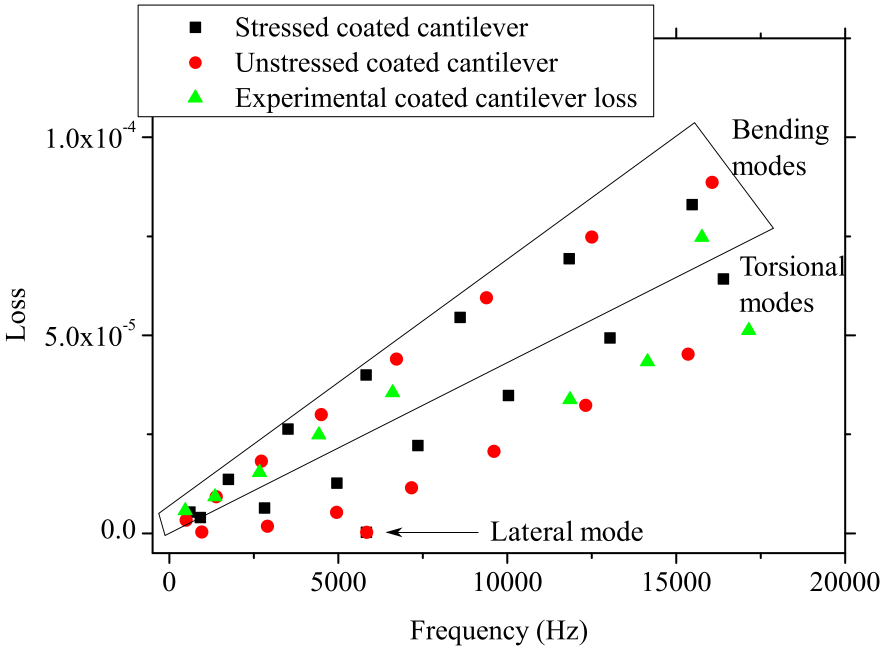

Figure 4-25: Stressed and unstressed coated cantilever's FEA thermoelastic loss. For comparison the experimental coated cantilever loss is also shown. The silicon cantilever was 34×10^{-3} *m long,* 5×10^{-3} *m wide and* 64×10^{-6} *m thick. The amorphous silicon coating was* 1×10^{-6} *m thick. For clarity, all of bending modes of the FEA models and experimental results are enclosed by a line.*

The addition of stress has a small effect on the bending mode losses but has a large effect on increasing the torsional mode losses. This does not agree with the experimental observations, as they suggest that the addition of stress should have a larger effect upon the bending mode losses and a smaller effect upon the torsional mode losses. However, it is difficult to quantify the significance of this result due to errors associated with the stress fits. Whilst it can be said that overall the polynomial fits gave a good representation of the FEA stress pattern, some directions of stress at specific locations in the coated cantilever were not fitted well as shown in Figure 4-26.

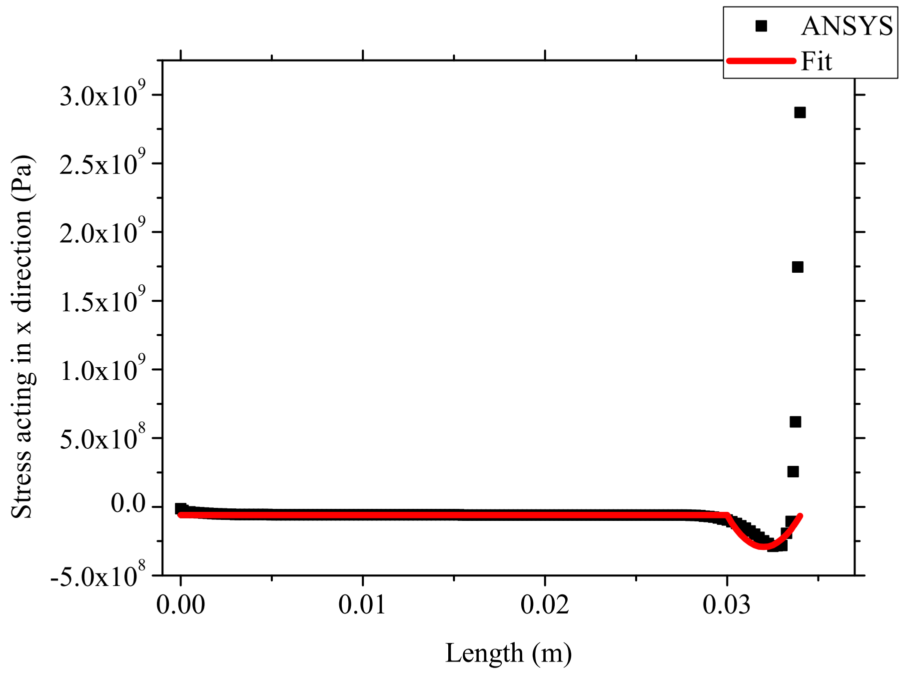

Figure 4-26: ANSYS and fitted stress values along the length of the coated cantilever. The stress's direction was also along the length of the coated cantilever. Positive stress represents tensile stress whilst negative stress represents compressive stress.

Attempts were made to fit the stresses better. However, only small incremental improvements could be made as a result of localised changes in stress within the coated cantilever that were too random to fit (particularly around the edges of the coated cantilever). Fitting the stress pattern was further complicated by COMSOL experiencing rounding problems when higher order polynomials were used in the fitting process. Whilst slightly increasing the quality of fits did not significantly change the thermoelastic loss results, the fits were not exact and the outcomes must therefore be questioned.

However, it may be possible to identify how the stress caused from coating deposition affects thermoelastic loss by use of a simple argument. Consider a stressed, curved, coated cantilever resting on a table. At rest, the coated cantilever is in equilibrium (as otherwise it would move). The coated cantilever's stresses in all directions must therefore be equally balanced i.e. the tensile or compressive stress in the coating must be balanced by a stress acting in the cantilever. If we model a coated cantilever in a *thermoelastic analysis* and apply an arbitrary tensile stress to the cantilever acting along its length, it can be observed that both the frequencies and loss of the coated cantilever change due to the stiffness changing. This is shown in Figure 4-27. However, in the same figure it can be observed that if we now also apply a compressive stress to the coating and increase it in magnitude, both the frequencies and thermoelastic loss of the coated cantilever shift back to their original unstressed values. When the force caused by the stress acting in the cantilever is perfectly balanced by the force caused by the stress in the coating (-1.92×10^8 Pa in the coating and 3×10^6 Pa in the substrate), both the loss and frequency of the coated cantilever perfectly match the values in the unstressed case. This same effect was observed to occur when stresses were applied and balanced in all of the other cantilever directions. This result suggests that the stresses in the coating and cantilever were not properly incorporated in the COMSOL model used to plot Figure 4-25. It implies that the net force on the coated cantilever was not zero, resulting in different thermoelastic loss values for the stressed model compared to unstressed model.

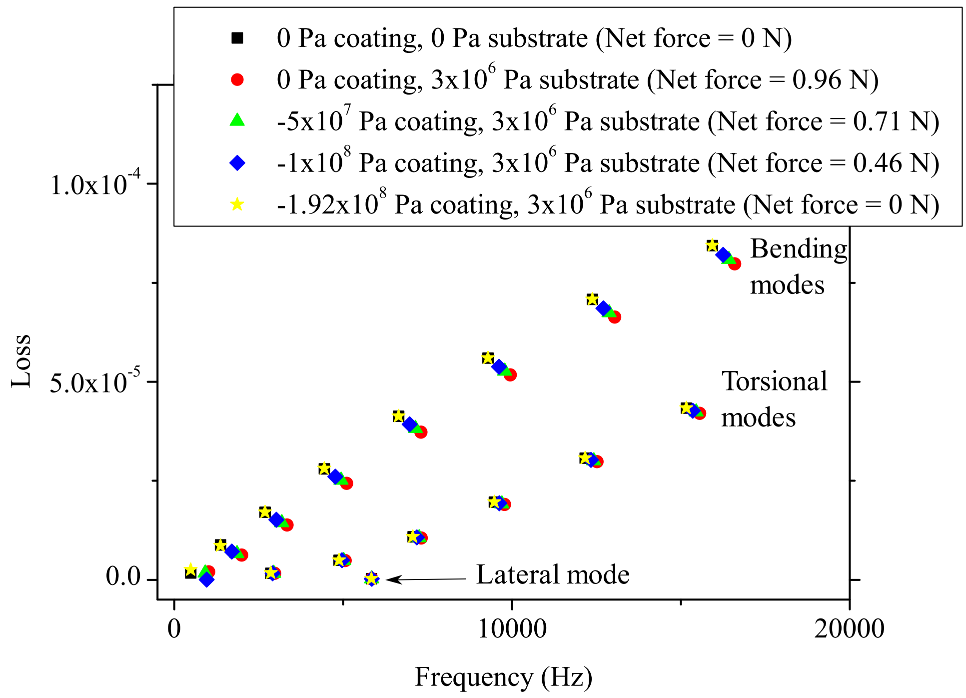

Figure 4-27: Varyingly stressed coated silicon cantilever's thermoelastic loss. The stress was acting along the length of the coating and cantilever. The net force of the coated cantilever is listed in the legend. The silicon cantilever was 34×10^{-3} m long, 5×10^{-3} m wide and 64×10^{-6} m thick. The amorphous silicon coating was 1×10^{-6} m thick.

To conclude, due to theoretical arguments and FEA models it was shown that the thermoelastic losses of an unstressed and stressed coated cantilever (in equilibrium) are the same and therefore it can be concluded that mechanical loss experiments whose aim is to determine the loss of a coating are unaffected by the effect of stress upon thermoelastic loss.

4.6 Effect of curvature on thermoelastic loss

This section presents an investigation into the effect of curvature on thermoelastic loss. FEA is used to examine whether cantilever curvature can explain the results shown in Figures 4-1 and 4-2, where it was observed that some of the measured bending mode losses of the coated cantilever were less than the thermoelastic losses of the uncoated (and uncurved) cantilevers. The section will conclude by presenting the results from an experiment whose aim was to verify the computational results (of the effects of curvature upon thermoelastic loss).

4.6.1 Investigating the effect of curvature upon thermoelastic loss using FEA

The coated cantilever for which the losses are shown in Figure 4-1, exhibited significant curvature. The cantilever's radius of curvature (R) was measured (by Raymond Robie [135]) and this value was then used to calculate the stress (σ_c) of the coating via Stoney's equation [158] (which relates coated cantilever radius of curvature to the coating's stress):

$$\sigma_c = \frac{B_s a^2}{6tR} \ , \tag{4.13}$$

where B_s is the biaxial modulus of the cantilever, t is the thickness of the coating and a is the thickness of the cantilever. Using a graphical tool, the vertical displacement of the cantilever was then calculated using the value for radius of curvature. The radius of curvature, stress of coating and vertical displacement were determined to be: 0.115 m, 1.40×10^9 Pa and 5.1×10^{-3} m respectively. Using a static structural analysis in COMSOL and following the same method outlined in Section 3.6.2.1 (for ANSYS), the vertical displacement of the cantilever and stress of the coating were replicated computationally. The curvature of the coated cantilever was then exported into a thermoelastic loss analysis, the coating geometry deleted and a thermoelastic loss analysis ran on the uncoated curved cantilever. The results are shown in Figure 4-28 along with the FEA determined uncurved, unstressed thermoelastic loss of the uncoated cantilever and the measured coated cantilever losses. To ensure that the FEA modal losses of the uncurved, unstressed and uncoated cantilever replicated experimentally measured uncoated losses as closely as possible, each of the FEA bending mode losses were increased by 4 % whilst 0 % was added to the torsional modes (see Sections 3.4.1.2 or 4.4.3 for rationale). These results will be discussed in conjunction with those shown in Figure 4-29.

The thermoelastic loss of the second curved coated cantilever, for which the measured coated losses are shown in Figure 4-2, was also estimated using FEA. Unfortunately, this coated cantilever had been immediately heat treated after the loss measurements. This resulted in stress relief and a change in curvature and therefore it was not possible to obtain the stress or the vertical displacement prior to heat treatment. However, since this cantilever was approximately the same thickness, coated at the same time and with the same material as the first cantilever, it was assumed that both coatings had the same stress. Using Stoney's equation, the curvature of this coated cantilever was then estimated and its vertical displacement obtained. The COMSOL thermoelastic losses of the curved and uncoated cantilever, uncurved and uncoated cantilever, and the coated cantilever's measured losses are shown in Figure 4-29. Figure 4-29's bending mode losses of the uncurved and uncoated cantilever were increased by 4 % (see Sections 3.4.1.2 or 4.4.3 for rationale).

The first observation to be made about Figures 4-28 and 4-29 is that the FEA thermoelastic loss of the curved cantilevers is less than that of the uncurved cantilevers for a majority of resonant modes. This can perhaps be most clearly illustrated in Figure 4-30 which shows the ratio between the thermoelastic loss calculated for the curved cantilever and the thermoelastic loss for the straight cantilever, for each mode. Since a majority of the modes have a ratio of < 1, curvature can be seen to decrease a majority of modal thermoelastic losses. This decrease in thermoelastic loss for bending modes is particularly interesting, as it may explain why some of the measured losses of coated cantilevers are less than the thermoelastic loss of the uncurved and uncoated cantilevers. This is illustrated in Figure 4-31, which shows the ratio of measured coated loss to uncoated and uncurved thermoelastic loss, for each bending mode. Whilst the values of the ratios in Figures 4-30 and 4-31 are not directly comparable (since the measured coated cantilever ratio accounts for the loss of the coating and the other ratio does not) the trends of the different ratios can be compared. It can be observed that the FEA modelling is qualitatively replicating experimental evidence as the computational and *experimental* bending mode ratios follow a similar trend (where the lowest frequency modes have ratios > 1 and the higher modes < 1).

For the first cantilever (Figure 4-28), the effect of curvature reduces the thermoelastic loss to below the measured loss of the coated cantilever. As expected this new loss of the uncoated cantilever is now less than the loss of the coated cantilever (which consists of the sum of this reduced thermoelastic loss and the loss contribution from the coating). However, the COMSOL model does not fully explain the measured losses for the second cantilever. In this case, the curvature does not reduce the thermoelastic loss of every mode sufficiently (i.e. for some modes, the modelled loss of the cantilever substrate is still higher than the loss of the coated cantilever - an unphysical situational). This therefore implies that either the curvature of this cantilever was underestimated (which is possible as it was only estimated and not measured) or something else other than curvature is also contributing to the coated losses being so low in value.

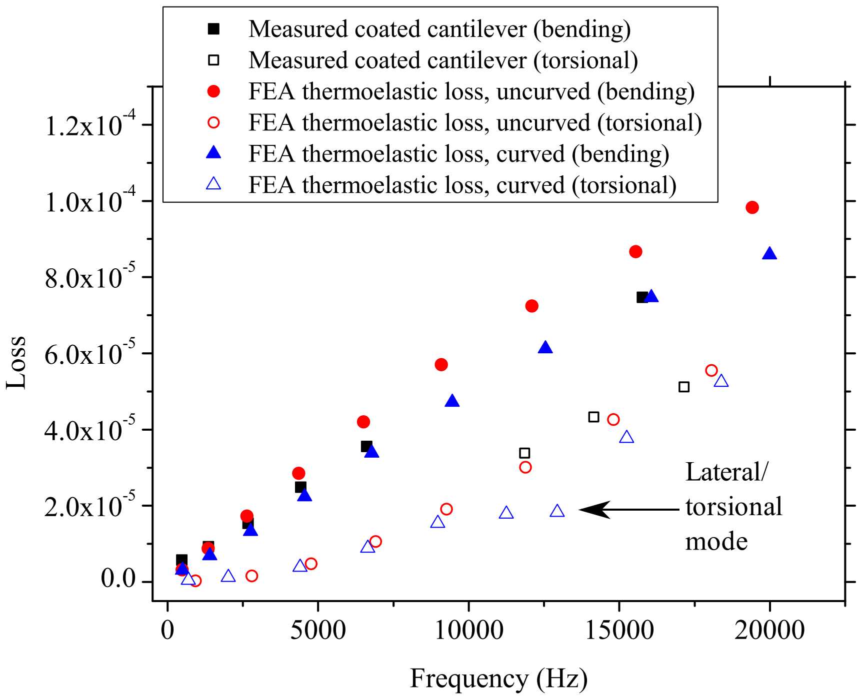

Figure 4-28: Measured mechanical losses of a 64×10^{-6} *m thick silicon cantilever coated with amorphous silicon. The thermoelastic loss of the uncoated and curved silicon cantilever is shown along with the thermoelastic loss of the uncoated and uncurved silicon cantilever. The thermoelastic loss was calculated using COMSOL. The FEA bending mode losses of the uncoated and uncurved silicon cantilever were increased by 4 % (as described in the text).*

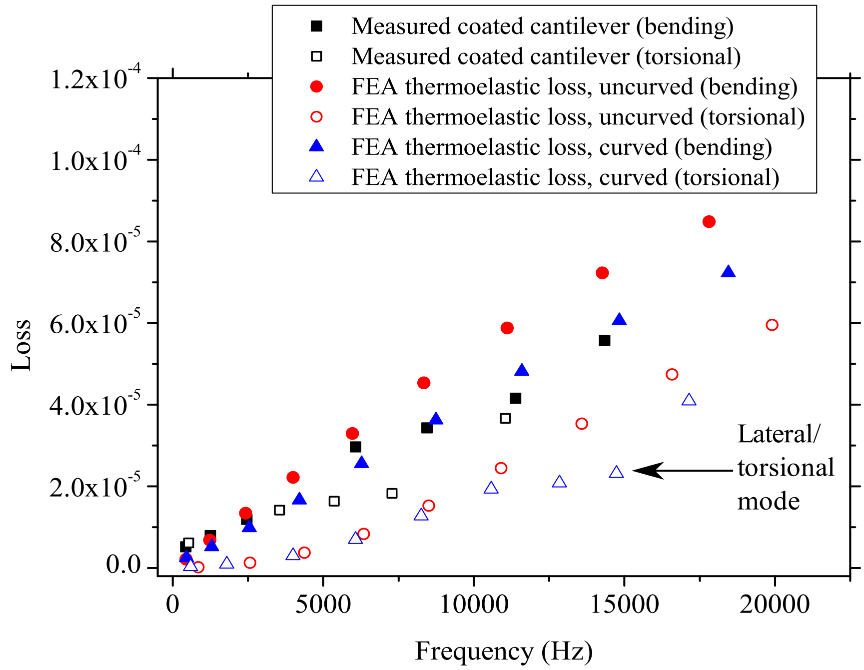

Figure 4-29: Measured mechanical losses of a 59×10^{-6} m thick silicon cantilever coated with amorphous silicon. The thermoelastic loss of the uncoated and curved silicon cantilever is shown along with the thermoelastic loss of the uncoated and uncurved silicon cantilever. The thermoelastic loss was calculated using COMSOL. The FEA bending mode losses of the uncoated and uncurved silicon cantilever were increased by 4 % (as described in the text).

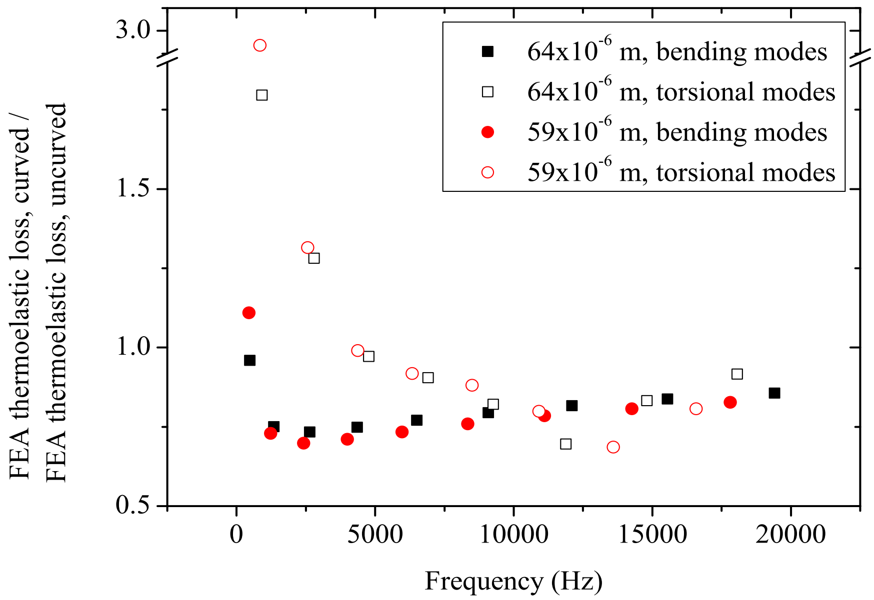

Figure 4-30: Ratio of each mode's curved FEA thermoelastic loss to its uncurved FEA thermoelastic loss. The ratio was calculated for both the 59×10^{-6} m and 64×10^{-6} m thick cantilevers respectively.

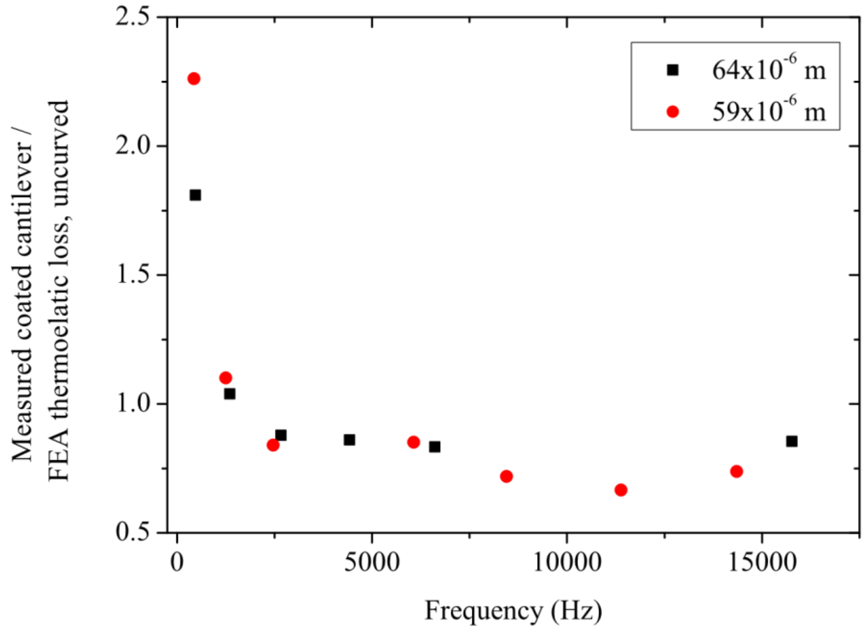

Figure 4-31: Ratio of each bending mode's measured coated cantilever loss to its uncoated and uncurved thermoelastic loss. The ratio was calculated for both the 59×10^{-6} m and 64×10^{-6} m thick cantilevers respectively.

An uncurved cantilever exhibits a lateral mode (not shown in Figures 4-28 and 4-29 but an example can be observed in Figure 4-27). When a cantilever is curved, the FEA modelling indicates that the lateral mode exhibits some twisting motion (like a torsional mode). This explains the extra mode of the curved cantilever (labelled "lateral/torsional mode" in Figures 4-28 and 4-29) when compared to its uncurved state. This quasi lateral/torsional mode exhibits a FEA thermoelastic loss which does not follow the trend of other torsional modes. Moreover, in the FEA models it appears that this mode couples to other torsional modes near it in frequency, lowering their thermoelastic loss (a similar effect was observed with the energy ratios in Figure 3-45 (b)). The curved cantilever's torsional modes between 10000-15000 Hz and 12500-15000 Hz in Figure 4-28 and Figure 4-29 represent the quasi torsional/lateral mode and the suspected coupled modes. Ignoring these specific losses, it can be generalised that the curvature of the substrate causes the thermoelastic losses of the bending modes to decrease more in absolute terms than torsional modes. This fits with the observations made by Pan et al [150] (discussed in the introduction) who saw an inversion of coated and uncoated bending mode losses when a highly stressed coating was applied to the cantilever but did not see the same effect for torsional modes.

It should be noted that when the *static structural* COSMOL model was used to replicate the measured vertical displacements of the cantilevers, the COMSOL model also curved the cantilever along its width (as would be expected from depositing a coating upon a cantilever at an elevated temperature and allowing to cool). The FEA thermoelastic losses of the curved cantilevers shown in Figures 4-28 and 4-29 therefore represent the losses of cantilevers which were curved along their width and length. Figure 4-32 illustrates the curvature along the width of the cantilever.

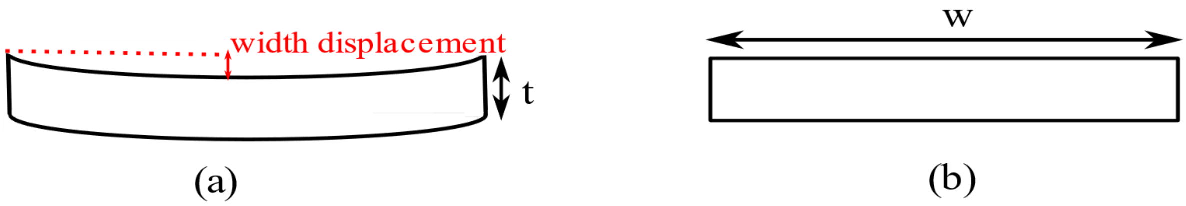

Figure 4-32: Illustration of the curvature along the width (w) of the cantilever. Image (a) represents the curvature at the end of the cantilever furthest from the clamping block. Image (b) illustrates the curvature of the cantilever at the end of the clamping block. Due to the size and weight of the clamping block, the face of the cantilever connected to the clamping block does not curve whilst the face of the cantilever furthest from the clamping block does. The curvature across the width gradually increases from the clamped end of the cantilever to the unclamped end. The letter t represents the thickness of the cantilever. The width displacement is the difference in height between the edge of a cantilever and its centre.

To investigate the effect of each type of curvature upon thermoelastic loss, cantilevers which were only curved either along their length or width were designed in COMSOL. It should be noted that the cantilevers modelled with only curvature across their width were designed to have a constant width curvature along their length due to FEA software limitations. Figure 4-33 shows that curvature along the width has a large effect on decreasing the thermoelastic loss of the bending modes and a small effect on decreasing the thermoelastic loss of the torsional modes. The more the cantilever is curved along its width, the greater the effect. Since curvature along the width increases the bending mode frequencies of the cantilever, it is suspected that the curvature along the width of the cantilever increases the stiffness of the cantilever along its length (and not along its width as it does not change the torsional mode's frequencies) resulting in the thermoelastic loss decreasing. It is not known why the lowest in frequency bending modes increase in thermoelastic loss due to curvature. Figure 4-34 shows that curvature along the length of the cantilever does not affect the cantilever's bending mode thermoelastic loss. Furthermore, there is no clear trend between curvature along the length of a cantilever and its torsional mode thermoelastic loss. At low frequencies, it appears that torsional mode losses may be decreasing as curvature increases but this is difficult to definitively conclude as the mode frequencies are also changing. At high frequencies, curvature does not seem to affect torsional mode thermoelastic loss or mode frequencies. At torsional mode frequencies in the middle of this range, the quasi lateral/torsional mode couples with other torsional modes lowering their thermoelastic loss. Since stiffness is known to change mode frequencies and the torsional mode frequencies of the cantilever can be observed to change with curvature, it is therefore hypothesised that curvature along the length of the cantilever is changing the stiffness of the cantilever resulting in a change in its thermoelastic loss. It should be noted that the FEA bending mode losses of the uncoated and uncurved cantilever in Figures 4-33 to 4-36 were not increased by 4 % as the losses in these figures were not being compared to experimental data but instead trends and patterns within the modelling were trying to be identified. Increasing the bending mode losses of the uncurved cantilever will act to skew these trends as can be seen in Figure 4-34, where if the 4 % were added, it would appear in the figure that COMSOL predicts that curvature along the length of the cantilever would decrease bending mode thermoelastic loss, when infact it does not.

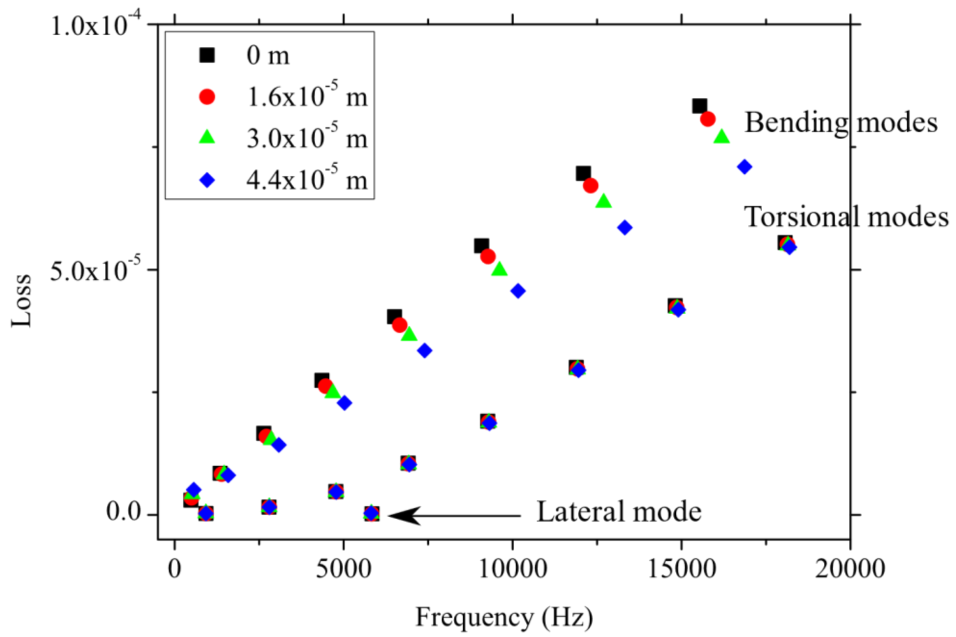

Figure 4-33: Effect of width curvature upon thermoelastic loss. Silicon cantilevers were curved with varying width displacements as listed in the legend. These width displacements were constant along the cantilever's length. The cantilever modelled was 64×10^{-6} m thick, 5×10^{-3} m wide and 34×10^{-3} m long.

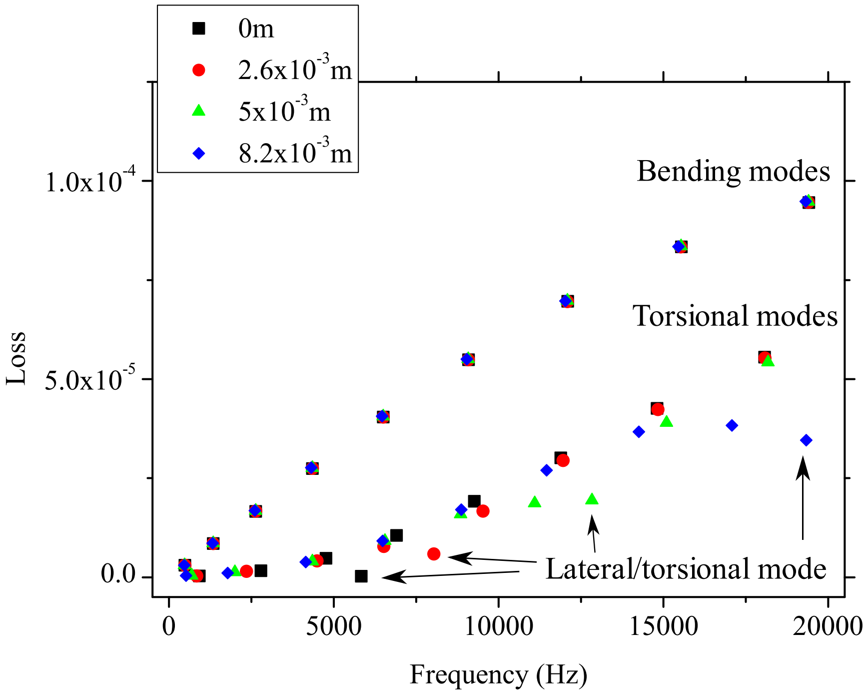

Figure 4-34: Effect of length curvature upon thermoelastic loss. Silicon cantilevers were curved with varying vertical displacements as listed in the legend. The cantilever was 64×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34×10^{-3} *m long.*

Figure 4-35 shows the effect of combining the curvatures along the cantilever's width and length upon thermoelastic loss. In this figure, the curvature along the width changes with length (as it does in reality). The width deflection in the legend was noted at the edge of the cantilever furthest from the clamp. The thermoelastic losses of the cantilever with a vertical displacement of 8.2×10^{-3} m and width deflection of 4.4×10^{-5} m are not shown due to a lack of convergence within the results.

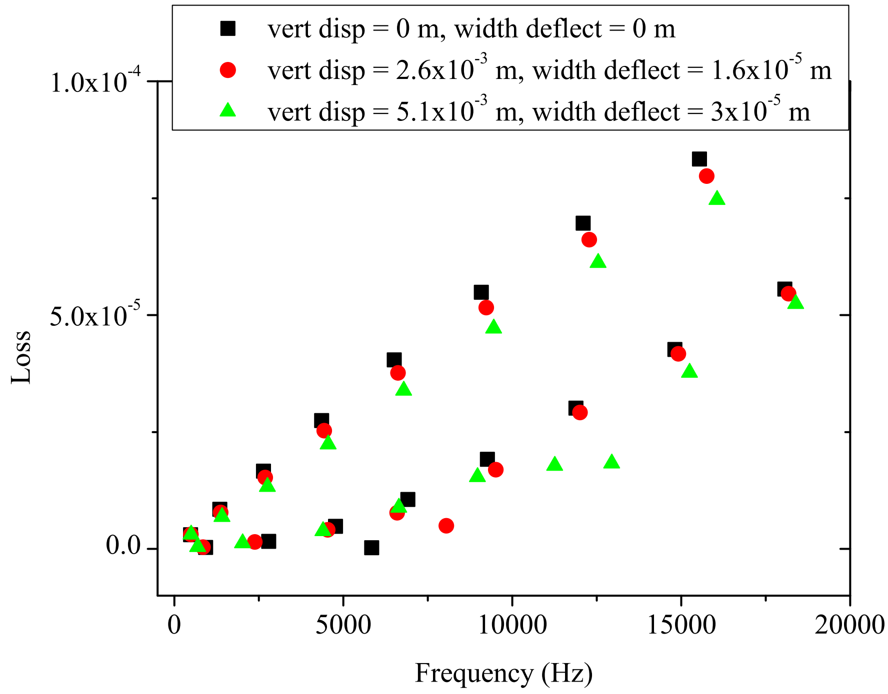

Figure 4-35: Effect of total curvature upon thermoelastic loss. Silicon cantilevers were curved with varying vertical displacements and width deflections as listed in the legend. The cantilever was 64×10^{-6} *m thick,* 5×10^{-3} *m wide and* 34×10^{-3} *m long.*

This figure shows that the FEA combined curvature model simply combines the losses of the separately curved cantilever models. Figure 4-36 shows the loss of the four highest frequency bending modes of the curved cantilevers shown in Figure 4-35, plotted against the width deflection (wd). Each mode's loss was regressed against its width deflection. The lines of best fit along with the equations which describe the fits are also shown in the figure. All of the fitted lines had a R^2 value of greater than 0.97. It can be seen that an increase in the width deflection of the cantilever will reduce the cantilever's bending mode losses approximately linearly (for the width deflections tested). Furthermore, it can also be observed from the fit's equations, that FEA predicts that the higher in frequency bending modes decrease more with an increase in width curvature.

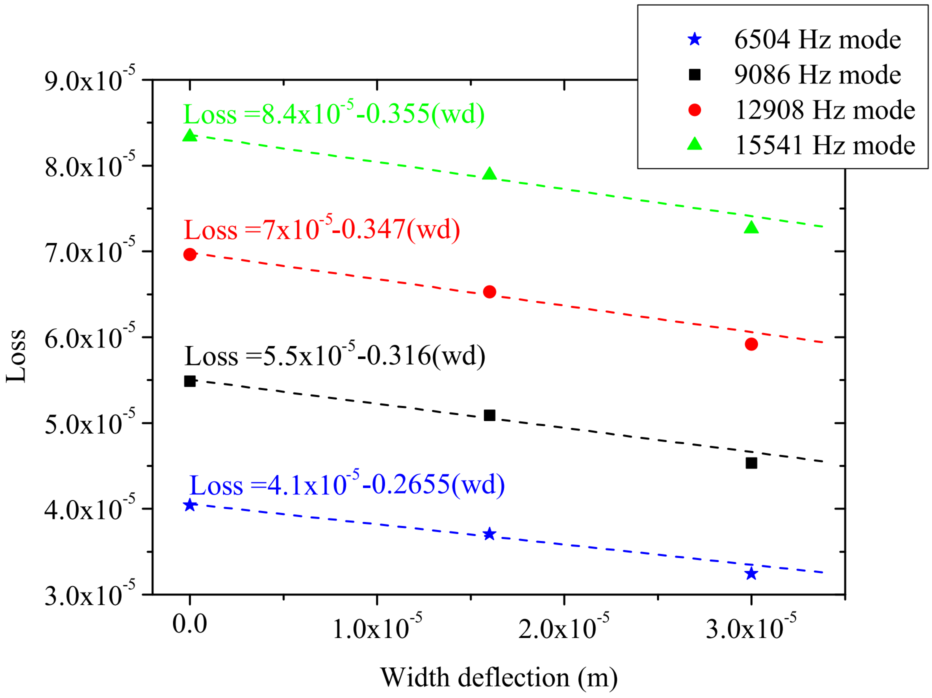

Figure 4-36: Loss of the four highest in frequency bending modes of the curved cantilevers shown in Figure 4-35 and their width deflections. Linear regressions and their equations of best fit are also shown in the figure.

To conclude, FEA modelling predicts that the curvature (caused by coating deposition) significantly changes the thermoelastic loss of a cantilever and can perhaps partially explain as to why measured coated losses from a curved coated cantilever are sometimes less than the uncoated and uncurved cantilever's losses.

4.6.2 Curvature effect upon coating loss

The COMSOL models potentially underestimate the effect on a cantilever's thermoelastic loss due to curvature (as shown in Figures 4-28 and 4-29). However even when using COMSOL's potentially underestimated values, the impact of cantilever curvature upon mechanical loss experiments is significant as the coating loss is proportional to the difference between the total loss of the coated cantilever and the loss of the cantilever substrate. Since curvature induced by the coating changes the loss of the cantilever substrate, it is clearly essential to correctly account for this in the coating loss measurements. This is best illustrated by a theoretical example.

Consider two coatings whose bending mode losses have been calculated when they were deposited upon silica cantilevers (which are not dominated by thermoelastic loss and will not therefore be affected by any curvature effects). Assume that the first coating (coating-1) was measured to have a bending mode coating loss of 4×10^{-4} whilst the second coating (coating-2) had a bending mode coating loss of 3.5×10^{-4}. Now suppose these coatings were deposited upon two identical silicon cantilevers with a thickness of 64×10^{-6} m. Assume that depositing coating-1 onto a silicon cantilever resulted in the coated cantilever to have a vertical displacement of 1.1×10^{-3} m (a typical vertical displacement due to coating deposition) and a corresponding maximum width deflection of 0.5×10^{-5} m. Now assume that depositing coating-2 resulted in a zero vertical displacement and zero width deflection. As a result of the coated cantilever curving due to coating-1, the loss of the cantilever is now approximately 6 % less than it's previously hypothetically measured uncoated loss at a bending mode of 9086 Hz (2 % predicted from Figure 4-36 and an extra 4 % due to COMSOL underestimating the bending modes of an uncurved cantilever's losses). Using the hypothetically measured uncoated and coated losses, the loss of coating-1 on the silicon cantilever was determined to be 3.2×10^{-4}. The bending mode loss of coating-2 is the same for the silicon cantilever as that measured on the silica cantilever (as the coated silicon cantilever does not exhibit curvature). Based upon the losses measured on the silica cantilevers, coating-1 has significantly higher loss than coating-2 whilst the opposite is true when using the results measured on the silicon cantilevers. This is of course contradictory and could lead to a lot of problems in identifying the optimum coating to use. However, by accounting for the effect of curvature upon the hypothetically measured uncoated loss and then calculating the loss of the coating, it can be observed that the loss of the coatings determined upon both of the substrates is now equal. Thus solving the problem. These calculations are summarised in Table 4-2.

This hypothetical example illustrates the significance of accounting for curvature affects upon thermoelastic loss when determining the loss of a coating deposited upon a silicon cantilever. The fact that the modelling potentially underestimates the reduction in thermoelastic loss (of a bending mode at 9086 Hz), further emphasises the importance of accounting for curvature effects. To conclude, loss measurements of coatings should not be made using curved silicon cantilevers unless the thermoelastic loss of the uncoated, curved cantilever can be accurately calculated.

	Silica cantilever with coating-1	Silica cantilever with coating-2	Silicon cantilever with coating-1	Silicon cantilever with coating-2
Hypothetically measured uncoated loss	N/A	N/A	5.5×10^{-5}	5.5×10^{-5}
Hypothetically measured coated loss	N/A	N/A	6.89×10^{-5}	7.02×10^{-5}
Energy Ratio (U_s/U_c) which account for stress and curvature	N/A	N/A	23.25	23
Calculated bending mode coating loss (using hypothetically measured uncoated and coated loss)	N/A	N/A	3.2×10^{-4}	3.5×10^{-4}
True uncoated loss (which accounts for curvature effect upon thermoelastic loss)	N/A	N/A	5.17×10^{-5}	5.5×10^{-5}
True bending mode coating loss (calculated using true uncoated loss and hypothetically measured coated loss)	4×10^{-4}	3.5×10^{-4}	4×10^{-4}	3.5×10^{-4}

Table 4-2: Theoretical calculation illustrating the importance of accounting for the effect of curvature upon thermoelastic loss. Two coatings (coating-1 and coating-2) were deposited onto one silicon and silica cantilever each. The loss of the coatings deposited upon the silica cantilevers was calculated and the loss of coating-1 was determined to be greater than that of coating-2. When the loss of the coatings deposited upon the silicon cantilevers was calculated, without accounting for the reduction in the uncoated loss (thermoelastic loss) of the silicon cantilever due to curvature, the loss of coating-2 was determined to be larger than the loss of coating-1. By accounting for the effect of curvature upon the uncoated loss of the silicon cantilevers, the loss of the coatings were determined to be the same as that measured on the silica cantilevers.

4.6.3 Experimentally verifying computational results

The results from the modelling indicate that stress in a coated cantilever does not change thermoelastic loss but that curvature does. An experiment was designed to verify these results.

The concept of the experiment was to coat one face of varyingly thick silicon cantilevers with a highly stressed coating (1×10^{-6} m thick) and coat another silicon cantilever on both faces with half the thickness of the same coating (5×10^{-7} m thick each side). Whilst all cantilevers would exhibit large stresses, the curvature of the coated cantilevers would be very different (the cantilevers coated on one side would curve by different amounts whilst the cantilever coated on both sides would not curve). By knowing the uncoated loss of the cantilevers (which is approximately its uncurved thermoelastic loss) and the bulk and shear losses of the coating (which had been previously measured upon a silica disc, that is unaffected by stress and curvature thermoelastic loss effects), an expected loss of the coated cantilevers can be determined under the assumption that curvature does not alter the thermoelastic loss of the cantilever. If the thermoelastic loss of a cantilever is indeed affected by curvature and not by stress, then the measured loss of the uncurved cantilever coated on both sides should agree with this expected coated loss (which will be referred to as the predicted uncurved coated loss). However, for the cantilevers coated on only one side, which do curve, the ratio between their predicted uncurved coated loss and measured loss should be > 1 (for modes which have a ratio of curved thermoelastic loss to uncurved thermoelastic loss of < 1) as the thermoelastic loss of the coated cantilever substrates is expected to decrease due to curvature resulting in the measured coated loss being less than the predicted uncurved coated loss.

Four cantilevers were coated with a highly stressed silica coating[1]. 3 cantilevers (50.5×10^{-6} m, 62×10^{-6} m and 90×10^{-6} m in thickness) were coated on one side with 1×10^{-6} m of silica. The vertical displacements of each coated cantilever were measured. By replicating the coated cantilevers' vertical displacements in a *static structural* COMSOL

[1] It was originally planned that the silicon cantilevers would be coated with amorphous silicon. This would have enabled a direct comparison to the previously discussed silicon cantilevers coated with amorphous silicon (their losses are shown in Figures 4-28 and 4-29). However due to a mistake by the coating vendor, a silica coating was deposited onto the cantilevers instead of a silicon coating.

model, the compressive stress of the coating was determined to be approximately 440 MPa. Using graphical software and knowing the geometries and vertical displacements of the three cantilevers, each of the cantilever's radius of curvature was estimated. Each cantilevers' radius of curvature was then plotted against their thickness. Regressing all of the cantilever radius of curvatures against their thicknesses using a squared function, an excellent fit can be observed, as shown in Figure 4-37. This excellent fit implies agreement with Stoney's theory. Using Stoney's equation (Equation (4.13)) and the estimated variable (8.4×10^7) obtained from the fit, the stress of the coating can be determined to be 460 MPa. This shows good agreement with the independent COMSOL method (described earlier) which estimated coating stress (to be 440 MPa).

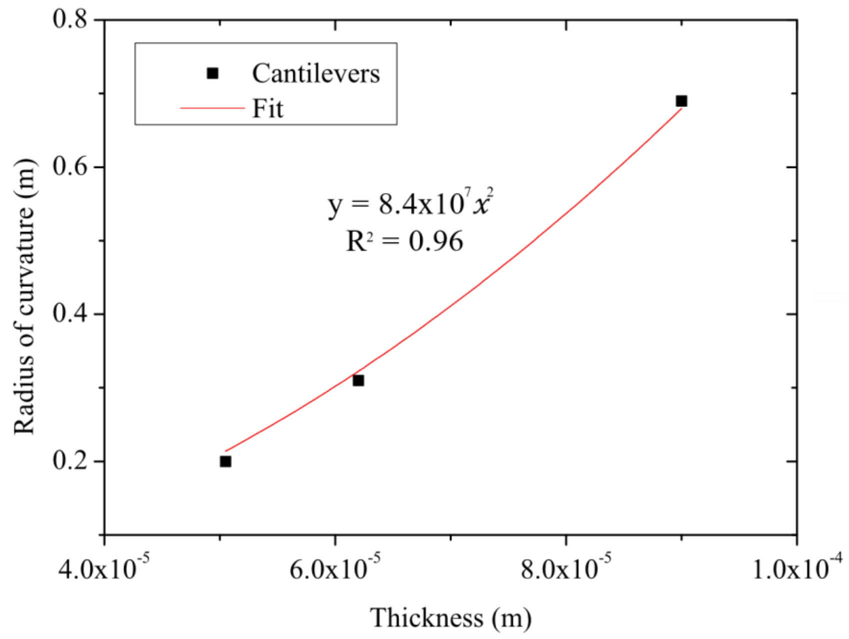

Figure 4-37: Radius of curvature and thickness of the three differently thick cantilevers which had been coated with 1×10^{-6} m of silica upon one side. The points are fitted by a squared function and the good fit implies agreement with Stoney's equation (Equation (4.13)). The equation for the line of best fit and its R^2 value are included in the figure.

The fourth cantilever coated (which was 74×10^{-6} m thick), had 5×10^{-7} m of silica deposited upon each of its sides. The losses of all four coated cantilevers were measured (using the setup described in Section 3.3.1.2). The uncoated loss of the uncurved 50.5×10^{-6} m, 74×10^{-6} m and 90×10^{-6} m thick cantilevers was estimated by calculating their thermoelastic loss using COMSOL and amending the calculated losses as described in Sections 3.4.1.2 and 4.4.3 (adding 4 % to the bending mode thermoelastic loss values and 0 % to the torsional mode thermoelastic loss values). The uncoated loss of the uncurved 62×10^{-6} m thick cantilever was measured pre-coating. As an example of the measured and calculated losses, Figure 4-38 shows the measured coated loss, calculated uncoated loss of the uncurved cantilever and the predicted uncurved coated loss for the 50.5×10^{-6} m thick cantilever. The measured coated losses shown in Figure 5-38 and used throughout this section have subtracted the loss mechanism caused by the mismatch in the coating's and substrate's materials, known as Fejer coating thermoelastic loss [109] (Section 3.6.4), to ensure fair comparison with the predicted coated losses which do not account for this loss mechanism (in their calculation). Each mode's Fejer coating thermoelastic loss represented less than 2 % of the measured mode's coated loss.

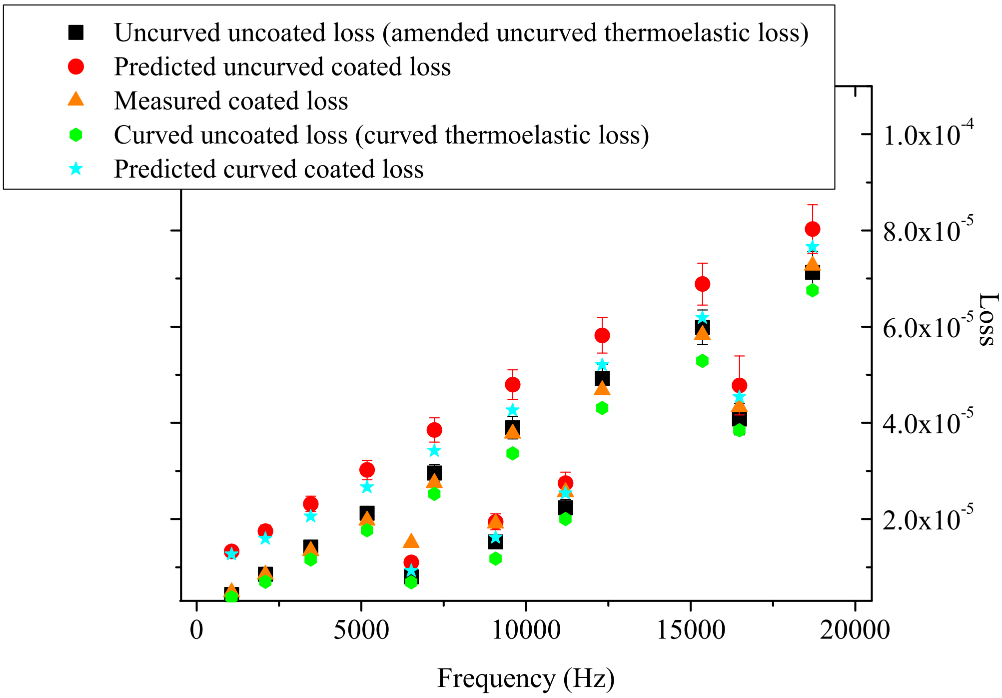

Figure 4-38: Uncurved uncoated loss, measured coated loss, curved uncoated loss, the predicted uncurved coated loss and the predicted curved coated loss for the 50.5×10^{-6} *m thick silicon cantilever. The uncurved uncoated loss was calculated using COMSOL (plus an additional 4 % loss to the bending modes) and is equal to the uncurved cantilever's thermoelastic loss. The measured coated loss was the coated loss of the coated cantilever measured in the lab minus Fejer loss. The curved uncoated loss is equal to the thermoelastic loss of the curved uncoated cantilever as calculated by COMSOL. The predicted uncurved coated loss was calculated using the uncurved uncoated loss of the cantilever whilst the predicted curved coated loss was calculated using the curved uncoated loss of the cantilever.*

Since the predicted uncurved coated loss is always greater than the measured coated loss for the cantilever's bending modes, then the ratio of predicted uncurved coated loss to measured coated loss should be greater than 1 for the 50.5×10^{-6} m thick cantilever's bending modes. For the cantilever's torsional modes, it can be observed that the predicted uncurved coated loss is sometimes above and sometimes below the measured coated loss. The ratio of predicted uncurved coated loss to measured coated loss for the cantilever's torsional modes should therefore be more similar to 1 than that obtained for the bending modes. These ratios are shown in Table 4-3 for the cantilevers of different thickness. This ratio was calculated separately for bending and torsional modes. Table 4-3 also lists the ratios for the 64×10^{-6} m and 59×10^{-6} m thick cantilevers (discussed in Section 4.6.1) which were also coated on one side with a 1×10^{-6} m thick layer of amorphous silicon.

Cantilever thickness (m)	Vertical displacement (m)	Ratio of predicted uncurved coated loss to measured coated loss (bending modes)	Ratio of predicted uncurved coated loss to measured coated loss (torsional modes)
50.5×10^{-6}	3.0×10^{-3}	1.59 ± 0.19	0.98 ± 0.11
59×10^{-6}	5.8×10^{-3}	1.42 ± 0.12	0.88 ± 0.11
62×10^{-6}	2.0×10^{-3}	1.27 ± 0.07	1.06 ± 0.03
64×10^{-6}	5.1×10^{-3}	1.32 ± 0.12	1.08 ± 0.11
90×10^{-6}	1.0×10^{-3}	0.96 ± 0.07	1.12 ± 0.10
74×10^{-6}	0	1.08 ± 0.09	0.95 ± 0.08

Table 4-3: List of the cantilever thicknesses, vertical displacements and their values for the ratio of predicted uncurved coated loss to measured coated loss. This ratio was calculated separately for the bending and torsional modes of the cantilevers.

The vertical displacement of the cantilevers coated with amorphous silicon (64×10^{-6} m and 59×10^{-6} m) do not follow the trend of the other cantilevers, coated with silica, due to the higher stress of the amorphous silicon. This difference in trend between the silicon cantilevers coated with silica and amorphous silicon can also be observed with the ratios as illustrated in Figure 4-39 which shows the vertical displacement and ratio of predicted uncurved coated loss to measured coated loss for the bending modes. It can be observed that due to different losses of the coating (or perhaps other effects) the ratios of the silicon cantilevers coated with silica and amorphous silicon are not directly relatable to one another as they appear to follow different trends.

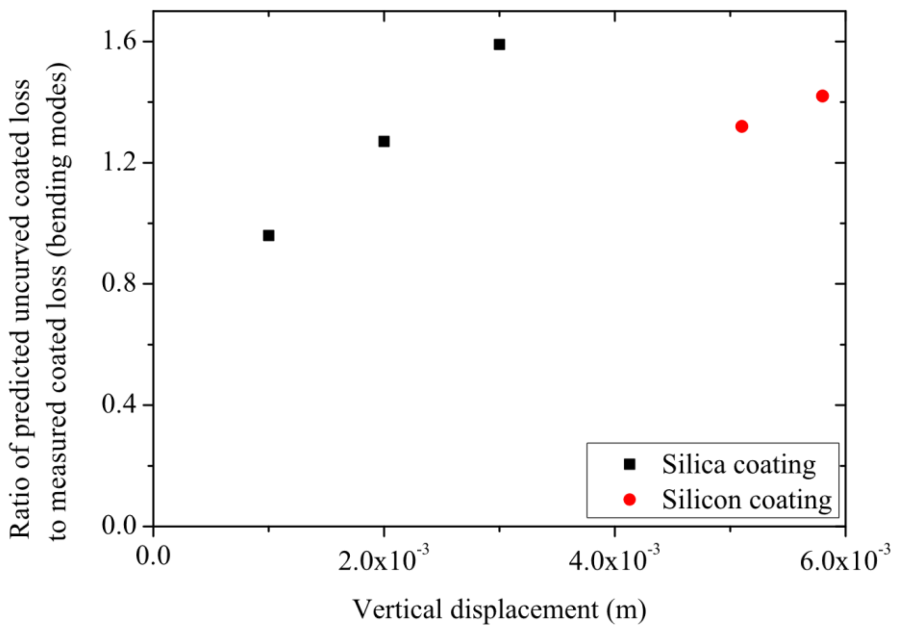

Figure 4-39: Ratio of the cantilevers predicted uncurved coated loss to measured coated loss as well as their vertical displacements. It can be observed in the figure that the silicon cantilevers coated with a silica coating and those coated with the silicon coating follow different trends.

It can be observed in Table 4-3 that the bending mode ratio is equal to 1 (within error) for the cantilever coated on both sides. This implies that the stress of the coated cantilever is not affecting the substrate's bending mode thermoelastic loss as the measured loss can be accurately predicted by assuming an uncurved uncoated loss for the cantilever. All of the other cantilever bending mode ratios are greater than 1 (as expected from the modelling), except the 90×10^{-6} m thick cantilever where the ratio is within error of 1. This is the least curved cantilever and therefore the least affected by curvature effects upon thermoelastic loss.

It can be observed in Table 4-3 that the torsional mode ratio is equal to 1 (within error) for the cantilever coated on both sides. This again implies that the stress of the coated cantilever is not affecting the substrate's torsional mode thermoelastic losses as the measured losses can be accurately predicted. However, the story for the torsional mode ratios of the other cantilevers is unclear as some are greater than 1 (as expected from the modelling) but others are equal to and less than 1 (which is not expected). The reason for these observations are unknown. However, it could be speculated that variations in coating thickness may result in twisting of the coated cantilever and that twisting of the coated cantilever may affect the substrate's thermoelastic loss for torsional modes as its mimics its resonant motion.

Another test of the COMSOL modelling is to directly compare the measured loss of the coated cantilevers to the predicted curved coated loss which is calculated using the curved thermoelastic uncoated losses. To compare these, the ratio of the quantities was calculated. If this ratio is closer to 1 than the previously calculated ratio of predicted uncurved coated loss (calculated using the uncurved uncoated loss) to measured coated loss (as shown in Table 4-3), then accounting for curvature effects upon thermoelastic loss enables greater prediction of the coated cantilever's measured losses. This would imply that the curvature effect upon thermoelastic loss is real. Figure 4-38 shows the uncoated loss of the curved cantilever, the predicted curved coated loss calculated using the curved thermoelastic loss values to represent the uncoated loss, the predicted uncurved coated loss calculated when uncurved thermoelastic loss values represent the uncoated loss and the measured coated losses of the 50.5×10^{-6} m thick cantilever. The predicted curved coated losses, which assume a curved thermoelastic loss as the cantilever uncoated loss, replicate the measured coated losses better than the predicted uncurved coated losses, which assume an uncurved thermoelastic loss as the cantilever uncoated loss, at bending modes but not for torsional modes. The ratio of predicted curved coated loss to the measured coated loss will therefore be closer to 1 than the ratio of predicted uncurved coated loss to the measured coated loss for the bending modes of this cantilever. However, the ratio of predicted curved coated loss to the measured coated loss will be further from 1 than the ratio of predicted uncurved coated loss to the measured coated loss for the torsional modes of this cantilever. These results are shown in Table 4-4 which lists the ratios of predicted curved coated loss to measured coated loss as well as the ratio of predicted uncurved coated loss to measured coated loss for all of the cantilevers. No errors are able to be calculated for the ratio of predicted curved coated loss to measured coated loss because no error can be quantified for the predicted curved coated loss as it is unknown how similar the curved thermoelastic loss values calculated by COMSOL replicate reality.

Cantilever thickness (m)	Ratio of predicted curved coated loss to measured coated loss	Ratio of predicted uncurved coated loss to measured coated loss
50.5×10^{-6}	1.45	1.59
59×10^{-6}	1.13	1.42
62×10^{-6}	1.19	1.27
64×10^{-6}	1.10	1.32
90×10^{-6}	0.91	0.96

(a) bending modes

Cantilever thickness (m)	Ratio of predicted curved coated loss to measured coated loss	Ratio of predicted uncurved coated loss to measured coated loss
50.5×10^{-6}	0.88	0.98
59×10^{-6}	0.81	0.88
62×10^{-6}	1.02	1.06
64×10^{-6}	0.90	1.08
90×10^{-6}	1.03	1.12

(b) torsional modes

Table 4-4: List of cantilever thicknesses, ratio of predicted curved coated loss to measured coated loss and the ratio of predicted uncurved coated loss to measured coated loss. The ratios were calculated separately for the bending (a) and torsional (b) modes of the differently thick cantilevers.

In Table 4-4 (a), it can be observed that the ratio which accounts for curvature effects upon thermoelastic loss at bending modes is closer to one than the ratio which does not account for curvature effects, for four of the five cantilevers. This suggests that the curvature effect upon a bending mode's thermoelastic loss is real and by accounting for it in our predicted curved coated loss calculation, the measured losses can be more accurately determined (in a majority of cases). This can be most clearly observed in Figure 4-40 which shows the two bending mode ratios against cantilever thickness. In this figure it is also clear that the more curved a cantilever is, the greater the difference in value between the two ratios (as the effect of curvature upon thermoelastic loss is larger).

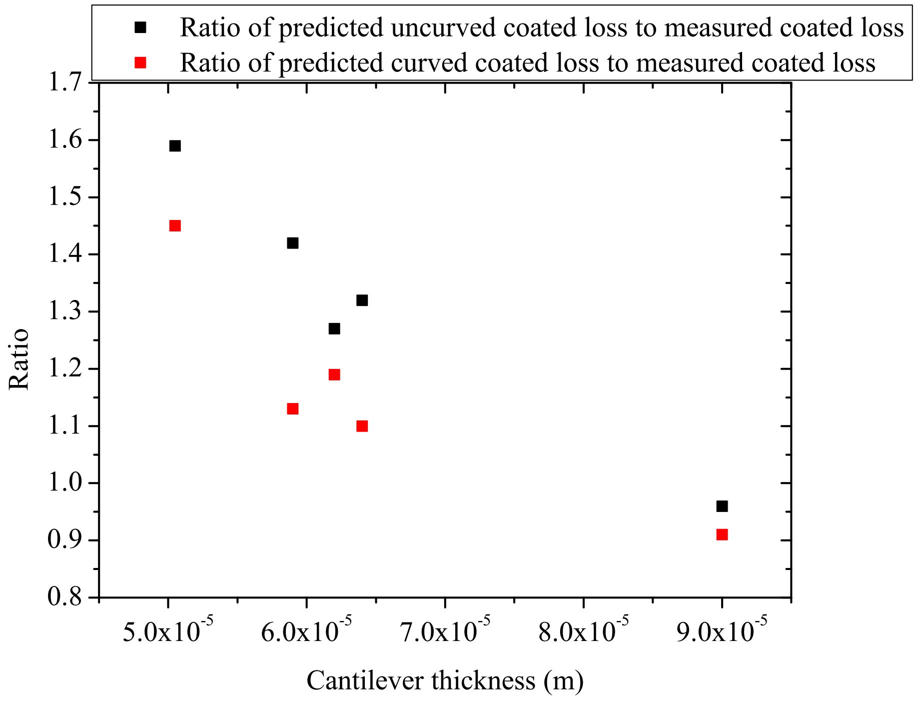

Figure 4-40: Ratio of the cantilevers' predicted uncurved coated loss to measured coated loss, against cantilever thickness. Plotted alongside are the ratios of the cantilevers predicted curved coated loss to measured coated loss. It can be observed in the figure that accounting for the effect of curvature upon thermoelastic loss helps to explain some of the difference between predicted uncurved coated loss and measured coated loss (for four of the five cantilevers).

However, it should be noted that the ratio which accounts for curvature is not equal to 1 for any of the cantilevers. This is important as if the effect of curvature alone upon thermoelastic loss explains the measured losses (and the COMSOL models accurately calculate the effect of curvature upon a bending mode's thermoelastic loss) then this ratio should be equal to 1. The fact that it is not equal to 1, hints that either an unmodeled property is affecting thermoelastic loss (D-shape of the cantilever) or the modelling is underestimating the effect of curvature upon thermoelastic loss.

When considering torsional modes (Table 4-4 (b)), the modelling of curvature is less successful in predicting the loss of the coated cantilevers. Only for two cantilevers does the curvature modelling result in a better prediction of the measured coated loss (i.e. curved ratio is closer to 1 than the uncurved ratio). This suggests that the modelled curvature effect upon a torsional mode's thermoelastic loss may not be real as accounting for its effect results in worse prediction of the measured losses for a majority of the cantilevers. In the future, the effect of a slightly twisted cantilever curvature (caused by coating thickness variations) upon a cantilever torsional mode's thermoelastic loss should be investigated as it could perhaps explain why the results are worse when curvature is accounted for.

It should be noted that in this subsection, coating/substrate energy ratios which account for stress and curvature effects were used in all calculations to determine the predicted uncurved and curved coated losses. Coating/substrate energy ratios which accounted for stress and curvature were used to determine the predicted uncurved coated losses as well as the predicted curved coated losses so that the only difference between the two set of values was the assumption of the cantilever's thermoelastic loss (curved or uncurved thermoelastic loss). By only varying the assumption of the substrate's thermoelastic loss, the effect of assuming a curved or uncurved thermoelastic loss upon replicating the measured losses could be identified. However, it should be noted that using unstressed and uncurved coating/substrate energy ratios had less than a 3 % effect upon the ratios quoted in Tables 4-3 and 4-4. Whilst it is known how a coated cantilever's vertical displacement affects the coating's bulk and shear energies (Section 3.6.3), it is unknown how width displacement affects them or how the combination of the two affect them. Furthermore, it is not known how stress affects bulk and shear losses. For these reasons, the bulk and shear energy ratios of uncurved and unstressed coated cantilevers were used in the calculations to determine the ratios presented in Tables 4-3 and 4-4.

To conclude, the experimental results imply that stress does not have an effect upon a cantilever's bending and torsional modes' thermoelastic losses but curvature (as modelled in COMSOL) might have an effect upon a cantilever's bending modes. Furthermore, the experimental results indicate that the curvature modelled in COMSOL does not fully explain the bending modes' measured losses and perhaps something else is also affecting the bending modes' thermoelastic loss.

4.7 Thermoelastic loss of a disc

Silicon discs are another kind of substrate (geometry) used in mechanical loss experiments. The mechanical loss of these substrates at room temperature are dominated by thermoelastic loss. It is therefore important to identify the effect of how a silicon disc's thermoelastic loss changes as a result of the curvature and stress caused by depositing a coating on its surface.

Before the effect of stress and curvature (caused by coating deposition) on a disc's thermoelastic loss can be investigated, a FEA model which is able to replicate experimental values as well as values from other models in the literature must first be built. By showing that an FEA model can replicate these models as well as experimental data, will give greater confidence in the results of the more complicated modelling required to investigate the effects of stress and curvature upon thermoelastic loss.

4.7.1 FEA model of disc

An FEA model can be shown to accurately determine thermoelastic loss when its values can replicate experimental losses as well as the values from other models in the literature.

The mechanical loss of a silicon disc with diameter 75.94×10^{-3} m and thickness 466.7×10^{-6} m was measured at room temperature. The disc was supported from its centre using a nodal support. The losses were measured by Lorenzini and published in the literature [159]. The thermoelastic loss of the disc was calculated using a FEA method developed at Jena university [159]. Using COMSOL method 2, the thermoelastic losses of the disc were estimated. The results are shown in Figure 4-41.

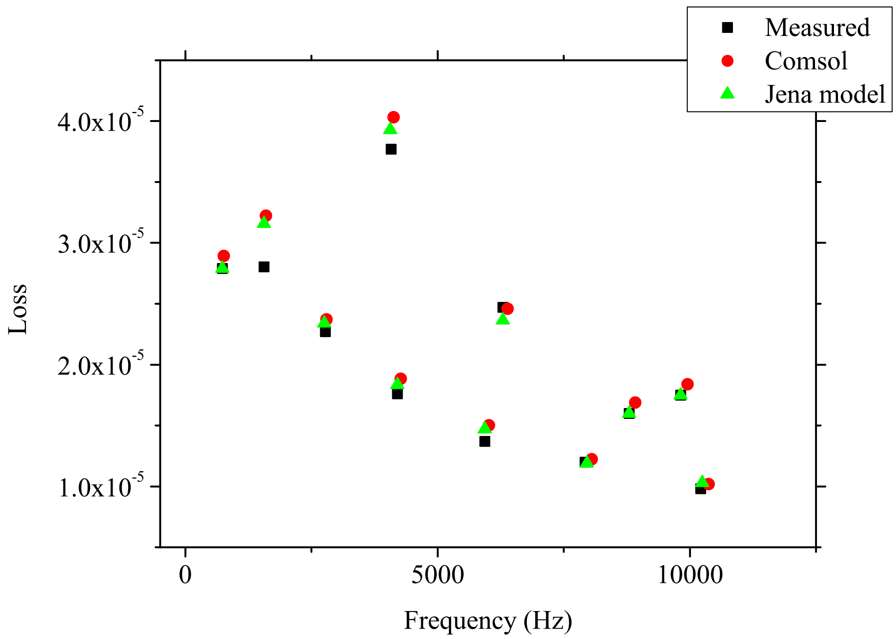

Figure 4-41: Calculated thermoelastic loss and measured mechanical loss of a silicon disc with diameter 75.94×10^{-3} *m and thickness* 466.7×10^{-6} *m.*

The experimental loss and COMSOL's thermoelastic loss approximately match one another (with an average difference of 5 %). Since the experimental losses of the disc were measured by Lorenzini [159], there was no access to the raw data. It was therefore not possible to investigate the mode losses which showed the greatest difference from the COMSOL model i.e. the second mode shown. No access was available to the disc either or to the errors in the disc's dimensions which would have been useful as it would have enabled a range of COMSOL thermoelastic loss values to be calculated. This would have been particularly useful as the COMSOL values are always overestimating the measured thermoelastic loss values implying that one of the disc's dimensions could be inaccurately specified. The average difference in loss between COMSOL method 2's values and Jena's was less than 2.5 % for all modes. To conclude, Figure 4-41 shows that COMSOL is on average able to replicate both the measured loss and other computational thermoelastic loss values in the literature within 5 %.

4.7.2 Stress and thermoelastic loss of a disc

The same argument which was made to investigate the effect of stress upon thermoelastic loss in the coated cantilever can again be made for a coated disc. When a coated disc bends, both the coating and disc bend together. The amount of bending the coated disc experiences is therefore dependent on both the stress in the coating and the stress in the disc. As a result, both the stress in the coating and the stress in the disc must be incorporated within a coated disc model when its thermoelastic loss is calculated. A coated disc at rest is in equilibrium. The stresses in the coating and the disc must therefore balance one another. Figure 4-42 shows three different stress patterns applied to the coated disc: no stress in either the coating or disc, a compressive stress in the substrate and no stress in the coating, and a compressive stress in the substrate and a tensile stress in the coating. When the disc and coating were stressed with a compressive and tensile stress respectively, the magnitudes of the stresses applied were specifically chosen such that their forces balanced one another. In Figure 4-42 it can be observed that when the stresses balance one another (radial stress of -5×10^6 Pa in the substrate and 2.8×10^9 Pa in the coating), there is no change in the thermoelastic loss of the coated disc compared to the unstressed coated disk. Therefore, it can be concluded that the effect of stress upon thermoelastic loss will not affect the accuracy of coating losses determined using mechanical loss experiments.

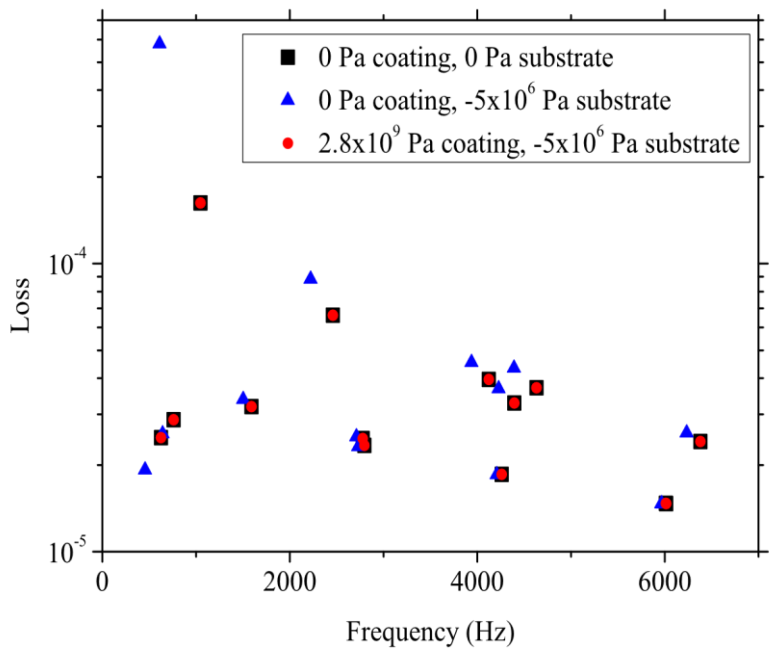

Figure 4-42: Varyingly stressed, coated, silicon disc's thermoelastic loss. The stress in the disc and coating was acting radially. The silicon disc had a diameter 80×10^{-3} *m and thickness* 500×10^{-6} *m. The amorphous silicon coating was* 1×10^{-6} *m thick. Positive stress represents tensile stress whilst negative stress represents compressive stress.*

4.7.3 Curvature and thermoelastic loss of a disc

When a coating is deposited upon a disc at an elevated temperature and the body is allowed to cool to room temperature, stresses develop in the coated disc which cause it to deform into a shape which resembles a spherical cap. This is illustrated in Figure 4-43.

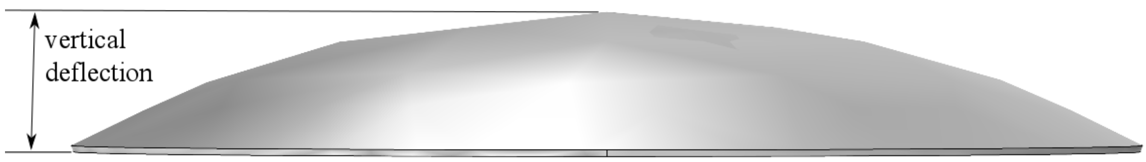

Figure 4-43: Curvature of a coated disc. Initially flat, the disc is bent into a spherical cap as a result of the coating deposition process. The vertical deflection of the disc refers to the vertical distance between the peak of the spherical cap and its edges.

Since the disc used in mechanical loss experiments is of much greater thickness than the coating, any deformation of the coated disc is extremely small (and not noticeably deformed). The thermoelastic loss of varyingly curved, uncoated discs was calculated and shown in Figure 4-44. The discs were curved using a static structural analysis. In the static structural analysis a coated disc was stressed and curved by cooling the coated substrate from an elevated temperature. As a result of the coating and disc having different thermal expansion coefficients the coated disc became stressed and curved. The geometry of the curved coated disc was then fed into a thermoelastic loss analysis, the coating geometry was deleted and the thermoelastic loss of the disc calculated. The silicon disc investigated had a diameter of 80×10^{-3} m and thickness 500×10^{-6} m. The value of curvature quoted in Figure 4-44 refers to the vertical deflection of the disc, i.e. the vertical distance between the peak of the spherical cap and its edges. The larger the value, the more curved the disc. A disc of these dimensions which has been coated with a 1×10^{-6} m thick coating of stress -300 MPa [54] (a typically measured coating stress) would exhibit a vertical displacement of approximately 40×10^{-6} m. It can be observed in Figure 4-44 that curving the disc (within 100×10^{-6} m) had no effect upon the disc's thermoelastic loss.

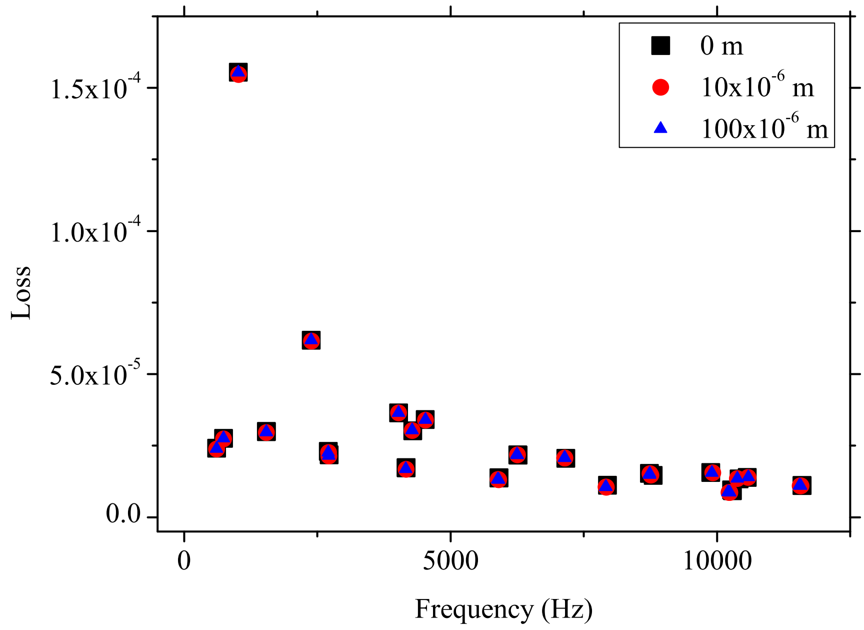

Figure 4-44: Calculated thermoelastic loss of an unstressed and vertically displaced silicon disc with diameter 80×10^{-3} *m and thickness* 500×10^{-6} *m.*

4.8 Conclusion

Depositing a coating upon a silicon cantilever results in the coated cantilever exhibiting stress and curvature. The effect of this stress and curvature upon a silicon cantilever's thermoelastic loss was investigated.

At the beginning of this chapter different methods of calculating thermoelastic loss were developed (ANSYS and COMSOL), tested (Zener and Liftshitz and Roukes) and compared to experimental values. It was observed that the theoretical models could not replicate the measured values as well as the FEA models due to mathematical assumptions which did not enable the theoretical models to capture a cantilever's width dependent oscillations. These more accurate FEA models formed the basis of the in-depth analysis into the effects of stress and curvature upon thermoelastic loss.

Experimental results have shown that the mechanical loss of some coated silicon cantilevers is less than their uncoated cantilever loss. In the literature, it has been speculated that this could be due to the curvature or stress of the coated cantilever reducing the cantilever's thermoelastic loss. FEA modelling within this chapter qualitatively showed that this reduction in coated loss compared to uncoated loss is most likely the effect of cantilever curvature and not stress affecting thermoelastic loss.

An experiment was designed to test the predictions of the FEA modelling. The results from the experiment partially agreed with the modelling as they both suggested that the stress caused by coating deposition does not affect the thermoelastic loss of a cantilever and that curvature most likely decreases a cantilever's bending mode thermoelastic loss. However, the experimental results and modelling disagree upon the effect of curvature on a torsional mode's thermoelastic loss. Overall, it is thought that the FEA modelling might not be accounting for all possible effects upon thermoelastic loss, particularly the D shape of the cantilevers and the effect of twisting along the cantilever length.

Depositing a coating upon a silicon disc results in stress and curvature of the disc. It was shown that the stress and curvature of the disc that is observed in the lab has no effect upon changing the disc's thermoelastic loss. This finding suggests that silicon discs could be an easier substrate to use in coating mechanical loss experiments as they are without the complications previously discussed with the cantilevers.

To conclude, it is clear care must be taken when using thin silicon cantilevers to measure coating loss, particularly in the temperature range of 150 K – 300 K, where the loss of these samples is dominated by thermoelastic loss which can be affected by curvature. Future work should investigate trying to explain the remaining differences between the FEA cantilever models and experimental results (perhaps caused by the D-shape of cantilevers or the effect of twisting along the cantilever's length).

5 Effect of stress and temperature on the optical properties of silicon nitride membranes at 1550 nm

5.1 Introduction

The Einstein Telescope (ET) is designed to have a factor of 10 sensitivity better than Advanced LIGO. To achieve this aim, new mirror coatings with low mechanical loss at low temperatures will have to be found. Using the current gravitational wave detector coating (made of a silica and tantala) [160] is not an option as it becomes a significant limiting factor in performance at cryogenic temperatures due to loss peaks around 20 – 30 K [161, 121]. One interesting material that could be potentially used in a cryogenic gravitational wave detector coating is amorphous silicon.

Amorphous silicon (aSi) is being considered as a possible material due to its very low mechanical loss ($< 2 \times 10^{-5}$) at temperatures below 30 K [113] . However, a gravitational wave detector coating made of aSi/silica would be dominated by the mechanical loss of the silica layers (the currently used low index material in a gravitational wave detector coating) [54]. Finding an alternative low-index material to silica – such as silicon nitride – is therefore a high priority.

Silicon nitride has been measured to have low mechanical loss at both room and cryogenic temperatures. The mechanical loss of silicon nitride has been studied by Chao et al. [162] and Kuo et al. [163], who investigated different compositions of silicon nitride deposited using plasma enhanced chemical vapor deposition (PECVD). Their results showed that a silicon nitride coating has lower mechanical loss than current gravitational wave detector coating materials at room temperature and at temperatures down to 10 K. Liu et al. [164] and Southworth et al. [165] found similarly low mechanical loss at cryogenic temperatures for a silicon nitride coating deposited via low-pressure chemical vapor deposition (LPCVD) upon a substrate (10^{-5} loss) and for a highly stressed substrate-free thin film (10^{-6}).

Along with low mechanical loss, low optical absorption (α) is also a requirement of gravitational wave detector coatings (in order to prevent thermal deformation of the mirrors and maintain a low mirror temperature). Whilst the absorption of silicon nitride has been measured in the past, this chapter will investigate the effect of stress (S) and temperature (T) upon the product of optical absorption and $\frac{\mathrm{dn}}{\mathrm{dT}}$ (where n is refractive index) for low stress ($<$ 250 MPa) amorphous silicon nitride membranes. The topic of this chapter was motivated by the observations made by Liu et al. [164] and Southworth et al. [165] who jointly saw a decrease in the mechanical loss of silicon nitride as its stress was increased. Since incorporating more highly stressed silicon nitride in a gravitational wave detector coating will result in lower mechanical loss, it is therefore essential to investigate the effects of stress upon the membrane's optical absorption.

This work presented in this chapter was published in Frontiers in Materials. The paper was called "Effect of stress and temperature on the optical properties of silicon nitride at 1550 nm" [166] .

5.2 Brief outline of study

Photothermal common path interferometry (PCI) was used to measure the $\alpha \times \frac{dn}{dT}$ of a silicon nitride membrane (shown in Figure 5-1) whilst its stress and temperature was varied using laser heating. (The membrane's $\alpha \times \frac{dn}{dT}$ was measured and not just it's absorption, as the PCI technique measures a signal proportional to $\alpha \times \frac{dn}{dT}$ and since the dependence of stress and temperature upon $\frac{dn}{dT}$ is unknown, the absorption of the membrane can't be extracted). The absorption signal was measured under varying compressive and tensile stresses. In the first measurement, the membrane's $\alpha \times \frac{dn}{dT}$ was measured when the stress and temperature of the membrane was increased by heating the frame of the membrane with a 532 nm laser. In the second set of measurements, the membrane's $\alpha \times \frac{dn}{dT}$ was measured when its temperature increased and its stress decreased by directly heating the membrane with a 1550 nm laser. To calculate the magnitude of the stress and temperature changes caused by the laser heating, experimental measurements and FEA models were used. Determining the change in the temperature and stress caused by laser heating enabled estimates of the membrane's thermal conductivity and thermal expansion to be made. Once the temperature and stress changes were known, ordinary least squares was used to identify the relationships between the membrane's stress and $\alpha \times \frac{dn}{dT}$, and temperature and $\alpha \times \frac{dn}{dT}$.

5.3 Description of the silicon nitride membrane samples

The membranes used in this study were fabricated by Norcada [167]. A 2×10^{-6} m thick silicon nitride coating was deposited (via LPCVD) upon a crystalline silicon (cSi) substrate with dimensions 10×10^{-3} m $\times$ 10×10^{-3} m $\times$ 500×10^{-6} m (length $\times$ width $\times$ thickness). An area of 5×10^{-3} m $\times$ 5×10^{-3} m was then etched off from the cSi substrate leaving a silicon nitride window in the centre of the cSi frame as shown in Figure 5-1.

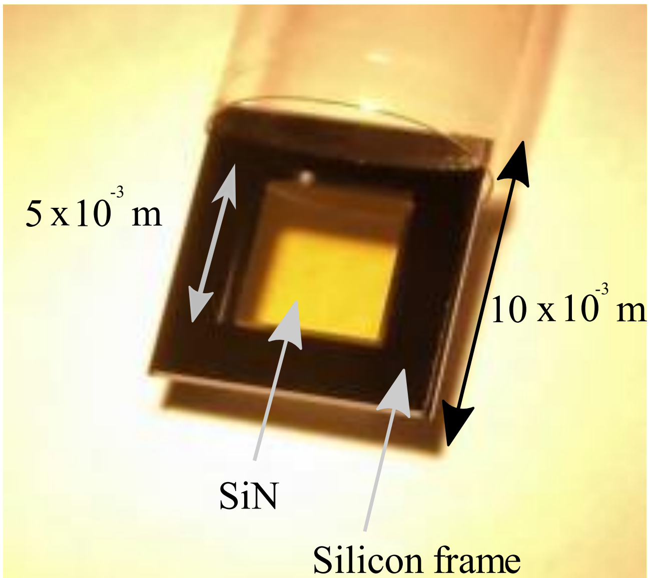

Figure 5-1: Photograph of a silicon nitride membrane. The 5×10^{-3} *m* $\times$ 5×10^{-3} *m* $\times$ 2×10^{-6} *m silicon nitride window is enclosed by a* 10×10^{-3} *m* $\times$ 10×10^{-3} *m* $\times$ 500×10^{-6} *m silicon frame.*

5.4 Measurement of $\alpha \times \frac{dn}{dT}$ for different temperatures and stresses

One membrane was attached to a steel screw using epoxy. This screw was then screwed into a cylindrical steel post which acted as a mount. PCI was used to measure the $\alpha \times \frac{dn}{dT}$ of the membrane [168]. A basic schematic of the setup used is shown in Figure 5-2.

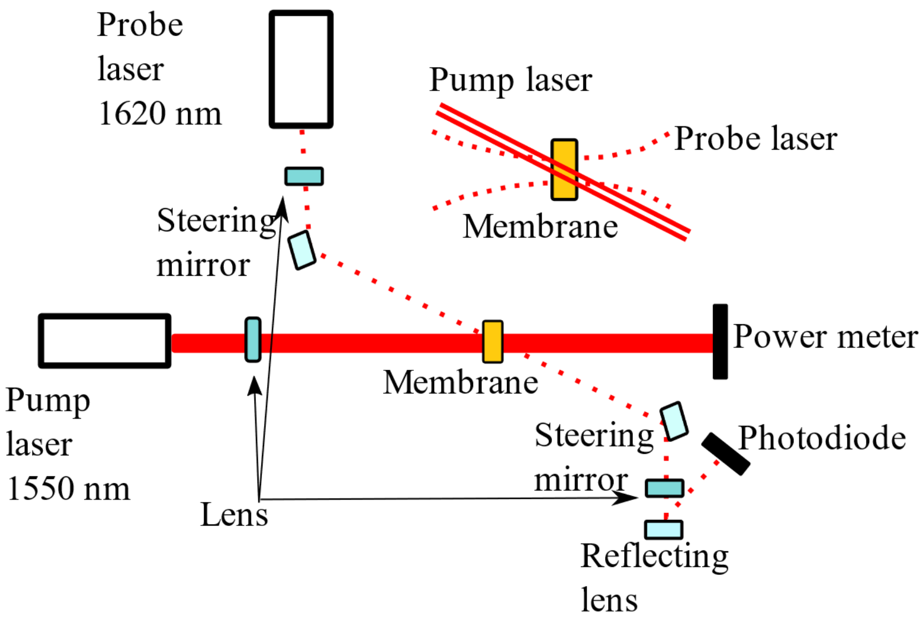

Figure 5-2: Schematic of the PCI setup. A high-power pump beam passes through the membrane inducing a phase shift in the probe beam. Since the probe beam is larger in radius, only a part of the probe beam experiences a phase shift. The phase shifted and non-phase shifted parts of the beam interfere with one another at a maximum value of one Rayleigh length from the beams crossing point and is measured by the photodiode. This value is proportional to the absorption of the membrane.

Two laser beams were passed through the membrane at the same time. A high powered, 1550 nm, pump laser beam was absorbed and created a thermally-induced optical length change in the membrane which is proportional to $\frac{dn}{dT}$. The second laser, a low power, 1620 nm, probe beam crossed the pump beam within the sample and was affected by the optical length change. The probe beam has a larger radius than the pump beam (as illustrated in Figure 5-2), and the part of the probe beam which overlapped with the pump beam experienced a phase shift. The phase-shifted part of the probe interfered with the unaffected part of the beam. At one Rayleigh length from the beams crossing point, the interference is at its maximum. The interference pattern at this point was imaged onto a photodiode sensor, generating a signal proportional to the membrane's $\alpha \times \frac{dn}{dT}$. A fused silica substrate of known absorption was then used as a calibration, allowing the absorption of the membrane to be determined.

Similarly to Steinlechner [112] , etalon effects were observed when measuring the silicon nitride membranes. To ensure the full transmission of both laser beams through the membrane (and the removal of the etalon effects) a crossing angle of 14 degrees between the pump and probe beams was used [112].

5.4.1 Changing the stress by heating the membrane

In the first measurement series, the change in the membrane's $\alpha \times \frac{dn}{dT}$ was monitored as the power of the 1550 nm pump laser beam was varied. This beam was positioned on the centre of the membrane (position 1), as shown in Figure 5-3.

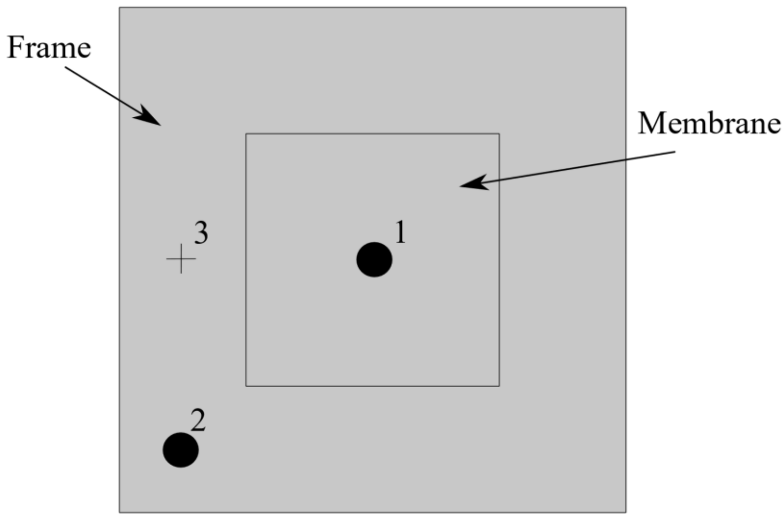

Figure 5-3: Geometry of the membrane and frame. The 1550 nm laser was focused upon the membrane at position 1. The 532 nm laser was focused upon the frame at position 2. The temperatures listed in Table 5-1 were measured at positions 1 and 3.

The initial power of the pump laser beam was 200 mW, as this was the minimum power required to create a measurable thermal effect within the membrane. The power of the pump beam was incrementally increased and $\alpha \times \frac{dn}{dT}$ was measured once the absorption signal had been shown to settle (indicating that the system had reached a steady state). $\alpha \times \frac{dn}{dT}$ was found to increase approximately linearly with the pump power, as shown in Figure 5-4.

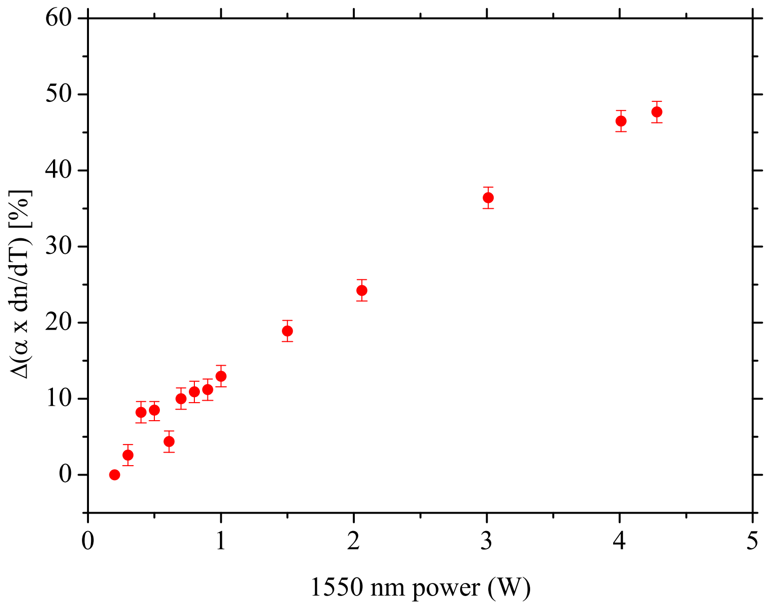

Figure 5-4: Change in the membrane's $\alpha \times \frac{dn}{dT}$ when the 1550 nm pump laser beam power was varied. The change in the membrane's $\alpha \times \frac{dn}{dT}$ was measured at position 1 in Figure 5-3.

Each change in $\alpha \times \frac{dn}{dT}$ was assumed to have an error of ± 1.4 % caused by misalignment during the measurement series. The size of this error was estimated by measuring the change in the absorption of the calibration sample before (21.4 %) and after the membrane measurements (20.8 %). The relative error of the power measurements was assumed to be negligible as the random variation in the laser power and power meter measurements were both very small (< 1 %). The measurements taken were insensitive to systematic errors in the laser power, power readings and material properties as relative changes in $\alpha \times \frac{dn}{dT}$ were investigated and not absolute values.

5.4.2 Changing the stress by heating the frame

In a second measurement series, the change in membrane's $\alpha \times \frac{dn}{dT}$ (at position 1 in Figure 5-3) was measured at constant 1550 nm power, while the cSi frame was heated with varying powers of a 532 nm laser (at position 2 in Figure 5-3). Figure 5-5 shows the change in value of $\alpha \times \frac{dn}{dT}$ when varying the 523 nm laser powers.

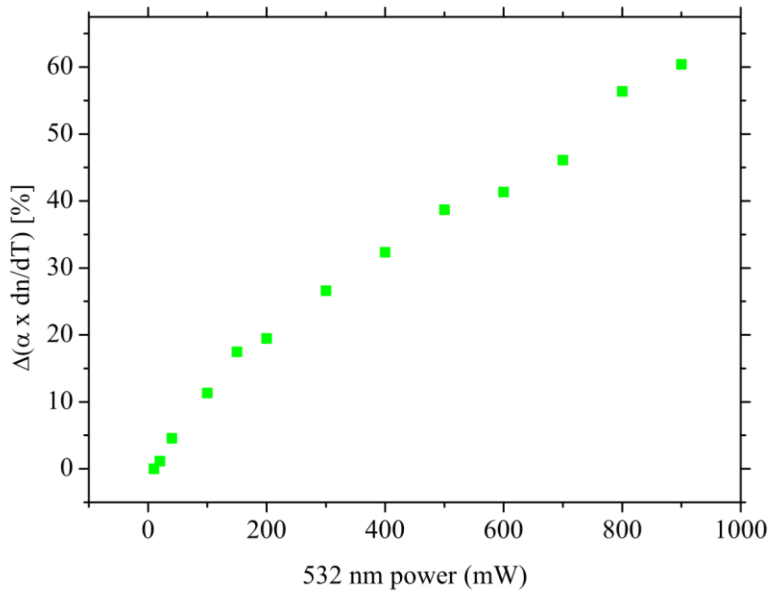

Figure 5-5: Change in the membrane's $\alpha \times \frac{dn}{dT}$ *when the 532 nm laser beam power is varied and the 1550 nm pump power is held constant. The measurements were taken at position 1 in Figure 5-3. The 532 nm laser beam was located at position 2 in Figure 5-3.*

The error associated with each change in $\alpha \times \frac{dn}{dT}$ was determined to be negligible as no change in the calibration signal was observed before and after the membrane measurements. For the same reasons as discussed earlier, the errors associated with the material properties, pump power and power meter did not couple into the change in $\alpha \times \frac{dn}{dT}$.

5.5 Modelling membrane temperature and stress due to laser heating

Using PCI, the change in a membrane's $\alpha \times \frac{dn}{dT}$ as a function of 1550 nm and 532 nm laser power is known. However, in order to determine how the membrane's $\alpha \times \frac{dn}{dT}$ changes with respect to the membrane's temperature and stress, FEA is required. FEA was used to determine the changes in the temperature and stress of the membrane in response to laser heating of both the frame and of the membrane itself. A number of FEA models were built, with experimental measurements carried out to provide inputs to the models and to verify the modelling results. A summary of all of the FEA models used in this chapter is presented in Figure 5-6.

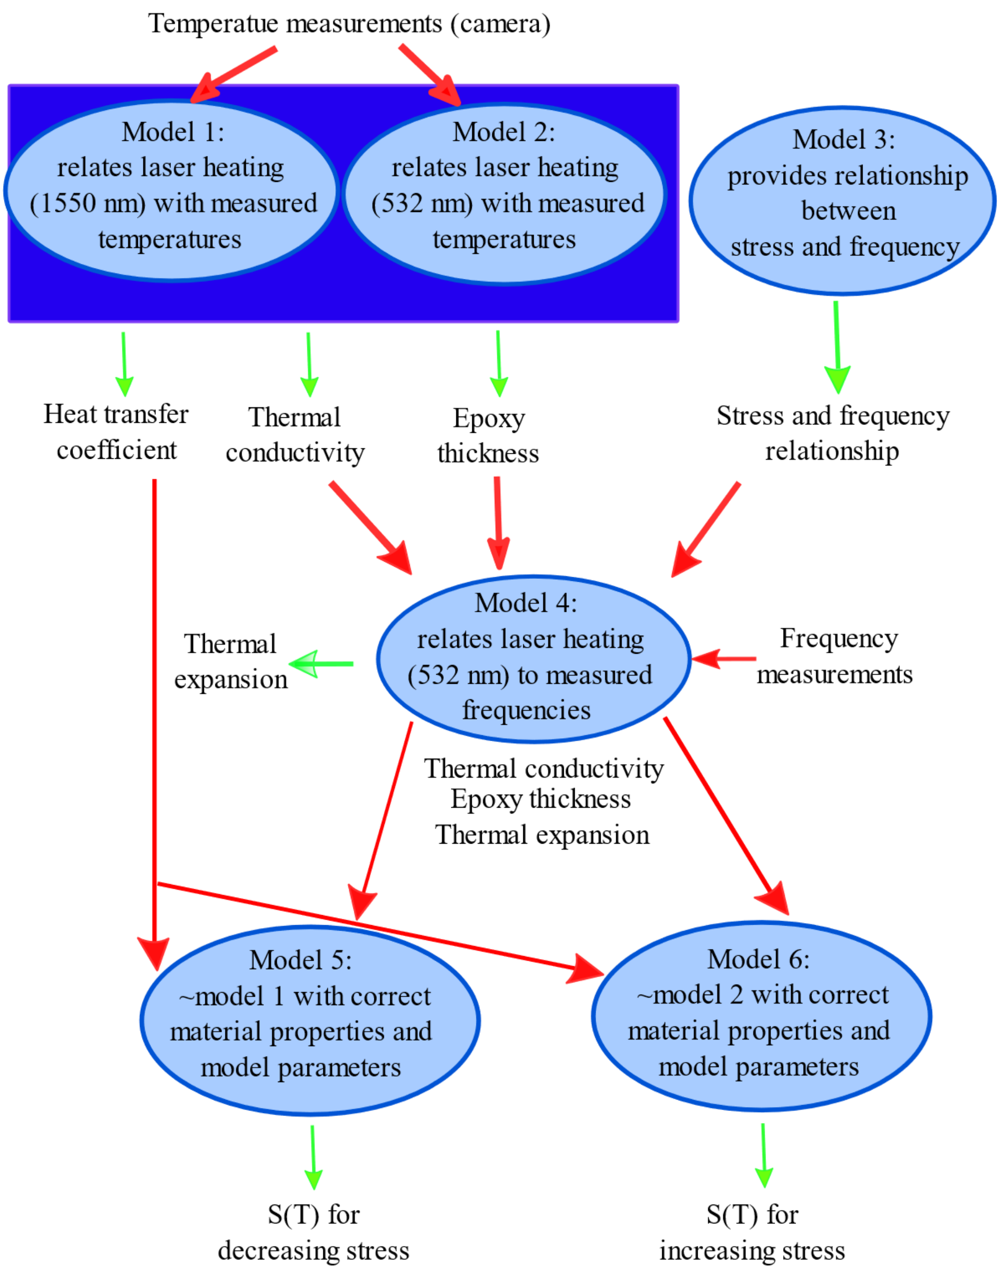

Figure 5-6: Schematic of the different FEA models used, their relationship to one another as well as their inputs (identified by red arrows) and outputs (identified by green arrows). Model 1 and model 2 are used in conjunction with one another to determine the heat transfer coefficient of air, the thermal conductivity of the membrane and the epoxy thickness used to attach the membrane to the screw. Model 3 determines the relationship between the membrane's resonance frequency and its stress. Model 4 determines the thermal expansion of the membrane by using frequency measurements as well as outputs from models 1,2 and 3. Models 5 and 6 use the newly determined values of thermal conductivity, thermal expansion, heat transfer coefficient and epoxy thickness to determine the stress and temperature changes when the membrane is directly heated with the 1550 nm laser (model 5) and when the frame is heated using a 532 nm laser (model 6).

Models 1 and 2 are used in conjunction with one another to replicate experimental temperatures of the membrane and frame when the membrane is directly heated with the 1550 nm laser (model 1) and when the frame is heated with the 532 nm laser (model 2). By replicating the experimental temperatures, these models determine values for the heat transfer coefficient of air, the thermal conductivity of the membrane and the epoxy thickness used to attach the membrane to the screw. Model 3 is a purely hypothetical model which is used to determine the relationship between the membrane's resonance frequencies and it's stress. This relationship is then used to convert experimentally observed changes in a membrane's resonant frequencies (as the membrane's frame is heated with varying powers of the 532 nm laser) to changes in stress. Model 4 replicates these changes in stress of the membrane due to laser heating and in doing so, determines a value for the membrane's thermal expansion coefficient. Models 5 and 6 use the newly determined values of thermal conductivity, thermal expansion, heat transfer coefficient and epoxy thickness to determine the stress and temperature changes when the membrane is directly heated with the 1550 nm laser (model 5) and when the frame is heated using a 532 nm laser at room pressure (model 6). Due to the mismatch in thermal expansion between the membrane and frame, heating the frame leads to an increase in membrane stress whilst heating the membrane leads to a decrease in stress.

5.5.1 Calibration of the FEA membrane temperature profile

Before the changes in membrane stress due to laser heating (described in Section 5.4) can be determined, it is essential to show that measured temperatures of the membrane and frame can be replicated computationally (when the membrane is directly heated with the 1550 nm laser and the frame is heated with the 532 nm laser, as described in Section 5.4). This is an important first step as it illustrates that the thermal characteristics of the mounted membrane can be correctly modelled. Building upon this model, the unknown stress changes due to laser heating can then be calculated.

Using the PCI setup described in Section 5.4, the temperature of the mounted membrane (at position 1 in Figure 5-3) was measured as the power of the 1550 nm laser (which directly heated the membrane at position 1) was varied (when the probe beam was turned off). The temperature of the membrane was measured using a Thermovision A40M [169] thermal camera. The temperature of the membrane as a function of 1550 nm laser power is shown in

Figure 5-7. The temperatures were taken when the membrane was observed to reach a steady state.

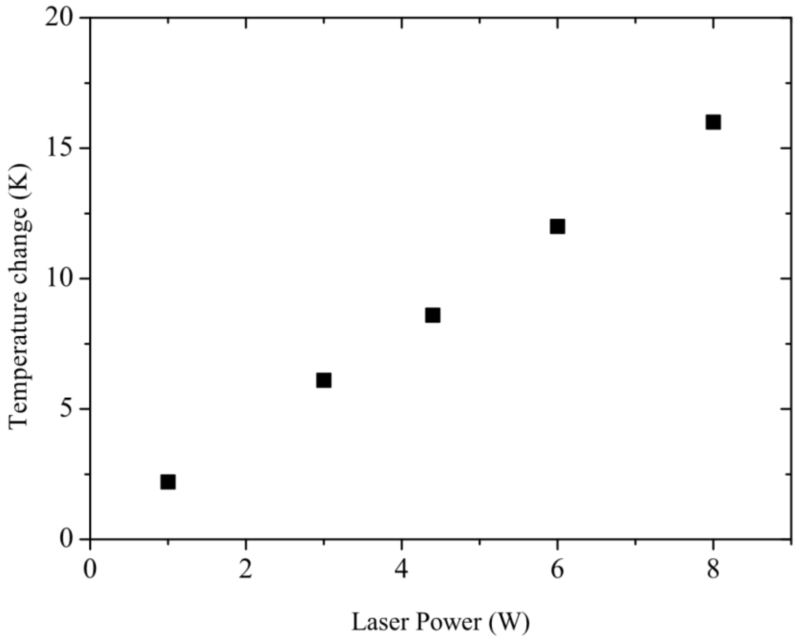

Figure 5-7: Measured temperature changes at position 1 (in Figure 5-3) when the membrane is directly heated with varying powers of a 1550 nm laser.

The temperatures of the membrane (at position 1 in Figure 5-3) and frame (at position 3, in Figure 5-3) were measured at different magnitudes of 532 nm laser power, when the laser heated the frame at position 2 (in Figure 5-3). These temperatures and laser powers are listed in Table 5-1. The temperatures were noted when the frame and membrane were observed to reach a steady state.

	Membrane temperature (K)		Frame temperature (K)	
Power (mW)	Measured	Model	Measured	Model
20	290	291	290	291
200	296	296	305	305
500	303	303	328	329

Table 5-1: Experimental and computational temperatures of the frame (position 3 in Figure 5-3) and membrane (position 1 in Figure 5-3) for different 532 nm laser powers. The computational temperatures were obtained for an FEA model which assumed a membrane thermal conductivity of 23 W/mK, a heat transfer coefficient of 12 W/m^{2}K and a 6×10^{-5} m thickness of epoxy layer.

To replicate these two sets of experimental temperatures using FEA, two *steady state* COMSOL models each within their own *heat transfer in solids* interface were built. The first model (model 1 in Figure 5-6) aimed to replicate the measured temperatures when the 1550 nm laser directly heated the membrane. The second model (model 2 in Figure 5-6) aimed to replicate the measured temperatures when the 532 nm laser was aligned upon the frame.

In model 1, the geometry of the membrane, screw, epoxy and steel post were specified, and the heat applied to the membrane from the laser (Q_{las}) was modelled as a surface heat flux. Approximating the heat source as a surface heat flux (and not volumetric) was a valid approximation as the thermally thin membrane temperature profile (has a minimal thermal gradient through it's thickness) which could be replicated using a surface heat flux [170]. The heat flux approximating the laser beam in the model was defined to have a Gaussian distribution such that it replicated the 70×10^{-6} m diameter (1/e^2) beam profile. In order to replicate the experimental heating power, a silicon nitride absorption value had to be assumed within the model. Since it can be observed in Figure 5-7 that there exists an approximately linear relationship between the 1550 nm laser power and the measured temperature changes of the membrane, a constant value was assumed for the membrane absorption. For modelling purposes, the absorption of the membrane was assumed to be 337 ppm as this was the calculated absorption of this membrane at its lowest PCI pump power (when a value of $\frac{dn}{dT} = 4\times 10^{-5}$ K^{-1} [112] was used). An image of the geometry built in COMSOL along with the heat flux applied to the membrane is shown in Figure 5-8.

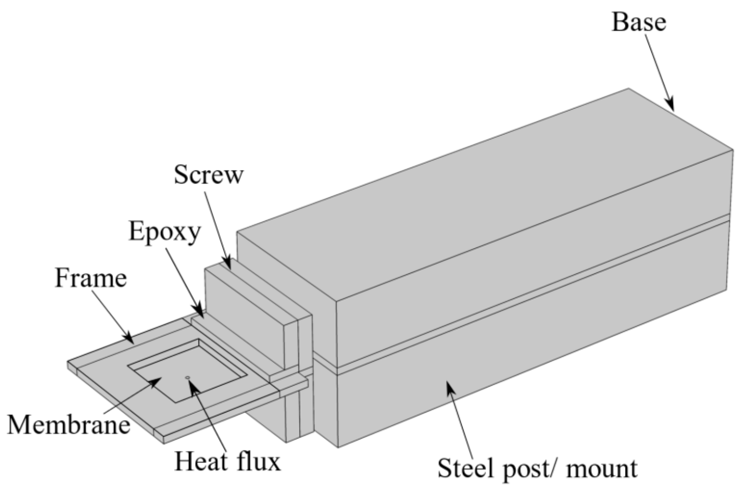

Figure 5-8: Geometry of the mounted membrane built in COMSOL.

It can be observed in Figure 5-8 that the cylindrical steel mount and screw in the experimental setup were approximated as being cuboids in the model. This was because of meshing problems associated with very small dimensions of different geometries coming into contact. This change in geometry was assumed to have a negligible effect on the results since the individual components were modelled to have an almost identical cross-sectional area and volume as that of the original bodies. The base of the mount was fixed at room temperature thus simulating an infinite conductive heat sink (which replicated the laboratory conditions where the mounted membrane was fixed to a metal support of much larger size).

Radiative cooling was incorporated within the model by specifying each material's emissivity values (ϵ) [171]. Radiative losses (Q_{rad}) were estimated using:

$$Q_{\text{rad}} = \epsilon(T^4 - T_{\text{amb}}^4) \,, \tag{5.1}$$

where T is the temperature of the material and T_{amb} is the ambient temperature. Heat losses due to convection (Q_{conv}) were also incorporated within the model using:

$$Q_{\text{conv}} = h(T - T_{\text{amb}}) \,, \tag{5.2}$$

where h is the heat transfer coefficient. In the model, the heat transfer coefficient was initially assumed to be 15.5 W/m^{2}K (as this is the middle of the range quoted in the literature for natural convection (6 – 25) W/m^{2}K [172]). Having defined the geometry, material properties, inward and outward heat fluxes, the FEA model could then be solved for temperatures using a version of the heat equation [173]:

$$\rho C_{\rho} \frac{\partial T}{\partial t} = Q + \nabla.(k\nabla \text{T}) \,, \tag{5.3}$$

where ρ is the density, C_{ρ} is the specific heat capacity at constant pressure, Q represents the heat sources and heat losses. t is the time and k is the thermal conductivity.

Model 2 was identical to model 1, with the exception that the surface heat flux was now applied to the frame (position 2 in Figure 5-3), to mimic the 532 nm heating power. The heat flux was defined to have a Gaussian spatial profile with the same 1/e^2 diameter as the laser beam (1.5×10^{-3} m). 96 % of the initial laser power incident upon the frame was assumed to be absorbed because a 2×10^{-6} m thick silicon nitride layer (with a refractive index of 2.05 at 532 nm [174]) on cSi (with a refractive of 4.15 at 532 nm [175]) makes a good anti-reflection coating with a reflectivity of 4 % at normal incidence. Since the cSi frame has a very short absorption depth of 1.3×10^{-6} m at a wavelength of 532 nm, capturing the transparency of the frame was not important and modelling the heat as being applied to the surface was appropriate.

Both of the FEA models were solved and their meshes were checked for convergence. To compare the FEA modelled temperatures to the experimental temperatures, the modelled temperatures had to be averaged over a circular area with diameter 900×10^{-6} m, as this was determined to be the pixel size of the Thermovision A40M thermal camera. Significant differences were observed between the FEA and experimental temperatures. However, in both of the FEA models a best estimate value was assumed for the conductivity of the membrane (16.5 W/mK), heat transfer coefficient (15.5 W/m^{2}K) and thickness of epoxy layer which separates the membrane and screw in the direction of the base (5.5×10^{-5} m) as no precise values were known. Unlike the other material properties and geometric dimensions used in the FEA models, these parameters had a very large uncertainty associated with them. The best estimates of these parameters were calculated by taking the median of their range of possible values (for the thermal conductivity this range was 4.9 - 30 W/mK [176] [177], heat transfer coefficient 6 - 25 W/m^{2}K [172] and epoxy thickness 1×10^{-5} - 1×10^{-4} m, where the epoxy thickness upper and lower bounds were estimated using a digital calliper). It was observed that by varying these three parameters simultaneously in both models, the three sets of measured temperatures could be matched almost exactly for only one set of parameters. The experimental temperatures were replicated best when: the thermal conductivity of the membrane was assumed to be (23 ± 3) W/mK, the heat transfer coefficient was assumed to be (12 ± 2) W/m^{2}K and the thickness of the epoxy layer was assumed to be $(6 \pm 0.75) \times 10^{-5}$ m (all within their ranges of possible values). The similarity of the FEA temperatures (which used these parameter values in their model) and the measured values can be observed in Figure 5-9 and Table 5-1. Figure 5-9 also shows the FEA temperatures calculated by model 1 when *other sets* of conductivities, heat transfer coefficients and epoxy thicknesses were used. These *other sets* of parameters were specifically chosen as when these sets of values were used in model 2, the calculated FEA temperatures replicated the measured values (as closely as the FEA temperatures listed in

Table 5-1). Figure 5-9 therefore shows that only one set of parameter values enables both FEA models (1 and 2) to replicate the measured temperatures. This set of parameter values is therefore the best estimate for these parameters, and was used for the rest of the modelling and analysis.

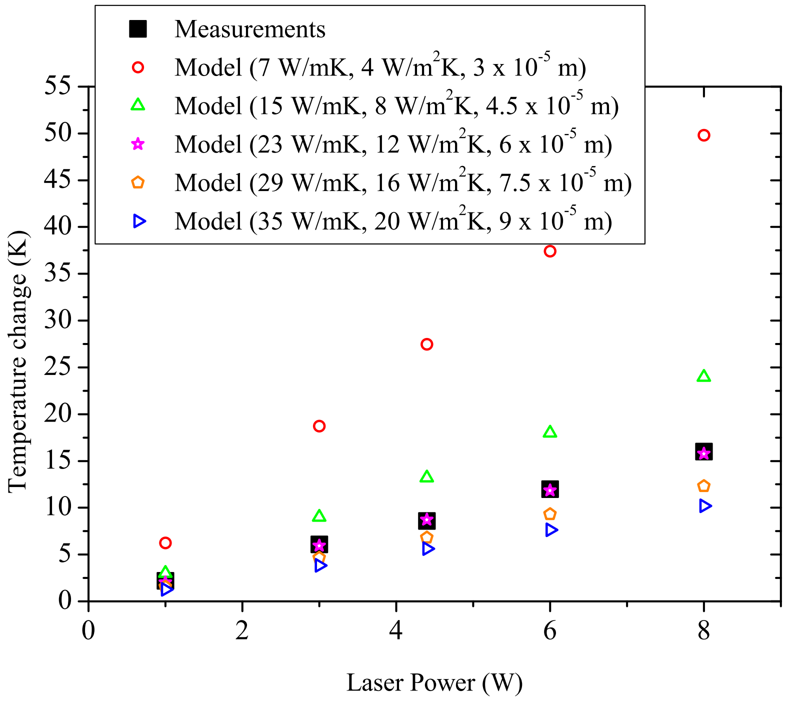

Figure 5-9: Measured and FEA temperature changes at position 1 (in Figure 5-3) when the membrane is directly heated with varying powers of a 1550 nm laser. The FEA temperatures were calculated when different values of membrane thermal conductivity, heat transfer coefficient and thickness of epoxy layer were assumed.

The best estimate for the thermal conductivity of the membrane depends on a number of parameters in the 1550 nm and 532 nm laser heating models (absorption of membrane, emissivity of frame and membrane etc). To find the error in the membrane's thermal conductivity, a philosophy of varying these parameters within their error bounds to maximise the membrane temperature was followed. In this case, the total heat loss of the system also must be at a maximum to reduce this membrane temperature to be consistent with the measured values, allowing errors on the parameters governing the heat loss to be estimated. By an iterative process of varying the parameters in both models and comparing to measured temperatures, the possible range of values of membrane conductivity, heat transfer coefficient and epoxy thickness were obtained and from this, an error on these parameters was estimated. To summarise, the errors in the thermal conductivity, heat transfer coefficient and epoxy thickness were determined from the uncertainties in the experimental measurements and in the inputs to the FEA models.

In determining these errors, it was assumed that the error associated with: the silicon frame material properties, the relative error in the camera ($\pm$ 0.08 °C *[169]*) as well as the size of the error in the power readings was negligible. Finally, it is important to note that the errors determined were insensitive to the uncertainty in the membrane's specific heat capacity as this uncertainty affected the time at which the system reached equilibrium but not its final temperature.

5.5.2 Calibration of the FEA thermal stress effect via measurement of resonance frequencies

Having demonstrated in Section 5.5.1, that FEA is able to replicate experimental temperatures it is now important to show that FEA modelling can also replicate experimental changes in stress caused by laser heating. Once the modelling has been shown to be able to replicate temperature and stress changes due to laser heating, the temperature and stress changes due to the laser heating described in Section 5.4 will then be able to be calculated with reasonable accuracy.

When the temperature of a membrane and its frame is changed, the dimensions of both parts change (due to their thermal expansion coefficients), resulting in a change in the stress of the membrane. This change in stress can be determined by monitoring the change in resonant frequency of the membrane, which is related to its stress (S) via [178]:

$$f_{n_x,n_y} = \frac{1}{2L}\sqrt{\frac{S}{\rho} \times \left(n_x^2 + n_y^2\right)}\,, \tag{5.4}$$

where L is the length of the membrane's sides and n_x and n_y are integers used to define the membrane's modes which have $(n_x - 1)$ and $(n_y - 1)$ numbers of nodes in the x and y directions. The first mode of a square membrane is shown in Figure 5-10.

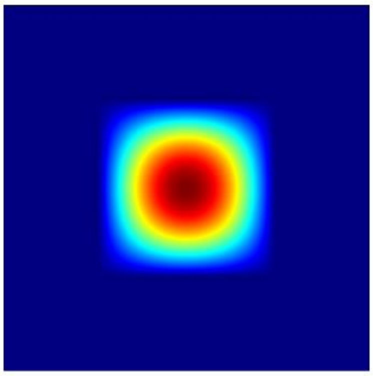

Figure 5-10: Membrane's first resonant mode. Pictured is the membrane and its frame. Red represents maximum displacement whilst blue illustrates minimum.

To calibrate the modelled thermal-stress effect, the change in a membrane's first mode frequency was measured as the frame was heated with a 532 nm laser at varying powers. The mount of the steel post used to support the membrane (described in Section 5.4) was removed and the screw was clamped between two steel blocks that were located within a vacuum tank as shown in Figure 5-11.

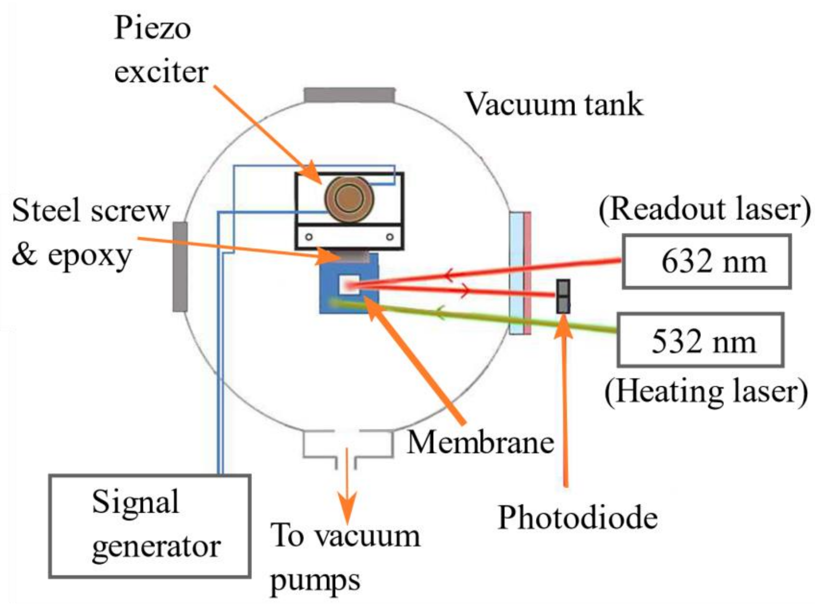

Figure 5-11: Experimental setup used to measure the change in frequency of a membrane's resonant modes when its frame was heated with a 532 nm laser.

The first resonance of the membrane was excited by a piezo-electric transducer that was positioned less than 2×10^{-3} m below the bottom of the membrane. A 5 mW, 633 nm laser beam was reflected by the membrane onto a split photodiode which generated a signal proportional to the amplitude of the membrane vibration. A signal analyser was used to measure the mode frequency. The 532 nm green laser was used to heat one corner of the membrane's frame (position 2 in Figure 5-3). The stress of the membrane was changed by varying the power of the green laser. The change in mode frequency of the heated membrane was noted when it had reached a steady state. Whilst it would have been optimum to calibrate the thermal stress model at room pressure values (as the laser heating discussed in Section 5.4 was at room pressure), it was not possible to excite the membrane's resonant mode frequencies at room pressure due to air damping. At a pressure of 10^{-4} mbar and without any laser heating, the frequency of the membrane's first mode was measured to be 31820 Hz. As the laser power was increased, the mode frequency was observed to increase, as shown in Table 5-2.

Laser power (mW)	1st mode frequency after laser heating (Hz)	Change in stress of membrane (MPa)
4	31839	0
40	31989	1.5
84	32177	3
218	32779	8
440	33827	17

Table 5-2: Membrane's first mode frequency when the frame was heated with varying powers of a 532 nm laser.The first mode frequency of the membrane without laser heating was measured at 31820 Hz. The FEA change in stress responsible for the measured frequency change is also listed.

COMSOL was used to determine the stress changes associated with the changes in the membrane's mode frequency (model 3 in Figure 5-6). The geometry of a membrane and its frame was built in an *eigenfrequency analysis* within a *solid mechanics* interface. To simulate the clamp, one side of the membrane's frame was fixed in space using a fixed constraint. The initial stress of the membrane was then specified, the model solved for its first mode frequency and the mesh was checked for convergence. The stress of the membrane was then incrementally increased and the first mode's frequency found. The relationship between the membrane's stress and mode frequency can be observed in Figure 5-12.

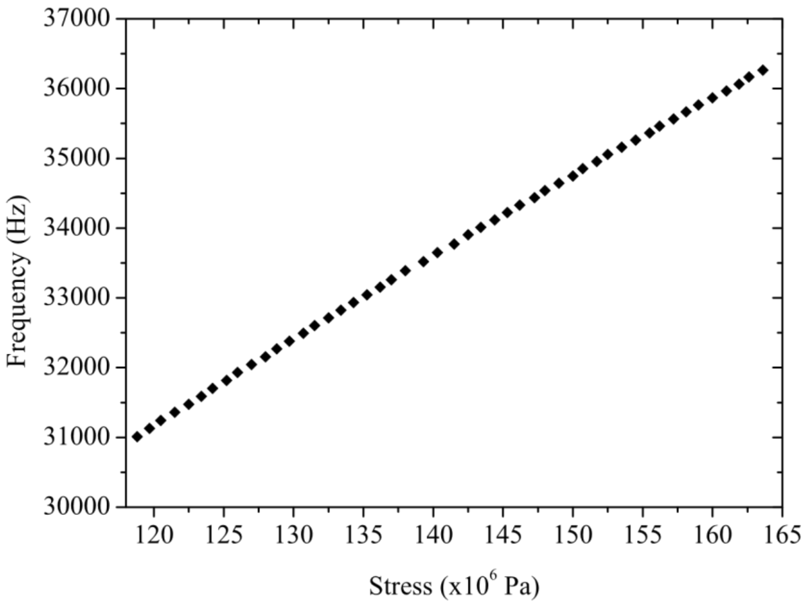

Figure 5-12: COMSOL determined relationship between the membrane's stress and the frequency of its first resonant mode.

The relationship presented in Figure 5-12 enabled the change in stress responsible for the measured frequency changes to be calculated. These changes in stress are listed in Table 5-2.

An FEA model (model 4 in Figure 5-6) was built to simulate the changes in stress caused by 532 nm laser heating. Whilst it would have been ideal to have measured the temperature of the frame and membrane at each resonant frequency, this was not possible as the thermal camera was not vacuum compatible. When the camera was positioned outside the vacuum chamber, the glass of the viewport interfered with any temperature measurements made. However, as shown in Section 5.5.1, model 2 can replicate the temperature profile of the membrane when its frame is heated with a 532 nm laser at room pressure. Model 2 (with the new best estimates of epoxy thickness, membrane conductivity and heat transfer coefficient) was therefore modified to build model 4.

The geometry and material properties of the membrane, frame, epoxy and screw remained the same in model 4, as in model 2. However, the geometry of the mount was modified to reflect the clamp used in this experiment. The effect of the clamp was the same as before, acting as a heat sink to the rest of the structure. To simulate vacuum conditions, the heat transfer coefficients for the structure were disabled, effectively turning off convective cooling. A multi-physics analysis was added to the model, allowing coupling between the thermal and mechanical properties. This enabled the strain resulting from thermal expansion to be calculated using:

$$\epsilon = \beta(T - T_{\mathrm{amb}}) \ , \tag{5.5}$$

which in turn allowed the heating-induced stress changes in the membrane to be determined. ϵ and β represent the material's strain and thermal expansion coefficients respectively.

The modelled stress change resulting from laser heating could now be compared to the experimental stress change calculated from the frequency shifts of the heated membrane. A discrepancy between the model and experimental values was observed. Noting that model 2 had already been able to replicate the temperature profile of the 532 nm laser heated membrane at room pressure and the simplicity of the stress/frequency relationship, other variables within the model were investigated to see whether they could explain this discrepancy. Upon varying the thermal expansion coefficient of the membrane, it was observed that for a value of $(1.4 \pm 0.2) \times 10^{-6}$ $1/\mathrm{K}$ the experimental and computational changes in stress matched one another very well. This new estimate of thermal expansion fell within the range of values quoted in the literature (between 1.4×10^{-6} and 3.7×10^{-6} $1/\mathrm{K}$ [179]). The similarity between the experimental and computational changes in stress (when using this thermal expansion coefficient) can be observed in Figure 5-13. Figure 5-13 shows the experimental and computational stress changes of the membrane at different laser powers. The computational stress changes were calculated using an array of different thermal expansion coefficients. It can be observed that the experimental stress changes can be best replicated assuming a value of 1.4×10^{-6} $1/\mathrm{K}$.

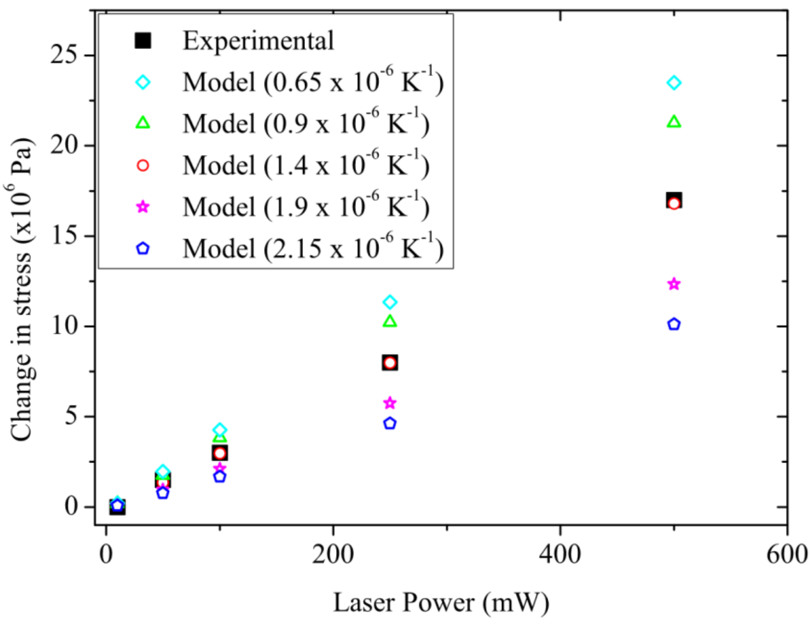

Figure 5-13: Experimental and computationally determined changes in stress of the membrane for different laser powers. The computational stress changes were determined using different thermal expansion coefficients in the FEA model.

The result for the thermal expansion of the membrane depends on a number of parameters in the COMSOL models labelled 3 and 4 (membrane's Youngs modulus, membrane's Poisson ratio etc). To find the error in the membrane's thermal expansion, model 3's parameters were varied within their error bounds to maximise the membrane's change in stress (which corresponded to the measured change in frequency). To replicate these maximum changes of stress in model 4, required a minimum value of thermal expansion. This value represented the lower bound value of thermal expansion. Repeating the process, the upper bound value of thermal expansion was determined. The difference between the highest (or lowest) value of thermal expansion and its best estimate value (which ever resulted in the larger value) is quoted here as the error. In this calculation, it was assumed that the error in the membrane's resonant mode frequencies was negligible (as it was < 0.1 %) and the error in the membrane's Youngs modulus and Poisson ratio was assumed to be 10 % (as no error was provided by the manufacturer).

5.6 Analysis: effect of stress and temperature on $\boldsymbol{\alpha \times \frac{dn}{dT}}$

Section 5.4 presented the measured change $\alpha \times \frac{dn}{dT}$ of a membrane when its frame and membrane were separately heated with varying powers of a 532 nm and 1550 nm laser. Section 5.5 showed that FEA can replicate thermal and mechanical measurements of a membrane and frame when they are heated with varying powers of laser light. Now variations of the FEA models (described in Section 5.5) will be used to calculate the stress and temperature changes of the membrane described in Section 5.4. Once the temperature and stress changes are known, ordinary least squares will be used to identify the relationships between the membrane's stress and $\alpha \times \frac{dn}{dT}$, and temperature and $\alpha \times \frac{dn}{dT}$.

5.6.1 Determining the stress and temperature changes due to laser heating

FEA models (models 5 and 6 in Figure 5-6) were built to determine the stress and temperature changes associated with the measured changes in $\alpha \times \frac{dn}{dT}$ when the membrane was directly heated with the 1550 nm pump beam and the frame was separately heated with the 532 nm laser beam (as described in Section 5.4).

Model 5 was used to determine the stress and temperature changes associated with the measured changes in $\alpha \times \frac{dn}{dT}$ when the membrane was heated with the 1550 nm pump beam. Model 5 was identical to model 1 with the exception that it used the newly determined best estimates of epoxy thickness, heat transfer coefficient, membrane thermal conductivity and membrane thermal expansion coefficient. Furthermore, a *thermal stress interface* which coupled the *heat transfer in solids* to a *solid mechanics* interface (described in Section 5.5.2) was added to this model to enable the temperature dependent stress changes to be calculated. Using the *steady state analysis*, the stress and temperature changes of the membrane and frame were calculated for the different laser powers. To determine the temperatures and stresses of the membrane which relate to the measured $\alpha \times \frac{dn}{dT}$ values, the FEA temperatures and stresses of the membrane were averaged over the Gaussian distributed, surface heat flux's $1/e^2$ diameter (which was used in the model to replicate the 1550 nm pump beam).

The temperatures and stresses were averaged over the surface heat flux's $1/e^2$ diameter as this represents the area over which the $\alpha \times \frac{dn}{dT}$ values are sampled in the PCI method [180]).

Model 6 was used to determine the stress and temperature changes associated with the measured changes in $\alpha \times \frac{dn}{dT}$ when the frame was heated with the 532 nm laser. Model 6 was identical to model 2 with the exception that it used the newly determined best estimates of epoxy thickness, heat transfer coefficient, membrane thermal conductivity and membrane thermal expansion coefficient. Furthermore, a *thermal stress interface* which coupled the *heat transfer in solids* to a *solid mechanics* interface (described in Section 5.5.2) was added to this model to enable the temperature dependent stress changes to be calculated. Using the *steady state analysis,* the stress and temperature changes of the membrane and frame were calculated for the different laser powers. The temperatures and stresses which relate to the measured $\alpha \times \frac{dn}{dT}$ values, were then estimated by averaging the FEA temperatures and stresses at the centre of the membrane (position 1 in Figure 5-3) over a circle with diameter 70×10^{-6} m (as this represents the area over which the $\alpha \times \frac{dn}{dT}$ values are sampled in the PCI method).

Figure 5-14 shows the changes in temperature and stress of the membrane at position 1 (in Figure 5-3) when the different 532 nm and 1550 nm laser powers were used to heat the membrane as the change in the membrane's $\alpha \times \frac{dn}{dT}$ was measured (see Section 5.4).

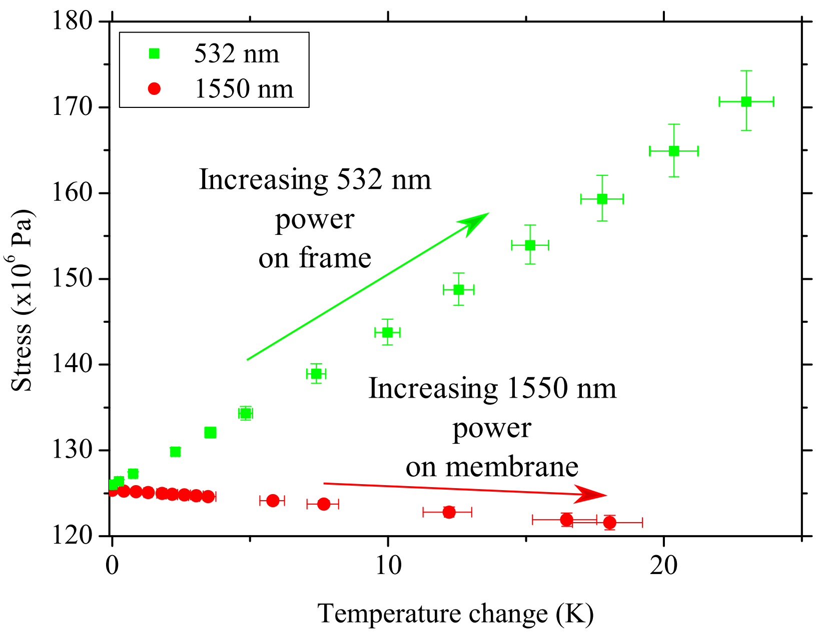

Figure 5-14: Change in the temperature and stress of the membrane at position 1 (in Figure 5-3) as the frame is heated with a 532 nm laser (green squares) and the membrane is heated with a 1550 nm pump laser (red circles).

It can be observed in Figure 5-14, that as the frame (position 2 in Figure 5-3) is heated by progressively larger powers of the 532 nm laser, the temperature of the membrane increases at position 1 (in Figure 5-3) and the membrane's stress at position 1 also increases. The stress of the membrane increases because as the frame is heated with the 532 nm laser, the frame increases in temperature more quickly than the membrane (due to the membrane's low thermal conductivity). As a result of the frame's greater temperature and thermal expansion coefficient (compared to the membrane), the frame expands more than the membrane, pulling on the membrane and thus increases its stress. This relationship between stress and temperature is noticeably different from directly heating the membrane with the 1550 nm laser. It can be observed in this case that as the laser power increases, the membrane's temperature increases and the stress decreases. In this example, the membrane's stress decreases as a result of the membrane rising in temperature more quickly than the frame (due to the membrane's low thermal conductivity) and therefore the membrane expands more quickly than the frame, thus the stress at position 1 decreases.

The error bars shown in Figure 5-14 were determined by re-running the COMSOL models (5 and 6) multiple times with model inputs which varied within their error bounds. The highest and lowest temperature and stress values from the COMSOL models were deemed to be the upper and lower error bounds.

5.6.2 Determining the relationship between stress and temperature upon $\boldsymbol{\alpha \times \frac{dn}{dT}}$

Deconstructing Figures 5-4, 5-5 and 5-14 enables the experimentally observed change in $\alpha \times \frac{dn}{dT}$ for both the 1550 nm and 532 nm lasers to be shown against their computationally determined stress and temperature values. Graphs of this form are useful as they help to identify the relationship between $\alpha \times \frac{dn}{dT}$ and temperature, and $\alpha \times \frac{dn}{dT}$ and stress. Earlier it was assumed for modelling purposes that the absorption of the membrane was constant during the 1550 nm laser heating. However, this was only an approximation as the R^2 value of the linear regression, laser power v temperature change (in Figure 5-7), is 0.99. Since the R^2 value is not equal to 1, it cannot be definitively concluded that there is not a small variation of absorption with temperature. To err on the side of caution, the 1550 nm temperatures and stresses are shown against the measured change in $\alpha \times \frac{dn}{dT}$ (and not just the change in $\frac{dn}{dT}$). Figure 5-15 shows the measured change in $\alpha \times \frac{dn}{dT}$ against the change in temperature for the two heating cases.

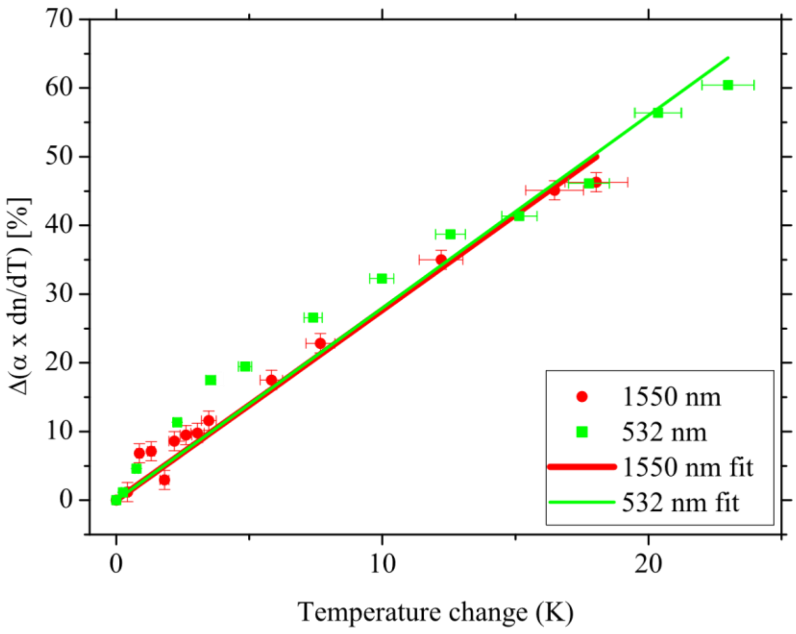

Figure 5-15: Measured change in the membrane's $\alpha \times \frac{dn}{dT}$ against it's change in temperature when the frame was heated with the 532 nm laser (green squares) and when the membrane was heated with the 1550 nm laser (red circles). All measurements and computational values were taken or calculated at position 1 in Figure 5-3. The green and red lines show the fitted changes in $\alpha \times \frac{dn}{dT}$ (as described in the text) for the 532 nm laser and 1550 nm laser respectively.

An increase in temperature results in an increase in $\alpha \times \frac{dn}{dT}$. For the same increase in membrane temperature, the membrane experiences approximately the same increase in $\alpha \times \frac{dn}{dT}$ when the frame is heated with the 532 nm laser and when the membrane is heated with the 1550 nm laser.

Figure 5-16 shows the measured change in the membrane's $\alpha \times \frac{dn}{dT}$ as a function of stress for when the frame is heated with the 532 nm laser and when the membrane is directly heated with the 1550 nm laser. It can be observed that $\alpha \times \frac{dn}{dT}$ increases when the stress is reduced and when the stress is increased.

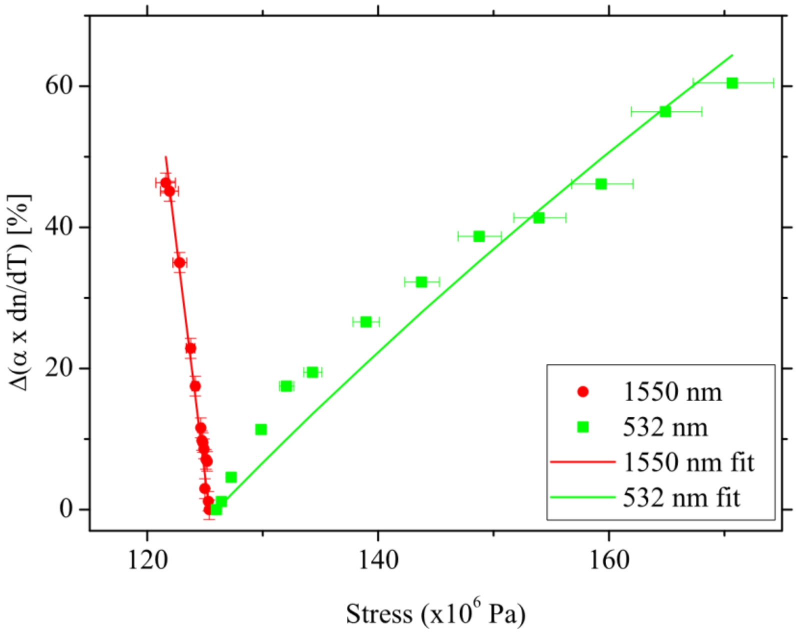

Figure 5-16: Measured change in the membrane's $\alpha \times \frac{dn}{dT}$ *as the stress of the membrane is changed using the 532 nm laser to heat the frame (green squares) and the 1550 nm laser to heat the membrane directly (red circles). All measurements and computational values were taken or calculated at position 1 in Figure 5-3. The green and red lines show the fitted changes in* $\alpha \times \frac{dn}{dT}$ *(as described in the text) for the 532 nm laser and 1550 nm laser respectively.*

From the results shown in Figures 5-15 and 5-16, it can be hypothesised that changing the membrane's stress had little to no effect upon the change in the membrane's $\alpha \times \frac{dn}{dT}$ as the membrane $\alpha \times \frac{dn}{dT}$ can be seen to have the same temperature dependence both when the stress of the membrane is increased and when it is decreased (by different amounts).

Having qualitatively established the relationship between $\alpha \times \frac{dn}{dT}$ and stress, and $\alpha \times \frac{dn}{dT}$ and temperature, these relationships were now quantitatively determined using the regression:

$$\Delta\left(\alpha \times \frac{dn}{dT}\right)_{T,S} = a \times \Delta T + \ b \times \Delta S\ , \tag{5.6}$$

where a and b are unknown parameters. Performing the regression, the values $a =$ (2.774 $\pm$ 0.214) % / K and b = (0.014 $\pm$ 0.133) % / MPa were obtained, where the errors represent 95% confidence intervals. The fitted $\alpha \times \frac{dn}{dT}$ values using the parameters a and b are shown as red (1550 nm) and green lines (532 nm) in Figures 5-15 and 5-16. The fit reproduces the measured points well. This is quantified by the regressions R^2 being equal to

0.98 which means that 98% of the variation in the measured values can be explained by the regression. The values of a and b were also tested for their significance. For a null hypothesis of each parameter being equal to 0, a had a p-value of 7×10^{-20} and b had a p-value of 0.835. This means that the null hypothesis can be rejected at a 1 % significance level for a (said another way this means with a 99 % confidence, a is not equal to 0). However, the null hypothesis could not be rejected at a 1, 5 or 10 % significance for b, implying that b is not statistically different from 0 (at a 1, 5 or 10 % level). Figures 5-17 and 5-18 show the separated temperature and stress effects upon $\alpha \times \frac{dn}{dT}$. These effects are described by:

$$\Delta\left(\alpha \times \frac{dn}{dT}\right)_T = (2.774 \pm 0.214)\ \%/\mathrm{K} \times \Delta T \quad \text{for} \quad 290\ \mathrm{K} \leq \mathrm{T} \leq 312\ \mathrm{K}\,, \tag{5.7}$$

and:

$$\Delta\left(\alpha \times \frac{dn}{dT}\right)_S = (0.014 \pm 0.133)\ \%/\mathrm{MPa} \times \Delta S \quad \text{for} \quad 121.5\ \mathrm{MPa} \leq \mathrm{S} \leq 171\ \mathrm{MPa}. \tag{5.8}$$

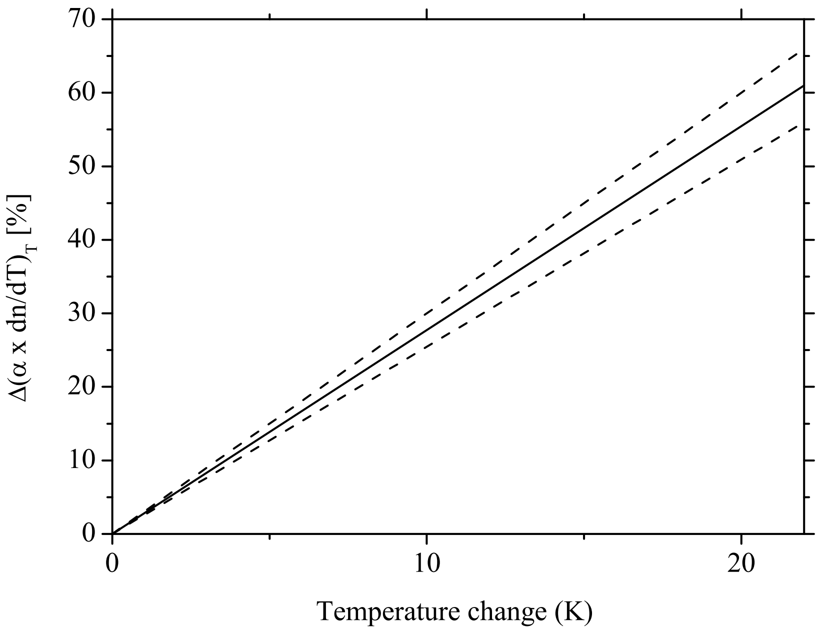

Figure 5-17: Relationship between the change in $\alpha \times \frac{dn}{dT}$ *and the change in temperature as numerically shown in Equation (5.7). The dashed lines represent the 95 % error confidence intervals.*

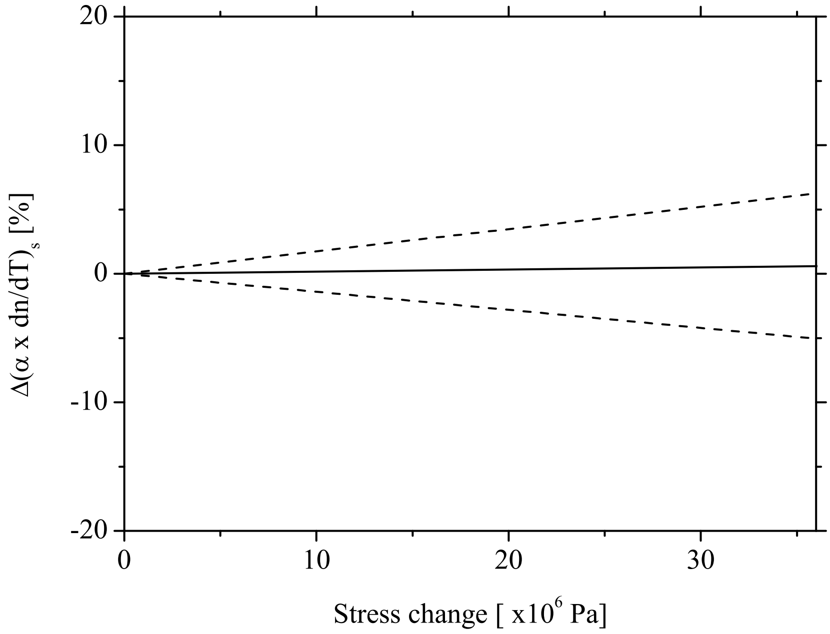

Figure 5-18: Relationship between the change in $\alpha \times \frac{dn}{dT}$ and the change in stress as numerically shown in Equation (5.8). The dashed lines represent the 95 % error confidence intervals. The best fit line runs approximately along the x-axis.

If it can be assumed that b is for a fact equal to 0 (which it can not), then the effect of stress upon $\alpha \times \frac{dn}{dT}$ can be excluded from Equation (5.6) and the regression repeated to determine the effect of only temperature upon $\alpha \times \frac{dn}{dT}$. In this case, α was determined to be 2.791 and the regressions R^2 equal to 0.98. This small change in α (when the effect of stress is not accounted for) helps to highlight the small effect (if any) of stress upon $\alpha \times \frac{dn}{dT}$.

5.7 Conclusion

The effect of stress and temperature upon the product of α and $\frac{dn}{dT}$ for a silicon nitride was investigated. Strong evidence that $\alpha \times \frac{dn}{dT}$ is independent of stress was found, while the relationship with temperature was found to be:

$$\Delta\left(\alpha \times \frac{dn}{dT}\right)_T = (2.774 \pm 0.214)\ \%/\mathrm{K} \times \Delta T \quad \text{for} \quad 290\ \mathrm{K} \leq \mathrm{T} \leq 312\ \mathrm{K}\,, \tag{5.9}$$

Temperature and mode frequency measurements, supported by finite element models of the system, yielded a membrane thermal conductivity of (23 ± 3) W/mK and a thermal expansion coefficient of $(1.4\ \pm\ 0.2) \times 10^{-6}\ 1/\mathrm{K}$.

These results suggest that using silicon nitride in a highly stressed state which is known to have a very low mechanical loss, should not present problems with excess optical absorption. Silicon nitride is therefore of great interest for further investigation as a coating material for use in gravitational wave detectors.

In the future, further studies should look at identifying the individual effect of stress upon α and $\frac{dn}{dT}$ as both of these relationships are important in understanding whether the thermal noise of a gravitational wave detector could be reduced using high or low stressed coatings. To separate the individual effect of stress upon α and $\frac{dn}{dT}$, a further experiment would have to be performed in conjunction with this work. This experiment would be to clamp a membrane and pass an angled beam of light through it as its stress is gradually changed. This stress would change the angle in which the light passes through the membrane as a result of its refractive index changing. This information informs us of the relationship between $\frac{dn}{dT}$ and stress. Using this result in conjunction with the work presented here would enable the individual effect of stress upon α to be determined.

Whilst this experiment was performed at room temperature, the interest in the results are predominately concerned with how silicon nitride behaves at cryogenic temperatures. However, due to a lack of knowledge regarding the thin film properties of silicon nitride at cryogenic temperatures, it is not possible to infer from the room temperatures results as to what the effects would be at cryogenic temperatures. Once silicon nitride's material properties are well known at cryogenic temperatures, this experiment should be repeated.

6 Silicon nitride membranes

6.1 Introduction

Coating loss measurements are often limited by substrate effects as they require comparing the loss of the coated and uncoated samples. Low loss substrates which are thin (closer in thickness to the coating) are desirable in order to improve measurements. The Institute of Gravitational Research has recently gained access to an operational coating chamber. The usual substrate geometries (discs and cantilevers) are not compatible with the coating chamber due to their size. In this chapter, investigations into the possibility of using thin, low-loss silicon nitride membranes as coating substrates for this chamber are presented. If silicon nitride membranes can be successfully used as substrates that can be coated and used in mechanical loss experiments, then this represents a fantastic opportunity for the group as it will eliminate the secrecy which currently shrouds the coating deposition process when external coating companies are contracted to do a job. By being able to exert full control over the deposition process, the effect of individual deposition parameters (such as temperatures, gases within the chamber, ion energies etc) upon a coating's bulk and shear loss can be more easily investigated.

This chapter will outline the criteria of what makes a good substrate in mechanical loss experiments before using this criteria to compare membrane performance against other commonly used substrates: silicon cantilevers, silica cantilevers and silica discs. The bulk and shear losses of a silica coating deposited upon a membrane and cantilever substrate will then be calculated and conclusions will be drawn as to whether membranes could/should be used in future experiments to calculate the bulk and shear losses of a coating. An investigation into the effect of heat treatment upon the optical absorption of silicon nitride membranes at both 1064 and 1550 nm will then be presented. Lastly, the possibility of using a silicon nitride coating in the Einstein Telescope (ET) will be examined by calculating the silicon nitride's coating thermal noise and absorption, and comparing these values to the detector's requirements.

Part of the work presented in this chapter (Section 6.7.2 onwards) was published in Physical Review Letters. The paper was called "Silicon-based optical mirror coatings for ultrahigh precision metrology and sensing" **[181]**.

6.2 Criteria of an ideal substrate used in mechanical loss experiments

The ideal substrate used in mechanical loss experiments (explained in Section 3.3) has:

1. no sensitivity to clamping conditions,
2. the same mechanical loss as nominally identical substrates,
3. as low a mechanical loss as possible.

1. Before the mechanical loss of a coating can be determined, the loss of the substrate which the coating is deposited upon must be known. To measure the loss of a substrate, the substrate must be supported in some way (in Chapter 3 the cantilevers were clamped, whilst the disks were suspended using silica fibres) which does not create additional mechanical loss. By re-clamping/ re-suspending the substrate and taking more loss measurements, it is possible to estimate the true loss of the substrate. However, re-clamping/re-suspending the substrates is not optimum and comes at a cost, most notably it is not time efficient and it leads to a higher probability of damaging the substrate. Therefore, the ideal substrate to use in mechanical loss experiments is one which is insensitive to clamping conditions.

2. Having to measure every substrate before a coating is deposited upon its surface is extremely time consuming. The ideal substrate is therefore one which can exhibit the same mechanical loss as nominally identical substrates as this would eliminate the need to measure all samples.

3. In a mechanical loss experiment, the loss of the substrate and coated substrate is measured, before the loss of the coating can be determined. If the substrate has a very high loss, the addition of a coating will only change the loss of the uncoated sample by a miniscule amount (as the coating is much thinner than the substrate). This small change in loss between coated and uncoated samples is difficult to measure and results in large errors. By reducing the loss of the uncoated substrate, the measurement becomes more sensitive to the effect of the coating. It is therefore ideal to have a substrate with as low a loss as possible.

These criteria will now be used to assess the performance of membranes as a substrate in mechanical loss experiments compared to other commonly-used substrates such as: silicon cantilevers, silica cantilevers and silica discs.

6.3 Performance of membrane substrate

6.3.1 Sensitivity to clamping conditions

One edge of a membrane's frame (Figure 5-1) was clamped using a steel clamp contained in a vacuum chamber (as described in Chapter 3) and the mechanical loss of the membrane's modes were measured. The membrane was re-clamped multiple times and the measurements repeated. This process was repeated for another three nominally identical membranes. Figure 6-1 shows the losses of one of the membranes that was clamped multiple times. The spread in modal losses shown in Figure 6-1 was observed to be similar for each of the membranes measured.

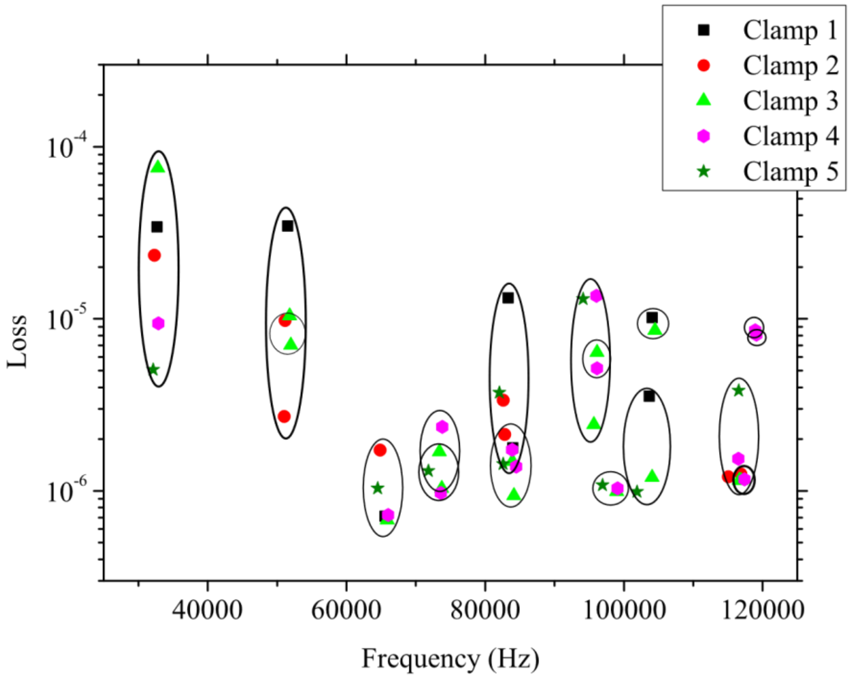

Figure 6-1: Measured losses of an uncoated membrane. The membrane was clamped 5 times and the losses of the various modes measured each time. Black ellipsoids are used to indicate all losses measured for a particular mode.

To determine the sensitivity of the membrane loss to clamping conditions, a standard deviation of the loss measurements for each mode was calculated. The average of all of these standard deviations (for all modes and all membranes) was then calculated and is listed in Table 6-1. This value is called the membrane's sensitivity value.

Sample	Sensitivity value
Membrane	3.290×10^{-6}
Silica cantilever	2.913×10^{-6}
Silicon cantilever	6.393×10^{-6}
Silica disc (wire suspension)	4.7×10^{-8}
Silica disc (GeNS)	3.1×10^{-8}

Table 6-1: Sensitivity values of different substrates. Substrates with a higher sensitivity value show a greater spread in modal losses due to clamping. Silicon cantilevers are the most sensitive to clamping conditions whilst silica disc are the least affected.

Sensitivity values were also calculated for silica cantilevers, silicon cantilevers and silica discs, and are shown in the table. The loss of 4 nominally identical silica cantilevers was measured. Each cantilever was clamped multiple times and each time its loss was re-measured. Two of the silica cantilevers showed little sensitivity to clamping (as shown in Figure 3-39) whilst the other two showed a greater sensitivity (as shown in Figure 6-2).

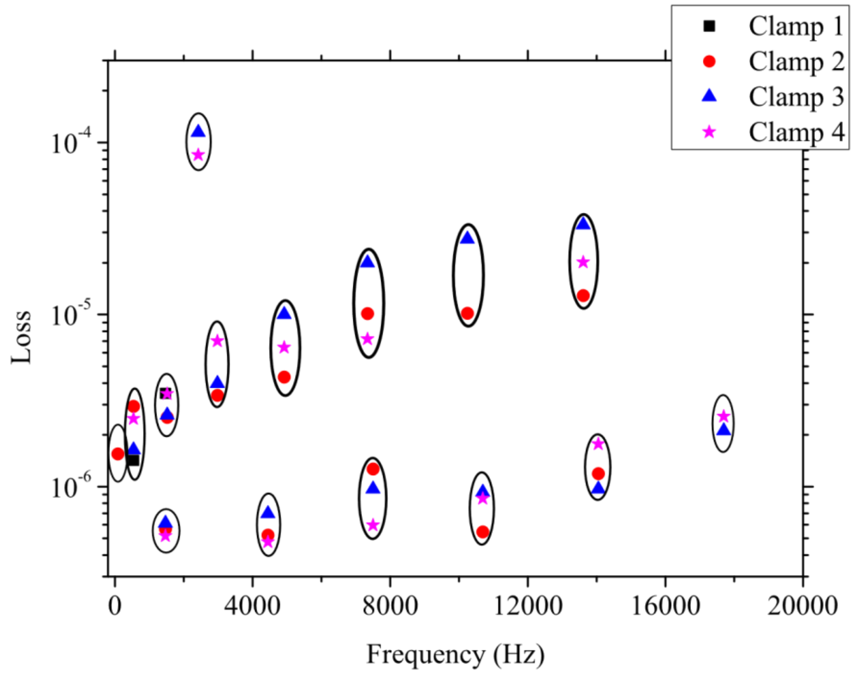

Figure 6-2: Measured losses of a silica cantilever. The cantilever was clamped 4 times and the losses of the various modes measured each time. Black ellipsoids are used to indicate all losses measured for a particular mode.

This process was similarly repeated for 4 nominally identical silicon cantilevers. Figure 6-3 shows the losses of one of these cantilevers.

The loss of a silica disc can be measured using either the suspension setup described in Chapter 3 or using a GeNS (Gentle Nodal Suspension) setup where the disc is supported using a ball bearing positioned at the centre of the disc [182]. The losses of 7 different silica discs were measured using the suspension setup (4 discs) and the GeNS setup (3 discs). In both setups, discs were re-clamped multiple times and the losses remeasured after each new clamp. Figure 6-4 shows the measured losses of one of the discs using the suspension setup.

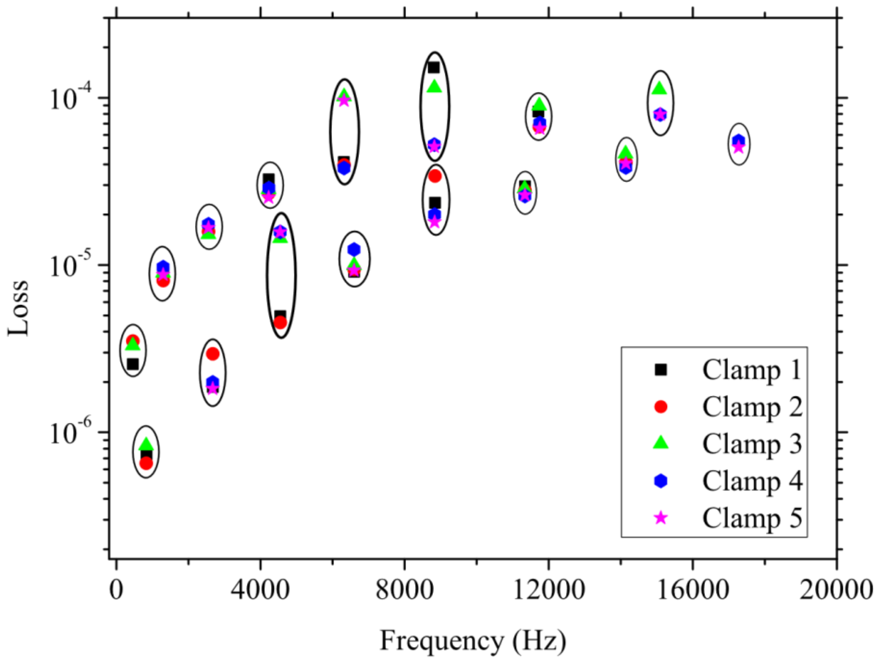

Figure 6-3: Measured losses of a silicon cantilever. The cantilever was clamped 5 times and the losses of the various modes measured each time. Black ellipsoids are used to indicate all losses measured for a particular mode.

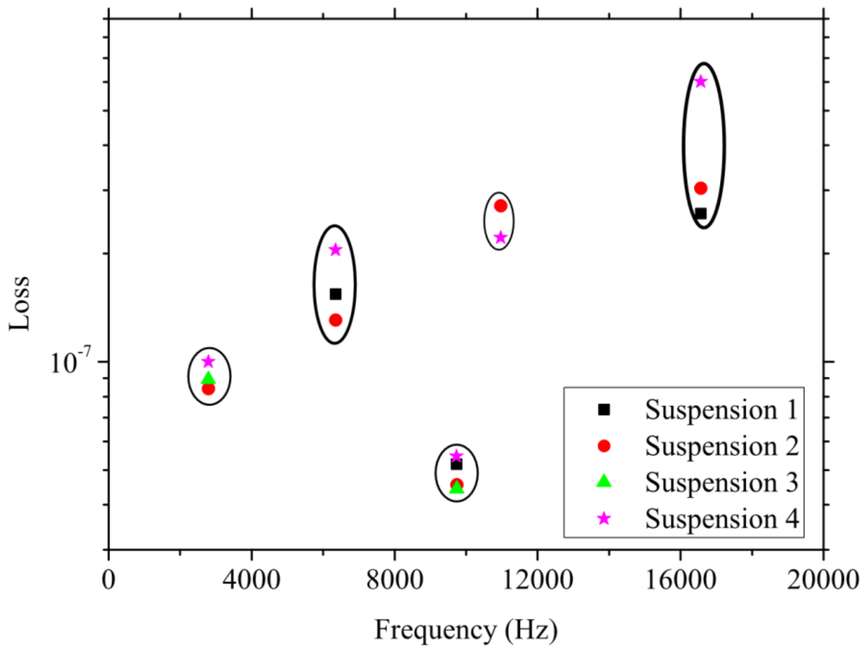

Figure 6-4: Measured losses of a silica disc (using the silica wire suspension setup). The disc was suspended 4 times and the losses of the various modes measured each time. Black ellipsoids are used to indicate all losses measured for a particular mode.

The losses of three silica discs that had each been suspended multiple times using the GeNS setup were measured. Figure 6-5 shows the measured losses for one of the discs.

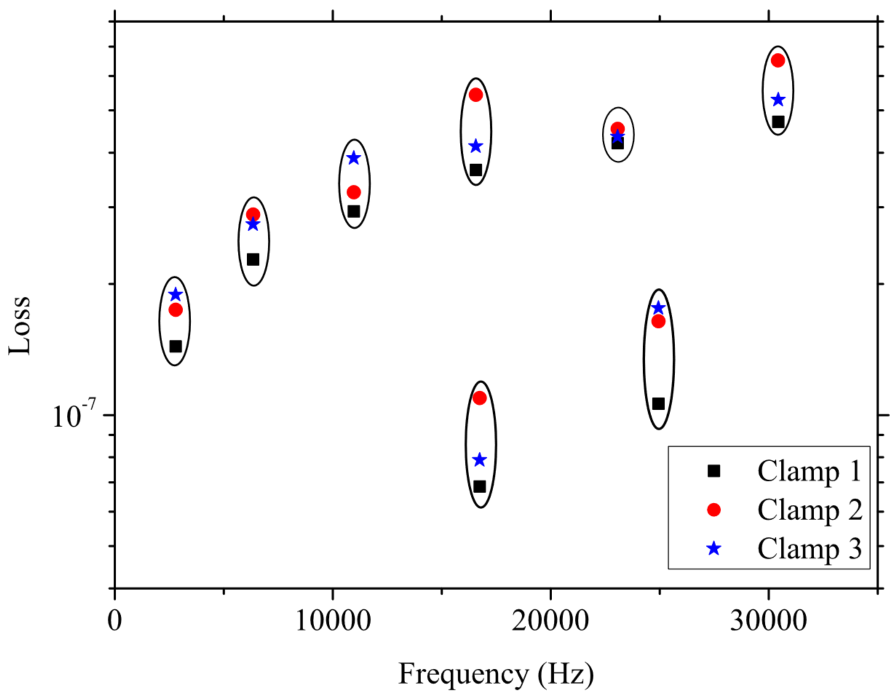

Figure 6-5: Measured losses of a silica disc (using the GeNS setup). The disc was suspended 4 times and the losses of the various modes measured each time. Black ellipsoids are used to indicate all losses measured for a particular mode.

Comparing all of the sensitivity values in Table 6-1, it can be seen that to minimise clamping effects in loss experiments, it would be optimum to use a silica disc (and GeNS setup). Silicon cantilevers were found to be the most sensitive substrate to re-clamping effects.

6.3.2 Similarity to nominally identical substrates

Figure 6-6 shows the lowest measured losses for a number of modes of 4 nominally identical silicon nitride membranes.

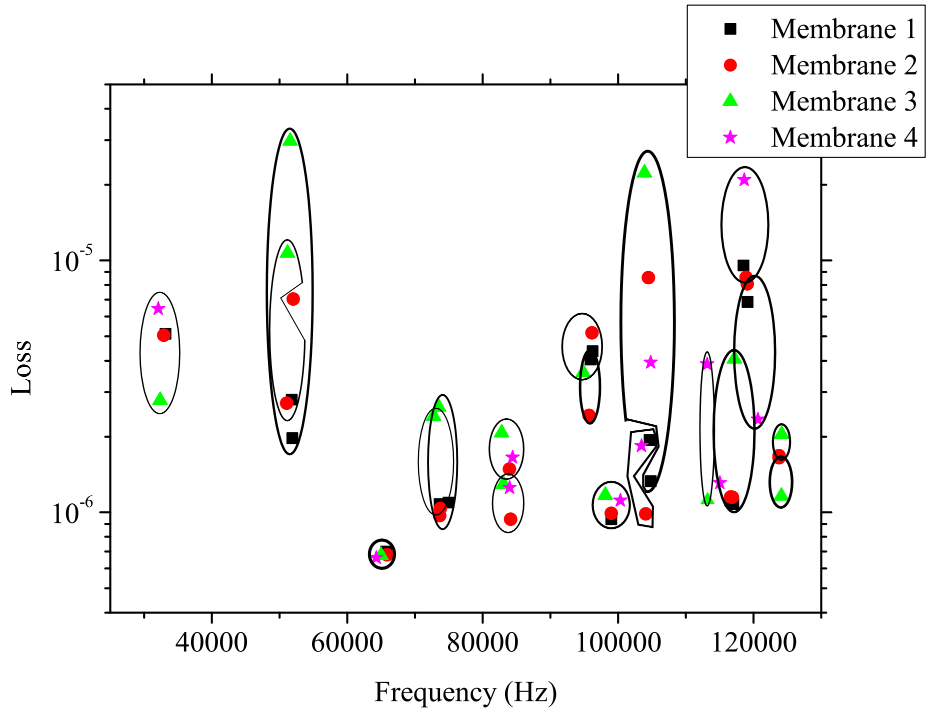

Figure 6-6: Lowest modal losses measured for 4 nominally identical membranes. Black outlines are used to group the same mode's measured losses.

To determine the similarity of the membranes' losses, a standard deviation of the loss of each mode was calculated. An average of all of the mode's standard deviations was then calculated. This value is listed in Table 6-2 and is called the dis-similarity value (as a higher value means the substrates are less similar). This calculation was similarly performed on 4 silicon cantilevers, 4 silica cantilevers, 4 silica discs (using the silica wire suspension technique) and 3 silica discs (using the GeNS setup). Figures 6-7. 6-8, 6-9 and 6-10 show the lowest measured losses for the nominally identical samples (silicon cantilevers, silica cantilevers, silica discs (using the silica wire suspension technique) and silica discs (using the GeNS setup) respectively). The dis-similarity values for these substrates are listed in Table 6-2.

Sample	**Dis-similarity value**
Membrane	2.104×10^{-6}
Silica cantilever	3.554×10^{-6}
Silicon cantilever	3.155×10^{-6}
Silica disc (wire suspension)	3.6×10^{-8}
Silica disc (GeNS)	2.5×10^{-8}

Table 6-2: Table listing the dis-similarity values of the different substrates. Silica discs measured using the GeNS setup show the greatest similarity in loss whilst the silica cantilever shows the least similarity.

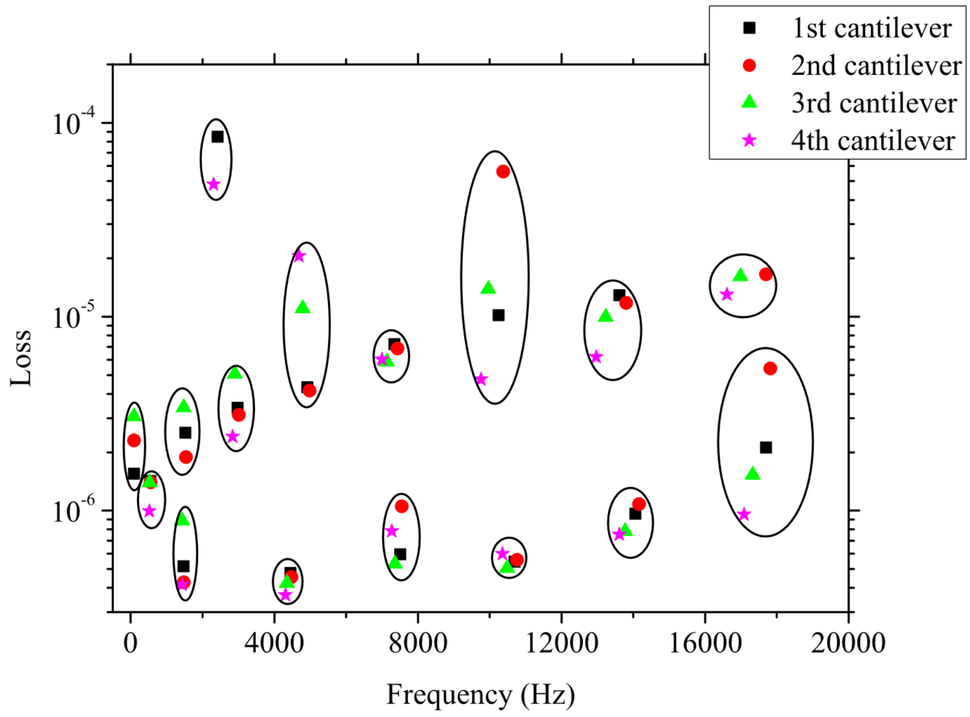

Figure 6-7: Lowest modal losses measured for 4 nominally identical silica cantilevers. Black ellipsoids are used to group the same mode's measured losses.

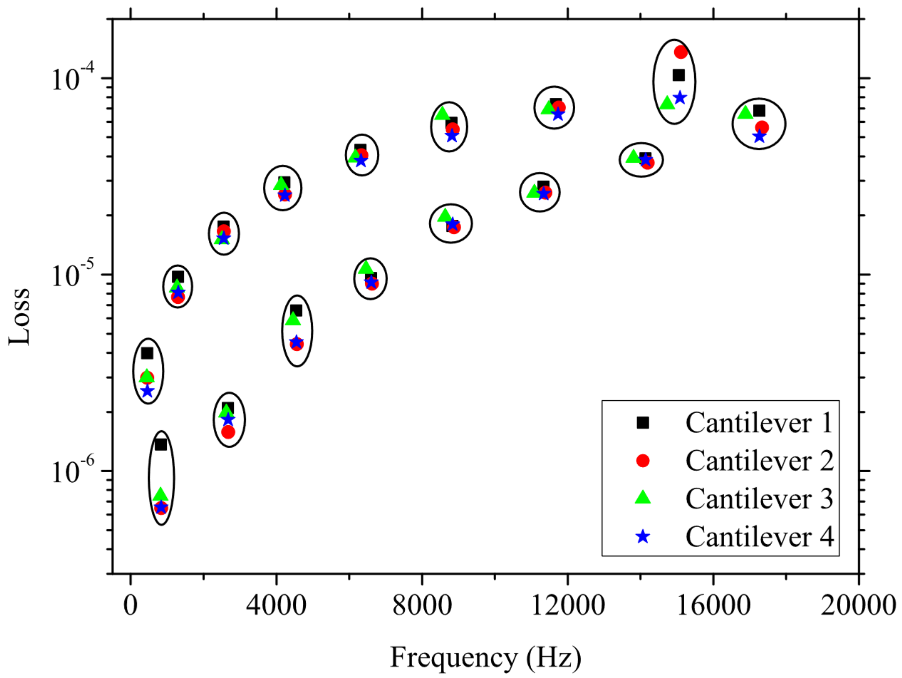

Figure 6-8: Lowest modal losses measured for 4 nominally identical silicon cantilevers. Black ellipsoids are used to group the same mode's measured losses.

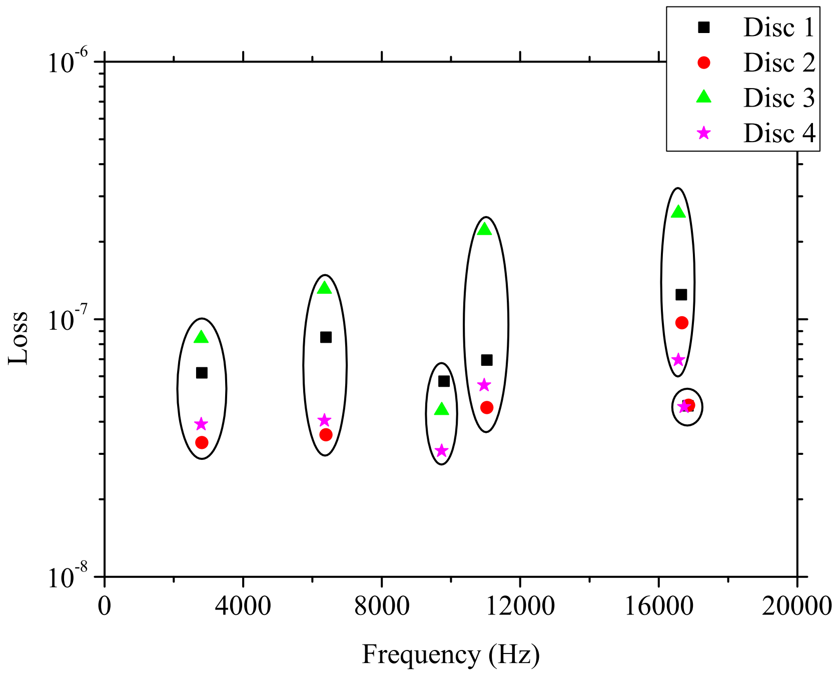

Figure 6-9: Lowest modal losses measured for 4 nominally identical silica discs. The losses were measured using the silica wire suspension setup. Black ellipsoids are used to group the same mode's measured losses.

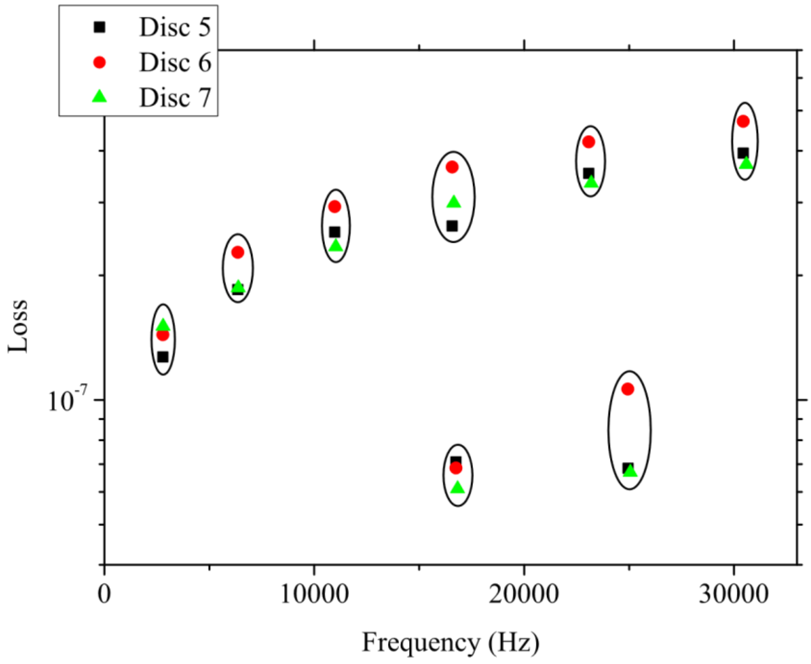

Figure 6-10: Lowest modal losses measured for 3 nominally identical silica discs. The losses were measured using the GeNS setup. Black ellipsoids are used to group the same mode's measured losses.

Table 6-2 shows that silica cantilevers exhibit the least similarity between their measured mechanical losses whilst silica discs (that were measured in the GeNS setup) showed the most similarity.

6.3.3 Magnitude of loss

Due to the frequency dependence of the different substrates' modal losses, it is difficult to rank which substrate is best to use in mechanical loss experiments based purely upon the magnitude of their losses.

An attempt to rank their performance can be made by first taking a simple average of the substrates' modal losses. Table 6-3 lists the average modal loss of the different substrates.

Sample	Average loss
Membrane	3.91×10^{-6}
Silica cantilever	8.92×10^{-6}
Silicon cantilever	3.010×10^{-5}
Silica disc (wire suspension)	8×10^{-8}
Silica disc (GeNS)	2.2×10^{-7}

Table 6-3: Table listing the average modal loss of the different substrates. Silica discs measured using the wire suspension setup have the lowest loss whilst silicon cantilevers have the highest loss.

The silica discs exhibit the lowest average modal loss. In Figure 6-6 to 6-10, it can be seen that the highest modal loss of a silica disc is lower than almost every other substrate's individual modal losses (with the exception of 3 silica cantilever modes). Furthermore, it can also be observed that silica discs have the lowest measured substrate modal loss. For the range of frequencies investigated, it can be argued that based upon the sole criteria of a substrate's magnitude of loss, silica discs are the best substrate to use in mechanical loss experiments. Different batches of silica discs were used in the silica wire suspension setup and in the GeNS setup, thus explaining the difference in average loss value. It is suspected that the larger average loss value of the silica discs used in the GeNS setup, is caused by some of the discs having unpolished edges (as polished edge discs have been shown to have lower loss than unpolished edge discs [144]).

In Table 6-3 silicon cantilevers exhibit the highest average modal loss. In Figure 6-6 to 6-10, the lowest silicon cantilever modal loss can be seen to be higher than the lowest modal loss of each of the other substrates. Furthermore, it can also be observed in these graphs that the silicon cantilever has the highest measured modal loss for all substrates. It can therefore be argued that silicon cantilevers are the worst substrate to use in mechanical loss experiments based solely upon the criteria of the magnitude of a substrate's modal losses (for these frequencies investigated). It is difficult to rank the performance of the membrane substrate relative to the silica cantilever, however their performance is somewhere between the silica discs and silicon cantilevers.

6.3.4 Conclusion

Based upon a substrate's sensitivity to clamping, similarity of loss to nominally identical substrates and magnitude of modal losses, silica discs can most likely make the claim to be the best substrate to use in mechanical loss experiments. Silicon cantilevers were shown to be the most sensitive to clamping, the second most dis-similar to nominally identical substrates and have the highest average loss. Based upon these results it can be argued that silicon cantilevers are the worst substrate to use in mechanical loss experiments. Membranes were observed to have the lowest average loss after silica discs, the second most sensitive substrate to clamping and have the second least variation of loss for nominally identical substrates (behind silica discs). Furthermore, one criteria of a substrate which was not discussed earlier but indicates that membranes would be the optimum substrate to use in mechanical loss experiments, is the thickness of the substrate. Thinner substrates make better substrates in mechanical loss experiments as more of the vibrational energy is stored within the coating. This means mechanical loss measurements are most sensitive to directly determining the loss of the coating and are less sensitive to measurement errors. Membranes are the thinnest of the substrates tested and for this reason are the most sensitive to measuring the loss of the coating. This greater sensitivity to the coating can be observed by calculating the average ratio of energy stored in the substrate and coating for a 500 um thick tantala coating deposited upon a silicon cantilever (54), silica cantilever (48), silica disc (860) and silicon nitride membrane (3). It can be observed that the membrane is the most sensitive to the loss of the coating as it has the lowest energy ratio as a greater proportion of the vibrational energy is stored in the coating,

To conclude, if any substrate can be selected to be used in a mechanical loss experiment, then preference should probably be shown in using silica discs before membranes. However, when membranes have to be used as a substrate (as in the case when the coating deposition chamber is very small), membrane performance is expected to be better than silicon cantilevers.

6.4 Reducing the membrane clamping effect

The loss results in Section 6.3 were measured with the membranes clamped, from the top and bottom, along one of their silicon frame edges. However, as described in Chapter 5, the silicon nitride membrane component extends over the silicon frame. Since the membrane itself was being clamped, it was suspected that the clamping process could be changing the membrane's mechanical loss as a result of "stick-slip" losses (which describe frictional losses associated with energy being lost from the oscillator into the clamp [183]). To limit the possible effect of clamping upon the membrane's losses, a membrane holder was designed. This holder enabled the membrane to be fixed in position (essential for loss measurements to be made) but did not require the membrane to be directly clamped or touched. Instead the sides of the membrane's frame (where no membrane is present) was held stationary by the holder. One of the holder's edges was then clamped using a steel clamp in a vacuum tank and the losses of the membrane measured. The holder that was tested is shown in Figure 6-11. Clamping the holder and designing its thickness to be > 1000 times the membrane thickness, helped to attenuate any "stick-slip" losses (as the magnitude of these losses are proportional to the ratio between the oscillator thickness and clamped objects thickness).

Using this holder had the effect of increasing the membrane's sensitivity value (making it more sensitive to clamping). Furthermore, the holder also had the effect of increasing the measured losses of the membrane resulting in them being at least 6 times larger than when they were not used. In summary, the spread in membrane losses due to clamping could not be reduced by using this holder. Other holders were also investigated. These had similar effects upon the measured losses. In one measurement, the membrane was left unclamped and just resting upon a metal block in the vacuum chamber. This too resulted in much larger mechanical losses being measured.

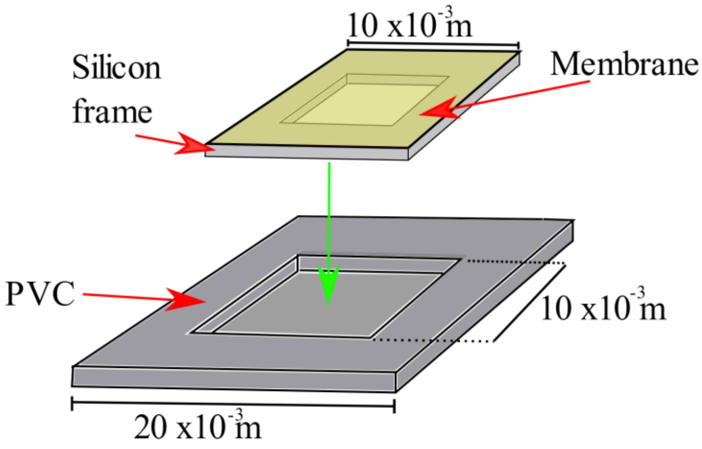

Figure 6-11: Schematic of vacuum compatible PVC membrane holder. The membrane was snapped into the holder which held it fixed.

6.5 Coating bulk and shear loss using a membrane

The ultimate test as to whether a membrane can be used as a substrate in mechanical loss experiments is to coat the membrane and another type of substrate with the same coating and calculate the coating's bulk and shear losses for both substrates.

A 500×10^{-9} m thick silica coating was deposited upon a membrane and silica cantilever. The losses of the coated membrane as well as its thermoelastic loss are shown in Figure 6-12. The thermoelastic loss of the silica coated membrane plus some offset is not equal to the measured coated losses.

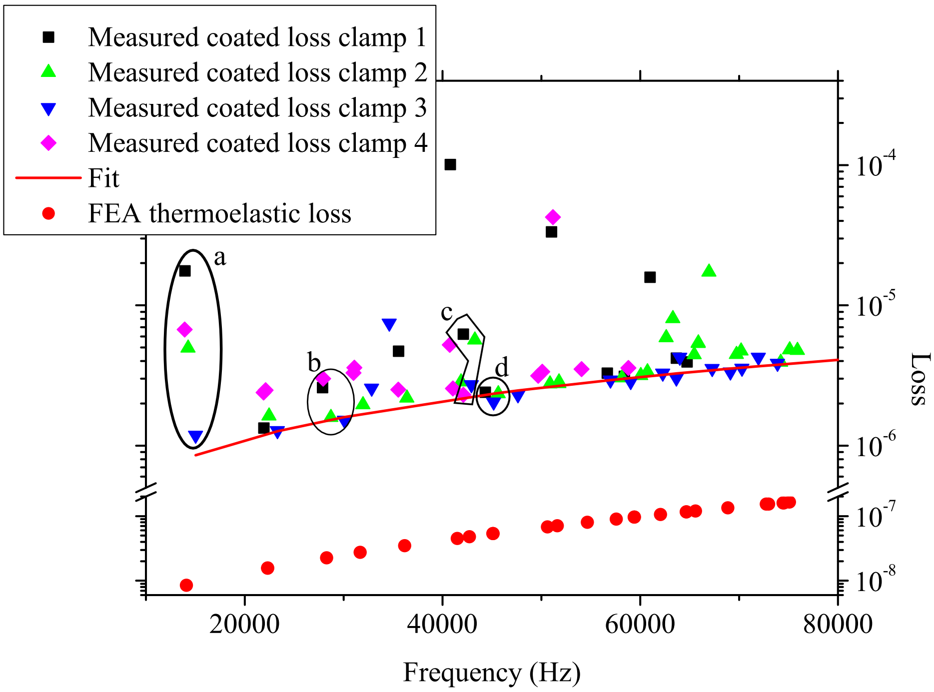

Figure 6-12: Coated membrane's measured losses. The coated membrane was clamped 4 times and its losses measured after each clamping. The 4 coated modes labelled a to d will be used to calculate the bulk and shear losses of the coating. The fitted line fits to the lowest measured modal losses. It too will be used to calculate the coating's bulk and shear losses. The FEA thermoelastic loss of the coated membrane's modes is also shown.

The uncoated losses of the membrane were not measured before the coating was deposited. To calculate the bulk and shear losses of the coating, it was assumed that the uncoated losses of the membrane were equal to the average of the previously measured four membranes' lowest modal losses (shown in Figure 6-6). Figure 6-13 shows the uncoated membrane's assumed losses (as well as it's calculated thermoelastic loss).

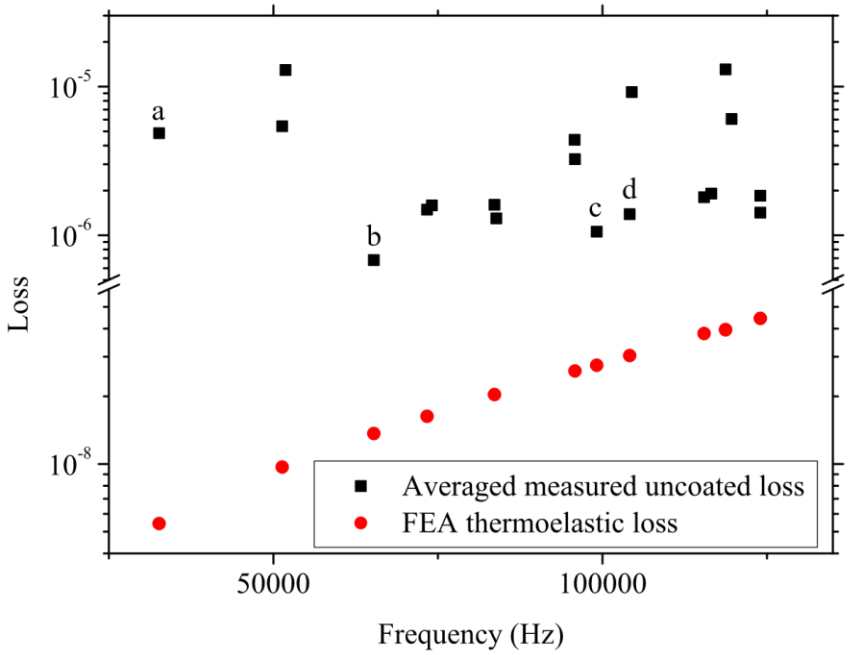

Figure 6-13: Assumed uncoated membrane losses. Each modal loss was calculated by averaging the 4 uncoated membrane modal losses shown in Figure 6-6. The 4 uncoated losses labelled a to d will be used to calculate the bulk and shear losses of the coating. The FEA thermoelastic loss of the uncoated membrane's modes is also shown.

The mode frequencies of the coated membrane are approximately 42 % of those of the uncoated membrane (see modes labelled a to d in Figures 6-12 and 6-13). This is due to the total net force of the coated and uncoated membranes being different from one another (as it can be observed in Equation (5.4) that the stress of a membrane determines it's mode frequencies). The coated membrane has a total net force (of 0.33 N) which is lower than that in the uncoated membrane (of 1.34 N). These net forces were determined by an FEA model, in which the net force was adjusted until the mode frequencies agreed with the measured values. This process, and the best match obtained to the experimental data, is illustrated in Figures 6-14 and 6-15.

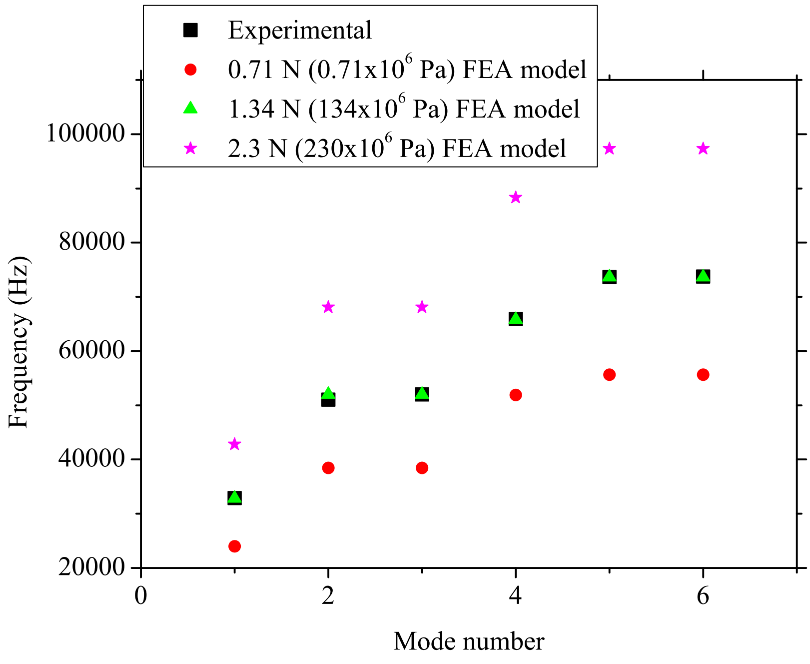

Figure 6-14: Measured and computationally determined uncoated membrane mode frequencies. The computationally determined mode frequencies were calculated for a range of differently stressed uncoated membranes. The total net force of the membrane and its stress are listed in the legend.

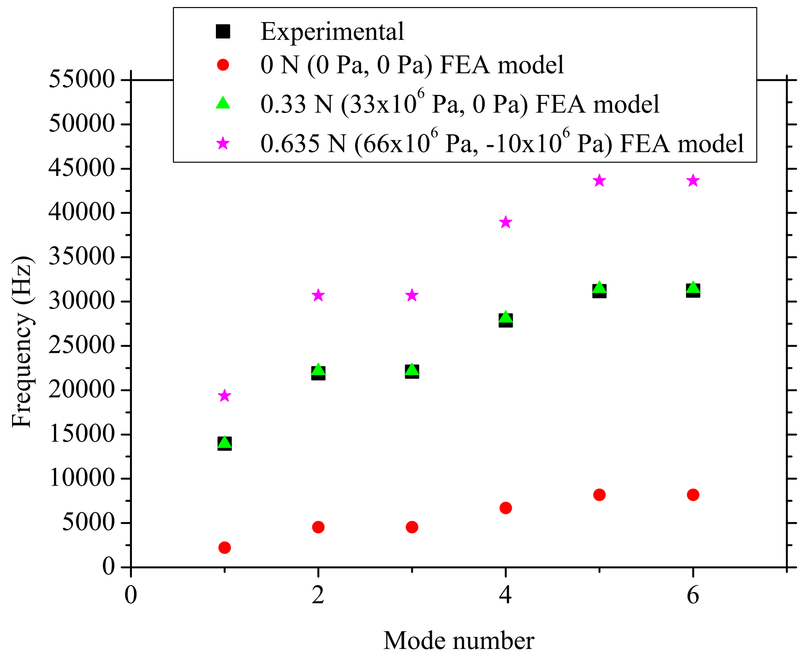

Figure 6-15: Measured and computationally determined coated membrane mode frequencies. The membrane was coated with a 500×10^{-9} *m silica layer. The computationally determined mode frequencies were calculated for a range of differently stressed coated membranes. The total net force of the coated membrane, the stress of the membrane and the stress of the coating are listed in the legend (in this order).*

To determine the mode-dependent silica coating losses and the coating's bulk and shear losses, the ratio of energies stored in the membrane and coating must be known. In Chapter 3, it was observed that this energy ratio for coated cantilevers could be affected by the stress and curvature of the coated cantilever. It is therefore necessary that when the membrane's energy ratios are calculated, the stress in the membrane and coating are properly accounted for (however, the effects of curvature were not investigated here, as there is no obvious membrane curvature and if it is curved, it is very small). Whilst it can be observed in Figure 6-15 that the net force in the coated membrane is 0.33 N (as it best replicates the measured mode frequencies), the exact individual stresses in the membrane and coating are unknown. In Figure 6-15 it was merely assumed that there was 33×10^6 Pa of stress in the membrane and 0 Pa of stress in the coating, however there are an infinite number of combinations of stress in the membrane and coating which can result in the same net force and same computational frequencies being obtained. Some of these combinations are listed in Table 6-4.

Net force in coated membrane (N)	Stress in membrane $\times 10^6$ (Pa)	Stress in coating $\times 10^6$ (Pa)
0.33	33	0
0.33	26.3	26.3
0.33	134	-404
0.33	0	132
0.33	300	-1068

Table 6-4: Selected combinations of membrane and coating stress which result in the coated membrane having the same net force of 0.33 N.

The coated membrane's energy ratios were calculated for the different combinations of membrane and coating stress listed in Table 6-4. The values are shown in Figure 6-16.

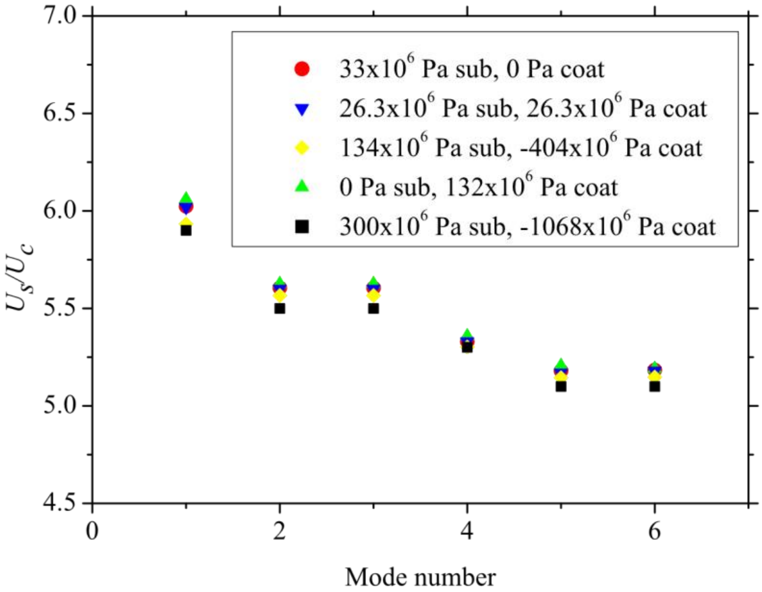

Figure 6-16: The energy ratio (between substrate and coating) calculated for different stresses in the coating and membrane.

It can be observed that for the range of stresses tested, the coated membrane's energy ratio changes by less than 4 % for all modes. Whilst coated membrane models which have larger absolute values of stress in the coating and membrane may exhibit a change in energy ratio greater than 4 %, it is unlikely that the silica coating in reality experiences a stress more negative than the model which assumed a value of -1068×10^6 Pa. This is expected as the same coating deposited upon silicon cantilevers (using the same deposition method) was measured to have a stress between -400×10^6 Pa and -700×10^6 Pa [54]. Since silicon nitride has a thermal expansion coefficient more similar to silica than silicon to silica (and the magnitude of the stress in the coating is related to the mismatch in the coating and substrate's thermal expansion coefficients), the stress of the silica coating on a silicon nitride substrate should be less negative than that deposited on a silicon substrate. Under the assumption that the stress of the silica coating is not less than -1068×10^6 Pa, the coating's mode dependent coating losses were calculated. The mode-dependent losses of the coating were determined using an average of the differently stressed model's energy ratios (shown in Figure 6-16).

Whilst many modes of the coated membrane were measured, only 4 mode-dependent coating losses were able to be determined (labelled **a** to **d** in Figures 6-12 and 6-13). This was because it was often not possible to measure the loss of the coated membrane's degenerate mode pairs meaning most of the coated losses could not be matched to their corresponding uncoated loss. Furthermore, if the loss of a degenerate mode pair was measured, then it was often only measured for one clamp, implying little confidence in the measured loss of the mode. Modes whose loss was only measured once (for one clamp) were not used to calculate the coating's mode dependent loss. The higher loss mode in a degenerate mode pair was also not used to calculate coating loss (as it was suspected that this loss exhibited higher loss due to coupling with the frame). One of the four calculated mode dependent coating losses were determined to be negative, implying that the assumed uncoated loss of the membrane was too large (**a** in Figures 6-12 and 6-13). The three non-negative mode dependent coating losses are shown in Figure 6-17 along with their one standard deviation random errors.

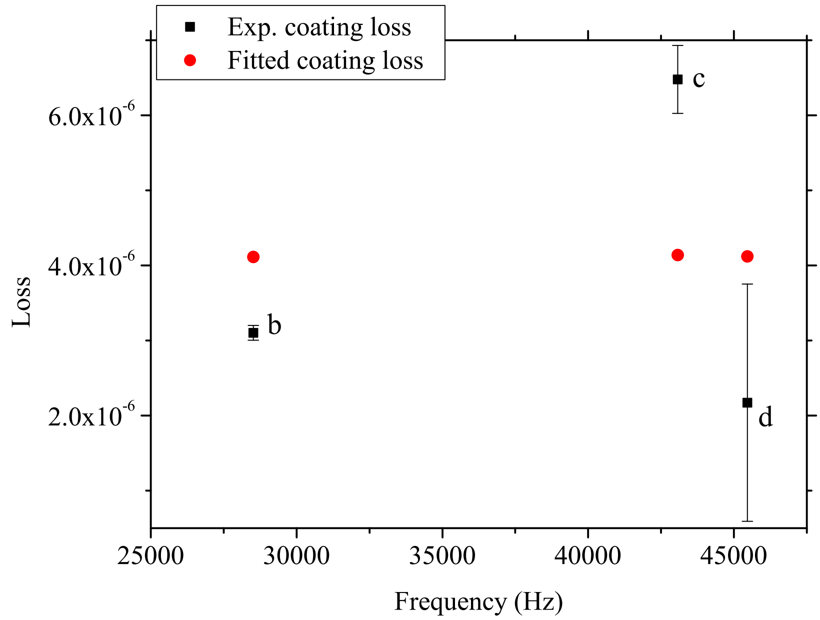

Figure 6-17: Experimentally determined loss of the silica coating deposited upon a membrane. The error bars represent the one standard deviation random error of the mode dependent coating losses. The best fit mode dependent coating losses were used to calculate the bulk and shear losses of the silica coating (described in Chapter 3).

To calculate the bulk and shear losses of the silica coating, the relative energy stored in bulk and shear motion had to be determined for each mode. However as described in Section 3.6.3, it is not possible to calculate the ratio of coating energy stored in bulk energy and shear energy when the FEA model being solved is stressed (as this leads to unphysical negative energies). These ratios were therefore calculated using an unstressed model. The calculated values of the silica coating's bulk and shear losses are therefore only relevant if these ratios do not change with coating and membrane stress. Figure 6-17 shows the quality of the fit used to determine the bulk and shear losses as well as the one standard deviation random errors. The silica bulk and shear losses were determined to be $(0.2 \pm 1.0) \times 10^{-7}$ and $(6.00 \pm 1.00) \times 10^{-6}$ respectively (where the errors listed are total errors, combination of systematic and random). For comparison, the same coating on a silica cantilever substrate was analysed, giving a bulk and shear loss of $(9.76 \pm 1.70) \times 10^{-4}$ and $(2.51 \pm 0.40) \times 10^{-4}$ (where the total errors are again listed). Figure 6-18 shows the experimental loss of the silica coating determined on the silica cantilever. The quality of the fit used to determine the bulk and shear losses of the coating are also shown in this figure.

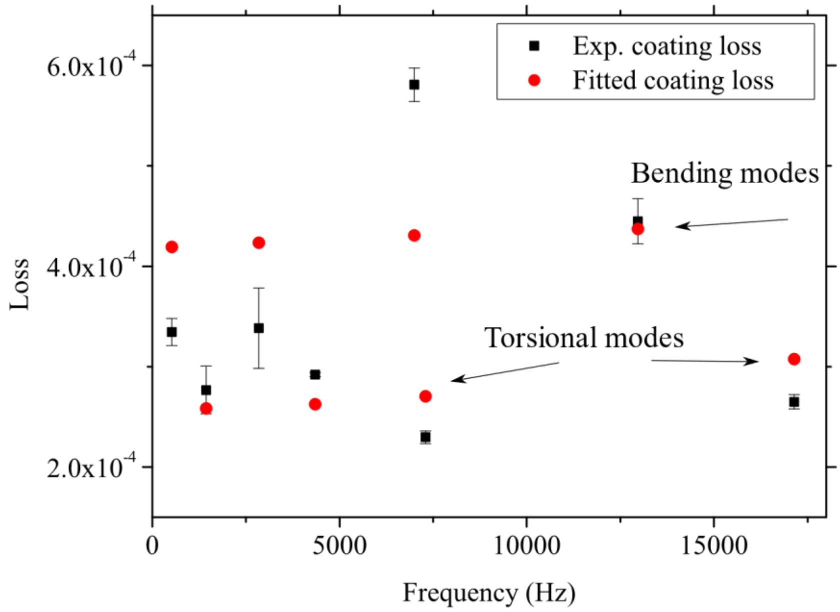

Figure 6-18: Experimentally determined loss of the silica coating deposited upon a silica cantilever. The error bars represent the one standard deviation random error of the coating losses. The fitted coating losses were used to calculate the bulk and shear losses of the silica coating (described in Chapter 3). The four modes with the highest losses are all bending modes, whilst the four with the lowest losses are torsional modes.

The two sets of bulk and shear losses are significantly different from one another. The bulk loss of the coating deposited upon the membrane substrate is approximately 50000 times less than that calculated with the silica cantilever substrate, whilst the shear loss determined for the coating on the membrane substrate is 40 times less than that obtained using the silica cantilever. Furthermore, the coating's shear loss is determined to be significantly larger than its bulk loss when the membrane is used as the substrate whilst the opposite is observed when a silica cantilever is used. It is suspected that the most likely reasons for this difference in results are: the ratio of coating energy stored in bulk energy and shear energy varying with membrane and coating stress; that the loss of the coating deposited upon the membrane should be calculated using uncoated and coated losses which are of the same frequency; that the uncoated losses of the membrane were overestimated as they are unknown and some effect of stress upon thermoelastic loss.

The bulk and shear losses of the silica coating deposited upon the membrane were recalculated using coated and uncoated membrane losses of the same frequency to determine whether this explains the difference in the bulk and shear losses presented earlier for the membrane and cantilever substrates. In Figure 6-12, there is a general trend of the lowest coated losses following a smooth upward trend with frequency. A line was therefore fitted through the lowest modal coated losses which was then used to estimate a coated loss at each of the measured uncoated mode frequencies. Using the same energy ratios as before (which assumed that the coated membrane had a net force of 0.33 N), the loss of the coating was determined at each uncoated mode frequency. The non-negative coating losses are shown in Figure 6-19. It should be noted that just as with the previous analysis, the losses of the coating were not determined using the higher loss mode in a degenerate mode pair. The one standard deviation random errors of the coating loss are shown. To determine these errors, it was assumed that the error in the coated loss was equal to the fitted lines 67 % confidence limit.

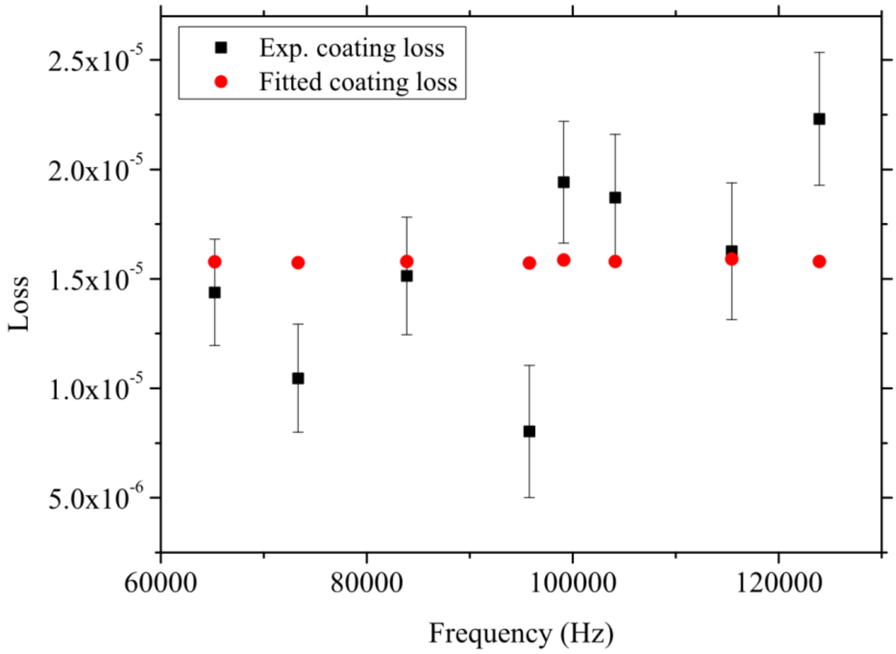

Figure 6-19: Experimentally determined loss of the silica coating deposited upon a membrane. The error bars represent the one standard deviation random error of the mode dependent coating losses. The best fit mode dependent coating losses were used to calculate the bulk and shear losses of the silica coating (described in Chapter 3).

The silica bulk and shear losses were determined to be $(0.2 \pm 1.0) \times 10^{-5}$ and $(2.2 \pm 0.5) \times 10^{-5}$ respectively (where the errors listed are total errors). Comparing these bulk and shear losses to those previously calculated for the coating deposited upon the membrane, it can be observed that the bulk losses are insignificantly different from one another, but the shear losses are significantly different from one another. Furthermore, comparing these new bulk and shear losses to those determined for the coating on the cantilever, it can still be observed that the bulk and shear losses for the different substrates are significantly different from one another. It can therefore be concluded that using coated and uncoated losses of the same frequency to determine the loss of the coating does not explain the difference in the bulk and shear losses of the coating deposited on the membrane and cantilever.

The bulk and shear losses of the coating deposited upon the membrane were re-calculated assuming a zero uncoated membrane loss, to see whether an overestimation in the uncoated loss of the membrane could explain the discrepancy in the different substrate's bulk and shear losses. The modal losses of the coating, shown in Figure 6-19, were therefore recalculated assuming a zero uncoated membrane loss. These new losses of the coating are shown in Figure 6-20. The one standard deviation random errors of the coating loss are shown. To determine these errors, it was assumed that the error in the coated loss was equal to the fitted lines 67 % confidence limit.

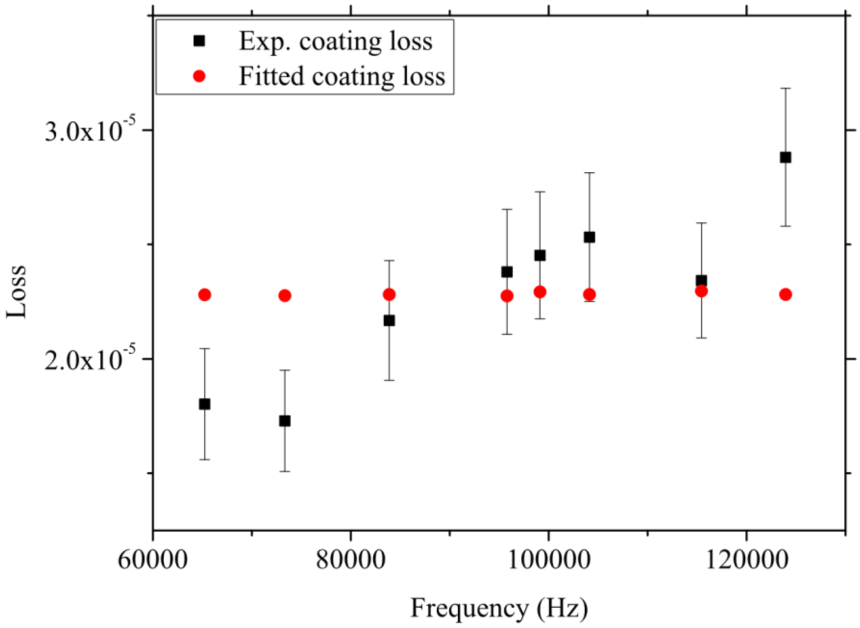

Figure 6-20: Experimentally determined loss of the silica coating deposited upon a membrane. The error bars represent the one standard deviation random error of the mode dependent coating losses. The best fit mode dependent coating losses were used to calculate the bulk and shear losses of the silica coating (described in Chapter 3).

The silica bulk and shear losses were determined to be $(0.8 \pm 2.0) \times 10^{-5}$ and $(2.95 \pm 1) \times 10^{-5}$ respectively (where the errors listed are total errors). This result is not significantly different to the result obtained previously which solely took into account the change in mode frequencies following coating deposition. Furthermore, these new bulk and shear losses are still significantly different from those obtained using a silica cantilever substrate. Overestimating the uncoated loss of the membrane can therefore not explain the difference in the different substrate's bulk and shear losses.

One previously hypothesised explanation was the effect of stress changing the uncoated membrane's thermoelastic loss when the coating is deposited upon it. In Chapter 4, FEA modelling suggested that the stress caused by the deposition of a coating onto a cantilever did not change the cantilever's thermoelastic loss as the net force of the coated cantilever and uncoated cantilever was the same (0 N). However, the coated membrane and uncoated membrane have different net forces as a result of the membrane being fully constrained by the frame (the frame imparts an equal and opposite force to the coated membrane). The thermoelastic loss of the uncoated and coated stressed membranes can be observed in Figure 6-12 and 6-13. It can be observed that the thermoelastic loss of the coated membrane is larger than the loss of the uncoated membrane as a result of stress relief increasing the thermoelastic loss. Using measured uncoated losses to calculate the loss of the coating would therefore overestimate the loss of the coating. However, since the measured losses of the membrane are not dominated by thermoelastic loss this error will be small. What this all means though is that stress effects upon thermoelastic loss cannot explain why the bulk and shear losses of the silica coating deposited upon the membrane are smaller than that calculated for the coating on the silica cantilever.

It is known in the literature that tensile stress can reduce the loss of silica [184]. Based upon the usual arguments of the mismatch in the coating and substrate's thermal expansion coefficients causing the stress in the coating and substrate, it is likely that the silica coating on silicon nitride is under compressive stress. However, this usual argument does not account for the effect of the membrane's frame. If the frame has the effect of inducing a tensile stress in the coating then it could be expected that this tensile stress would be larger than that of the coating deposited upon the silica cantilever (as the cantilever and coating have the same thermal expansion coefficients). If this is true, then the intrinsic loss of the silica coating deposited on the membrane could be less than the intrinsic loss of the coating deposited upon the silica cantilever. This would provide an explanation for the low loss of the coating calculated on the membrane. However, this should be further investigated.

As shown for other coatings investigated in this thesis (Chapter 3), it is not expected that Fejer thermoelastic loss has a significant effect upon the bulk and shear losses of the coating.

Until the difference in the bulk and shear losses of the silica coating deposited upon the membrane and silica cantilever can be explained, membranes should not be used as substrates in mechanical loss experiments.

6.6 Bulk and shear loss: Tantala coating

A 500×10^{-9} m thick tantala coating was deposited upon a membrane and silica cantilever. Attempts were made to determine the bulk and shear losses of the coatings. However, it was not possible to determine the bulk and shear losses of the coating deposited upon the membrane as the coated membrane's losses could not be measured. The coated membrane's losses could not be measured due to the existence of wrinkles across the coated membrane's surface which caused aberrations in the laser light that was used to monitor the membrane's motion (Chapter 3). These wrinkles were most likely formed by the creation of compressive stress in the coated membrane caused by coating deposition [167]. These wrinkles were not present on the silica coated membrane. Figure 6-21 illustrates what the wrinkles across the coated membrane's surface looked like.

Figure 6-21: Image illustrating the wrinkles formed across the coated membrane's surface [185].

6.7 Silicon nitride as a highly reflective coating material

Future gravitational wave detectors that will be operated at cryogenic temperatures (such as the ET) are expected to be limited in sensitivity by the thermal noise of the highly reflective mirror coatings.

Amorphous silicon is being considered as a possible material due to its very low mechanical loss ($< 2 \times 10^{-5}$) at temperatures below 30 K [113]. However, a gravitational wave detector coating made of amorphous silicon/silica would be dominated by the mechanical loss of the silica layers (the currently used low index material in a gravitational wave detector coating) [54]. Finding an alternative low-index material to silica is therefore important.

Silicon nitride has been measured to have a low mechanical loss at both cryogenic [164] and room temperature (shown in Figure 6-6). This section will present the results of an investigation in to whether an improvement in the ET's thermal noise and absorption can be obtained by using a coating made from silicon nitride and amorphous silicon instead of using the Advanced LIGO coating (made from silica and titania doped tantala). This will involve measuring the heat-treated absorption of silicon nitride as well as calculating the thermal noise of an amorphous silicon/ silicon nitride highly reflective coating stack.

The section will conclude by proposing an optimum solution to minimising the absorption and thermal noise in the ET. This solution is to construct a “multi-material” coating [115, 114] made from silica, titania doped tantala, silicon nitride and amorphous silicon. By constructing a multi-material coating where the detector laser light first interacts with the lower absorbing but higher loss bi-layers (silica and tantala) and then the lower loss but higher absorbing bi-layers (amorphous silicon and silicon nitride) underneath them, the absorption and thermal noise of the detector can be optimised. The thermal noise and absorption of this multi-material coating will be compared to the ET’s low frequency (LF) design specifications.

6.7.1 Membrane Absorption

It is important to investigate the effect of heat-treatment upon a silicon nitride membrane’s absorption as there is evidence in the literature of materials showing a reduced absorption with heat-treatment [186]. Heat-treated silicon nitride could therefore be beneficial in reducing a detector’s absorption (if it were to be used as a coating material). Furthermore, if amorphous silicon is to be used as a detector coating material it needs to be heat treated at 450 - 500 °C to reduce its absorption to a reasonable level [186]. If a detector coating is going to be made from both silicon nitride and amorphous silicon (as proposed earlier and will be investigated later) it is important to identify how this level of heat-treatment could affect silicon nitride’s absorption. Whilst the absorption of silicon nitride has been measured in previous research papers [112] as well as this thesis, this is the first time known to the author that the effect of heat-treatment upon membrane absorption has been investigated.

The absorption of a heat-treated silicon nitride membrane was measured at both 1064 nm and 1550 nm using photothermal common path interferometry or PCI for short (described in Chapter 5). The absorption of the membrane was measured as-deposited as well as after it had been heat treated at 200 °C, 400 °C, 600 °C, 800 °C and 900 °C for 3 hours. As mentioned in Chapter 5, it is essential for calibration purposes that the absorption of the membrane is measured when it is aligned relative to the pump and probe beams of the PCI setup at maximum transmission. In Chapter 5, this maximum transmission of the pump and probe laser beams was obtained by angling the membrane at 50 degrees to the probe beam and at 64 degrees to the pump beam. The membrane used for these heat-treated absorption measurements could not be aligned simultaneously to maximise both the pump and probe beams. This is due to the PCI setup easily enabling the angle between the pump beam and membrane to change, the angle between the probe beam and membrane to change but not the angle between the pump beam and probe beam to change. The angle between the pump and probe laser beams in the PCI setup is 14 degrees and out of sheer fortune this angle enabled a maximum transmission of the pump and probe beams to occur simultaneously for the membrane discussed in Chapter 5. However, this angle did not enable the pump and probe beams which passed through the heat-treated membrane to be maximised (it is suspected that this could be due to small changes in the membrane's optical thickness and geometry). To vary this angle between the pump and probe beams would require major changes to the PCI setup. These were not pursued. Instead two absorption measurements were made at each heat-treatment and pump beam wavelength, one when the membrane was aligned to give a maximum probe transmission and the other when the membrane was aligned to give a maximum pump transmission. It is important to note that the PCI signal is proportional to the membrane's absorption in both cases, and while there might be some uncertainty as to the precise value of the absorption, the two sets of values are useful in developing an estimate of membrane absorption and how it changes with temperature. The heat-treated absorption results are presented in Figure 6-22. It should be noted that some of the heat-treated steps only quote an absorption value for either a maximum pump or probe beam as this inability to maximise both the pump and probe beams simultaneously was not thought to be a significant issue during the time when the measurements were taken. A 10 % error has been assumed for each absorption value. This is because a calibration sample's absorption is measured before and after every measurement shown in Figure 6-22. If the sample's absorption changes by more than 10 % then the heat-treated measurements are re-taken.

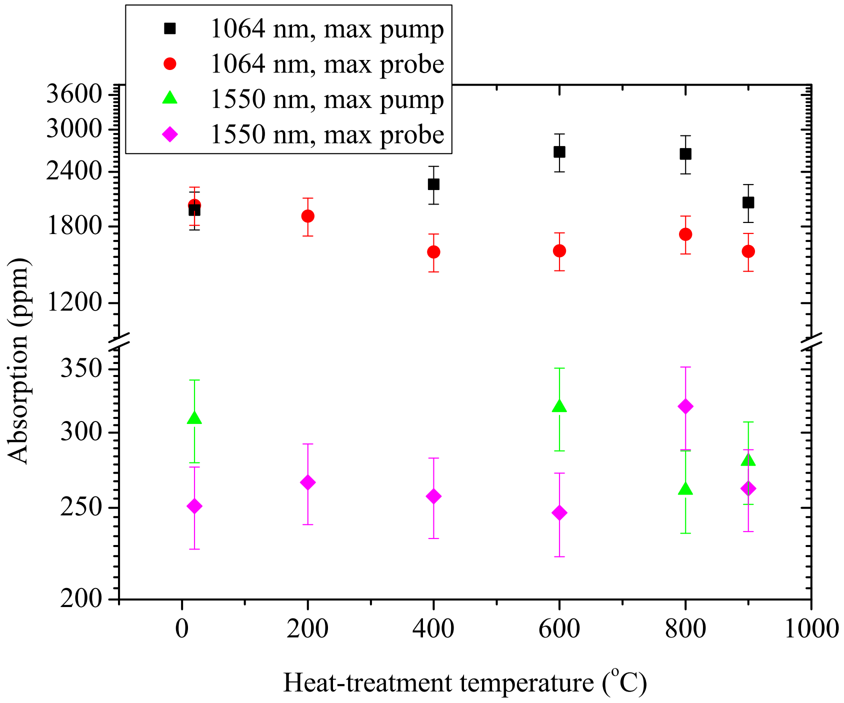

Figure 6-22: Measured absorption of a heat-treated membrane at both 1064 and 1550 nm wavelengths.

From the results shown in Figure 6-22, it is clear that the absorption of the membrane is greater at the 1064 nm wavelength than the 1550 nm wavelength. This is because silicon nitride's bandgap is approximately 5.5 eV [187] (corresponding to a wavelength of 220 nm). Hence as the laser beam wavelength is increased from 1064 nm to 1550 nm, the absorption of the material decreases as the number of photons which can cause the silicon nitride electrons to jump their bandgap decreases at higher wavelength. From the figure, it also appears that heat-treatment (at temperatures between 0 and 900 °C) does not significantly change the absorption of the membrane. However, this is perhaps not surprising since the membrane was fabricated at 900 °C (which was learned post-measurements). In the future, higher heat-treated temperatures should be tested to investigate whether they have an effect upon absorption.

6.7.2 Coating thermal noise theory

As described in Chapter 2, a highly reflective mirror coating is typically made from alternating layers of different materials. In this section, the thermal noise of a multi-layer coating deposited upon a substrate ($S_x(f)$) was calculated using a variation of Equation (2.26) [188]:

$$S_x(f) = \frac{2k_bT}{\pi^{\frac{3}{2}}f}\left(\frac{1-\nu^2}{w_0Y}\right)\left(\begin{array}{c}\phi_{\text{substrate}} + \frac{1}{\sqrt{\pi}}\left(\frac{d}{w_0Y_\perp}\right)\times \\ \left[\begin{array}{c}\left(\frac{Y}{1-\nu_\perp} - \frac{2\nu_\perp^2 YY_\parallel}{Y_\perp(1-\nu^2)(1-\nu_\parallel)}\right)\phi_\perp + \\ \left(\frac{Y_\parallel\nu_\perp(1-2\nu)}{(1-\nu_\parallel)(1-\nu)}\right)(\phi_\parallel - \phi_\perp) + \\ \left(\frac{Y_\parallel Y_\perp(1+\nu)(1-2\nu)^2}{Y\left(1-\nu_\parallel^2\right)(1-\nu)}\right)\phi_\parallel\end{array}\right]\end{array}\right), \tag{6.1}$$

where k_b is the Boltzmann constant, T is the temperature, f is the frequency, w_0 is the distance over which the electric field amplitude of the laser beam falls to $\frac{1}{e^2}$ of its maximum value,

$$\phi_\parallel = \frac{(Y_1\phi_1 d_1 + Y_2\phi_2 d_2)}{Y_\parallel(d_1+d_2)}\,, \tag{6.2}$$

$$\phi_\perp = \frac{Y_\perp(\phi_1 d_1/Y_1 + \phi_2 d_2/Y_2)}{Y_\parallel(d_1+d_2)}\,, \tag{6.3}$$

Y is the Youngs modulus of the substrate, ν is the Poisson ratio of the substrate,

$$Y_\parallel = \frac{Y_1 d_1 + Y_2 d_2}{d_1 + d_2}\,, \tag{6.4}$$

$$Y_\perp = \frac{d_1 + d_2}{d_1/Y_1 + d_2/Y_2}\,, \tag{6.5}$$

$$\nu_\parallel = \frac{\nu_1 + \nu_2}{2}\,, \tag{6.6}$$

$$\nu_\perp = \frac{\nu_1 Y_1 d_1 + \nu_2 Y_2 d_2}{d_1 Y_1 + d_2 Y_2}\,, \tag{6.7}$$

d represents the coating thickness, ϕ is the mechanical loss and the subscripts 1 and 2 refer to the alternating coating material properties. As described in Chapter 2, $\phi_\parallel$ and $\phi_\perp$ are physically interpreted as being the parallel and perpendicular mechanical losses of the coating.

Whilst it would have been ideal to calculate the thermal noise of a coating using Hongs' approach (outlined in Section 3.2), Harry's approach (Equation (6.1)) was used as bulk and shear losses for the coating materials later investigated were not known (Advanced LIGO silica coating and silicon nitride coating)

6.7.3 Reflectivity of mirror coatings

Equations (6.1) to (6.7) show that the thermal noise of a coating is dependent upon its thickness. The thickness of the coatings used in a gravitational wave detector are dictated by the detector's reflectivity requirements.

After the ET's beam splitter there exists a partially transmissive mirror known as the input test mass (ITM) mirror and a more highly reflective mirror, known as the end test mass (ETM) mirror (in each of the detector's arms). The ITM mirror has a higher transmission (and lower reflectivity) than the ETM mirror to enable light to be injected into the Fabry-Perot cavity formed by the ITM and ETM mirrors. To satisfy the ET-LF's technical specifications, the ITM mirror must have a reflectivity of 99.8 % whilst the ETM must have a reflectivity of 99.9995 % [189]. To meet these different mirror specifications, different thicknesses of a highly reflective coating are required to be deposited upon each test mass.

A standard highly reflective coating design uses alternating layers of two materials of different refractive index (where one of the materials has a high refractive index and the other a low refractive index). The thickness (t) of each layer is given by:

$$t = \frac{\lambda}{4n}, \tag{6.8}$$

where n is the refractive index of the material used in the layer and λ represents the wavelength at which high reflectivity is required. The reflectivity of a coating is dependent upon the number of pairs of alternating layers as well as the difference in the refractive indices of the layers. The reflectivity (R) of a coating comprised of an even number ($2N$) of these alternating layers is calculated using [190]:

$$R_{2N} = \left(\frac{n_s y - n_m}{n_s y + n_m}\right)^2, \tag{6.9}$$

where n_s is the refractive index of the test mass, n_m is the refractive index of the medium in front of the coating and $y = \left(\frac{n_H}{n_L}\right)^{2N}$, where n_H and n_L denote the refractive indices of the alternating high and low index layers respectively. If the coating is comprised of an odd number $(2N+1)$ of these alternating layers, then the coating's reflectivity is calculated using [190]:

$$R_{2N+1} = \left(\frac{n_H^2 y - n_m}{n_H^2 y + n_m}\right)^2 . \qquad (6.10)$$

From Equations (6.9) and (6.10) it can be observed that a coating can be made more reflective by increasing the number of layers or by increasing the difference in refractive index between the two coating materials. Since increasing the thickness of a coating results in an increase in thermal noise (Equation (6.1)), it is therefore more beneficial to maximise the difference in the two coating materials' refractive indices in order to increase coating reflectivity.

6.7.4 Thermal noise and absorption

Using the Advanced LIGO coating materials, silica and titania doped tantala, 18 bi-layers are required to achieve the ET-LF ETM reflectivity (calculated using Equation (6.9) and the material properties listed in Table (6.5)).

	Titania doped Tantala	Silica	Amorphous silicon	Silicon nitride
Refractive Index	2.05 [115]	1.44 [115]	3.48 [115]	2.17 [191]
Youngs modulus (GPa)	140 [192]	72 [193]	147 [134]	270 [191]
Poisson ratio	0.23 [189]	0.17 [193]	0.22 [194]	0.27 [195]
Loss	8.6×10^{-4} [196]	7.8×10^{-4} [121]	0.2×10^{-4} [113]	0.01×10^{-4}
Extinction Coefficient	N/A	N/A	3.5×10^{-5} [181]	4.3×10^{-6} [181]

Table 6-5: List of the tantala, silica, amorphous silicon and silicon nitride material properties used in the thermal noise and absorption calculations. Refractive indices are quoted at 1550 nm. The loss of the silica and tantala materials were measured at 20 K whilst the loss of the amorphous silicon and silicon nitride materials were measured at room temperature. The silicon nitride loss value was assumed from Figure 6-6. The extinction coefficient of amorphous silicon was determined at 20 K and 2 um. The extinction coefficient of silicon nitride was determined at 293 K and 2 um. The extinction coefficients of titania doped tantala and silica are unknown.

Using Equations (6.1) to (6.7) as well as the material properties listed in Table 6-5, the thermal noise of this coating deposited upon the silicon ETM was estimated to be 6.81×10^{-21} m/Hz$^{1/2}$ at 10 Hz (centre of peak sensitivity region) and 20 K (assuming a beam radius of 9 cm). In the case of the less reflective ITM coating, only 9 bilayers are required resulting in a coating thermal noise of 4.68×10^{-21} m/Hz$^{1/2}$. The total thermal noise (TTN) for the whole detector was then estimated to be 11.69×10^{-21} m/Hz$^{1/2}$, using:

$$\mathrm{TTN} = \sqrt{2 \times \mathrm{TNI}^2 + 2 \times \mathrm{TNE}^2} \ , \tag{6.11}$$

where TNI and TNE are the thermal noise contributions from the ITM and ETM respectively. This value of total thermal noise is a factor of 3.44 above the ET-LF requirement of 3.4×10^{-21} m/Hz$^{1/2}$. Whilst this coating does not meet the ET's thermal noise requirements, it does satisfy the ET's absorption requirement of 5 ppm [189] as this coating has a calculated absorption value of 0.5 ppm (if Advanced LIGO's coating absorption is extrapolated for the thicker coating layers required when using a 1550 nm laser and if identical extinction coefficients of the coating materials are assumed at 1550 nm and 1064 nm).

Now let's consider whether a coating made from alternating layers of amorphous silicon and silicon nitride can result in an improvement of thermal noise and absorption compared to the silica and titania doped tantala coated test mass. At 1550 nm, amorphous silicon has a very high absorption. In order to minimise this coating's absorption, it will be assumed that the ET-LF's wavelength is now configured to be 2 um (instead of 1550 nm). Assuming extinction coefficients of 3.5×10^{-5} for amorphous silicon and 4.3×10^{-6} for silicon nitride [181], a coating absorption of 31 ppm is calculated at 2 um (which is still much larger than the ET-LF's limit of 5 ppm). To calculate this coating's absorption (as well as the coatings which follow), each coating layer's electric field intensity and their extinction coefficient is required to be known. To estimate the electric field intensity in each layer and the coating's absorption, thin film modelling tools were used. The ET-LF's total thermal noise is calculated to be 7.61×10^{-22} m/Hz$^{1/2}$ which is a factor of 4.47 below the ET-LF's requirement (where 12 bilayers were required on the ETM resulting in a thermal noise of 4.45×10^{-22} m/Hz$^{1/2}$ and 6 bilayers on the ITM resulting in a thermal noise of 3.02×10^{-22} m/Hz$^{1/2}$). It should be noted that this value of total thermal noise is an approximation as a room temperature mechanical loss and extinction coefficient of silicon nitride was assumed in the calculations.

An amorphous silicon and silicon nitride coating is very promising from a thermal noise perspective. However, to allow such a coating to be used in the ET-LF detector, the optical absorption would need to be reduced by a factor of about 6. One way of solving the absorption issue is to use a "multi-material" coating [114, 115]. In these coating designs, a few layers of low-absorption silica and titania doped tantala are deposited on top of the coating stack to reflect the majority of laser power, reducing the power in the more highly absorbing amorphous silicon and silicon nitride layers below. This is a trade-off between absorption and mechanical loss, as the high loss of the silica and titania doped tantala layers will increase the thermal noise of the coating. To satisfy detector absorption requirements and ETM reflectivity requirements whilst also minimising the mirror's thermal noise, 12 amorphous silicon and silicon nitride bilayers are required to be deposited upon the test mass with a further 3 silica and titania doped tantala bilayers deposited on top. Figure 6-23 illustrates the structure of the "multi-material" and bilayer amorphous silicon/silicon nitride coatings deposited upon the ETM test masses.

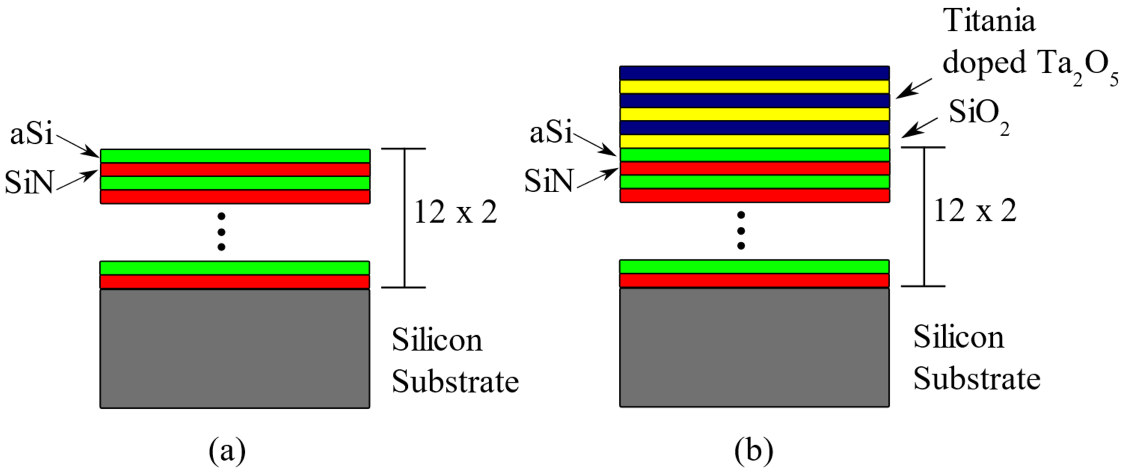

Figure 6-23: Structure of the purely bilayer (a) amorphous silicon and silicon nitride coating, and (b) the multi-material coating deposited upon an ETM silicon test mass.

The absorption of the multi-material coating is equal to 4 ppm and the thermal noise of the coated test mass equal to 2.8×10^{-21} m/Hz$^{1/2}$. To satisfy detector absorption requirements, ITM reflectivity requirements and minimise the mirror's thermal noise, 4 amorphous silicon and silicon nitride bilayers are required to be deposited upon the ITM with a further 3 silica and titania doped tantala bilayers deposited on top. The absorption of this coating is 4.1 ppm. The coated ITM's thermal noise is equal to 2.78×10^{-21} m/Hz$^{1/2}$. Using Equation (6.11) the total thermal noise of the detector is calculated to be 5.58×10^{-21} m/Hz$^{1/2}$, which is only 64 % above the ET's requirement. This represents a reduction of approximately 50 % compared to the coating made purely from titania doped tantala and silica. Table 6-6 summarises the total thermal noise and average absorption of the different coatings considered. Figure 6-24 shows the achieved strain sensitivity of the Advanced LIGO detector and the planned sensitivity of the ET-LF detector. Also shown on this plot are the strains associated with the total thermal noise of the ET-LF detector if the different coatings are used (where strain is equal to the total thermal noise divided by detector arm length, i.e. 10 km).

Coating	Total thermal noise (m/Hz$^{1/2}$)	Absorption (ppm)
Titania doped Tantala / silica	11.69×10^{-21}	0.5
Amorphous silicon / silicon nitride	7.61×10^{-22}	31
Multi-material	5.58×10^{-21}	4.05
ET requirement	3.4×10^{-21}	5

Table 6-6: Summary of the detector's thermal noise when different coatings are used at 10 Hz. Also included in the list are the absorptions of the coatings.

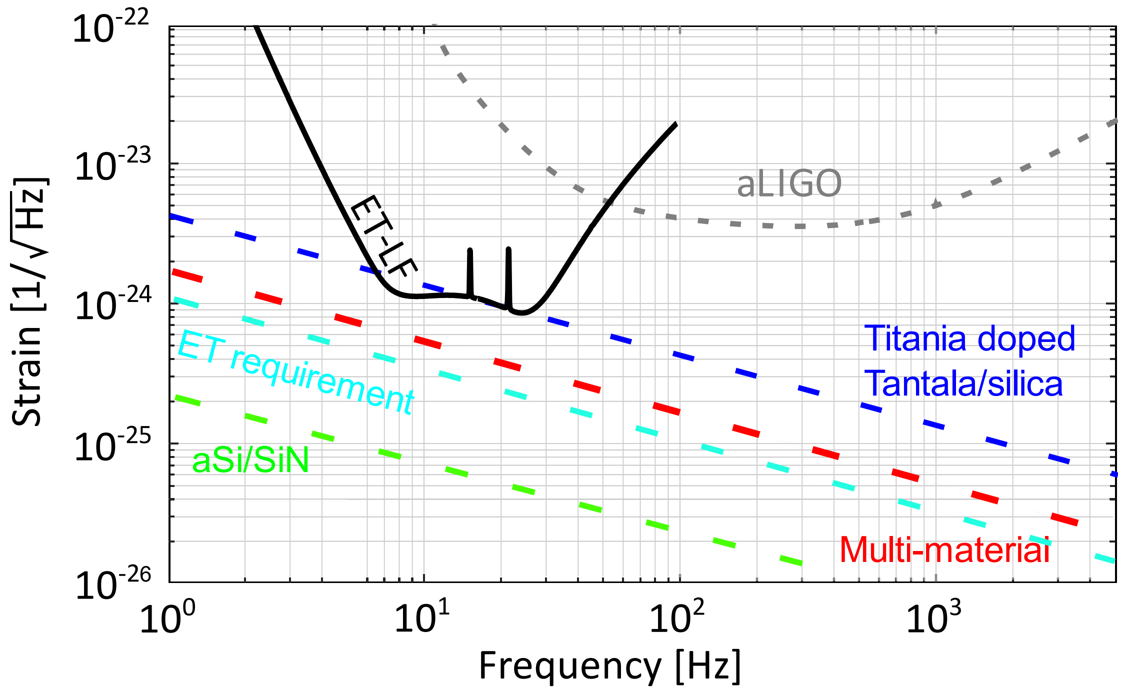

Figure 6-24: Strain sensitivity of the Advanced LIGO (aLIGO) and the ET-LF detectors. Also shown in the figure are the strains associated with the total thermal noise of the ET-LF detector if the different coatings listed in Table 6-6 are used.

To conclude, the multi-material coating satisfies the ET-LF absorption requirements and reduces the total thermal noise in the detector by almost 50 % compared to when the Advanced LIGO titania doped tantala coating is used. The detector's total thermal noise when the multi-material coating is used is only 64 % above the ET's requirement which could perhaps be reduced from further investigations into the effect of heat treatment upon the absorption of silicon nitride (a lower absorption of silicon nitride could result in the removal of a silica and titania doped tantala bilayer thus reducing the coating's thermal noise to perhaps detector requirements).

6.8 Conclusion

The use of silicon nitride membranes as substrates for coating mechanical loss measurements was investigated. One advantage of the membranes is that they are thin and therefore a relatively high proportion of the elastic energy associated with a resonant mode is stored in the coating. This ensures that the coating has a large effect on the loss of the coated substrate. Comparison to other substrates (silicon cantilevers, silica cantilevers and silica discs) showed that membranes are the second most sensitive substrate to the effects of clamping loss resulting in 'noise' in the coating loss measurement. However, membranes were also shown to have low mechanical loss and good repeatability, only surpassed by the silica discs. These results indicate that silica discs are the optimum substrate for mechanical loss measurements, but do not preclude the use of membranes - particularly in situations where high sensitivity to the coating loss is required or where large substrate areas cannot be coated. However, an experimental test using a coated membrane revealed another issue - a very large change in mode frequency and in membrane stress after coating deposition - which made analysis of the coating loss results very difficult. Significantly different bulk and shear losses were obtained from the membrane measurements and from measurements using a silica cantilever - possibly due to stress effects in the membrane case. Further work is required to fully understand these differences in bulk and shear loss before membranes can be reliably used for these experiments.

Based on measurements of the membrane loss and absorption, silicon nitride was considered as a potential material for use in highly reflective mirror coatings for cryogenically cooled gravitational wave detectors. A bilayer coating made from silicon nitride and amorphous silicon was shown to reduce the thermal noise of the ET-LF's detector by a factor of 15 compared to when a silica and titania doped tantala coating was used. However, a silicon nitride and amorphous silicon coating was shown to have the drawback of having an absorption (31 ppm) substantially above the ET-LF's requirement (5 ppm). Using a multi-material coating (made from silicon nitride, amorphous silicon, titania doped tantala and silica), the ET-LF's absorption requirements were achieved (4.05 ppm) and its thermal noise requirements almost obtained (64 % too large). Silicon nitride and amorphous silicon coatings are of great interest in the future as they could help in achieving the design sensitivity of cryogenically cooled gravitational wave detectors.

7 Conclusion

Gravitational waves were predicted by Einstein's General Theory of Relativity. The first directly detected gravitational wave was observed on the 14th of September 2015. Since that day five more gravitational waves have been observed providing new astronomical information. These detections were made using ground-based interferometric detectors which observe the effects of the waves by monitoring the differential changes in their arm lengths. Since gravitational waves are expected to change the detector's arm lengths by approximately 1×10^{-19} m, it is essential for all sources of noise to be exceedingly low. Coating thermal noise is one major limit to the sensitivity of current and future gravitational wave detectors at their most sensitive frequencies.

The research presented in this thesis focused upon improving the process currently used to determine the magnitude of a detector's coating thermal noise as well as identifying coating materials which can reduce thermal noise.

Recent literature indicates that an accurate determination of a coating's thermal noise requires knowledge of the bulk and shear mechanical loss angles of the coating. Two new methods to determine the bulk and shear losses of a coating were developed and applied to a range of coatings (amorphous silicon, tantala, silicon nitride). It was found that the bulk and shear losses of a coating could be significantly different from one another. This is believed to be the first published analysis which shows that the bulk and shear losses of a coating can indeed be different from each other. Furthermore, the results also showed that it is possible for the loss of a coating to be completely dominated by shear loss. Since coating thermal noise is less sensitive to shear loss than to bulk loss, this raises the interesting possibility of reducing coating thermal noise by using materials which are dominated by shear loss. The robustness of the methods used to determine the bulk and shear losses of a coating were tested by investigating their sensitivity to the stress and curvature of a coated substrate. It was observed that stressed and curved substrates have significantly different energy ratios for uncurved and unstressed substrates, and in order to determine the most accurate bulk and shear losses of a coating, this effect should be incorporated within the calculation process.

To determine the thermal noise of a coating in a gravitational wave detector, the mechanical loss of the coating must first be known. Recent research and observations have called into question the accuracy of coating losses when the coated substrate has exhibited significant deformation due to stress. This has particularly been observed in thin silicon cantilever substrates. It is speculated that the inaccuracy of these losses could be due to the effect of the coating causing the coated cantilever to curve and stress which in turn changes the silicon cantilever's thermoelastic loss. The results from FEA indicated that the thermoelastic loss of a silicon cantilever (which has been coated) is unaffected by stress but is affected by curvature. Experimental evidence appears to confirm these conclusions but they further imply that either the FEA does not fully account for the effect of curvature upon thermoelastic loss or an unmodelled effect is also affecting a silicon cantilever's thermoelastic loss. It is clear care must be taken when using thin silicon cantilevers to measure coating loss, particularly in the temperature range of 150 K – 300 K, where the loss of these samples is dominated by thermoelastic loss which can be affected by curvature. FEA was also used to investigate the effect of stress and curvature upon a silicon disc's thermoelastic loss. It was observed that the magnitude of stress and curvature caused by depositing a coating onto a silicon disc in the lab had no effect upon the substrate's thermoelastic loss.

One material which could potentially be used in future cryogenic gravitational wave detector coatings is silicon nitride as it can have a low mechanical loss. However, silicon nitride has quite a high absorption. Since the mechanical loss of silicon nitride has been shown to decrease as its stress increases, the effect of stress upon its absorption was investigated. Due to experimental limitations, the direct effect of stress upon absorption could not be quantified, instead the effect of stress upon the product of the membrane's absorption and its thermo-optic coefficient was determined. This product was found to be independent of stress. Whilst this result implies that absorption is likely to be independent of stress, this should be explicitly checked. It was also determined from this experiment that the thermal conductivity and thermal expansion of the low stress silicon nitride membrane used was (23 ± 3) W/mK and $(1.4 \pm 0.2) \times 10^{-6}\ 1/\mathrm{K}$ respectively. Accurate knowledge of these parameters is essential for estimating the magnitude of a mirror coating's thermoelastic noise.

The use of silicon nitride membranes as substrates for coating mechanical loss measurements was investigated. One advantage of the membranes is that they are thin and therefore a relatively high proportion of the elastic energy associated with a resonant mode is stored in the coating, giving high sensitivity to the loss of the coating. Compared to other substrates (silicon cantilevers, silica cantilevers and silica discs) membranes were shown to be the second most sensitive substrate to the effects of clamping loss resulting in 'noise' in the coating loss measurement. However, membranes were also shown to have low mechanical loss and good repeatability, only surpassed by silica discs. Whilst none of these results excludes silicon nitride membranes being used in mechanical loss experiments, they indicate a preference should be shown in selecting silica discs before membranes. However, it was also later observed that the bulk and shear losses of the same coating deposited upon a membrane and cantilever were significantly different from one another. Further work is required to fully understand these differences in coating loss before membranes can be reliably used in ring down experiments.

Future gravitational wave detectors, such as the ET, are expected to be limited in sensitivity by the thermal noise of the highly reflective mirror coatings. A mirror coating made from amorphous silicon/ silicon nitride was investigated to identify whether it could enable the ET to achieve its absorption and thermal noise requirements. From the literature, it is already known that in order to minimise the absorption of a silicon nitride/ amorphous silicon coating, 450 °C heat-treated amorphous silicon should be used. Since heat treating amorphous silicon changes its absorption, the effect of heat-treatment upon silicon nitride absorption was investigated. No degradation or improvement in performance was observed for annealing temperatures up to 900 °C. The thermal noise and absorption of the ET was calculated when an amorphous silicon/ silicon nitride coating was assumed to be used. These values of absorption and thermal noise were compared to those determined for the ET when the Advanced LIGO (silica/ titania doped tantala) coating was assumed to be used. The amorphous silicon/silicon nitride coating was shown to have a thermal noise that met ET design requirements and showed a factor of 15 reduction in thermal noise compared to the current Advanced LIGO coating. This is a huge reduction in thermal noise and shows that the thermal noise of the ET detector can be achieved. However, it was observed that the amorphous silicon/ silicon nitride coating had a higher absorption than the Advanced LIGO coating and that it exceeded the ET design requirements. To try to take advantage of the low thermal noise of the silicon nitride/amorphous silicon coating and the low absorption of the silica/titania doped tantala coating, a “multi-material” coating was proposed. This “multi-material” coating satisfied the ETs absorption requirement but not its thermal noise. However, unlike the other coatings investigated, this “multi-material” coating only narrowly exceeded design requirements and with a little further improvement (perhaps by heat-treating at higher temperatures) this coating could be used in the ET.

The development of coatings with lower thermal noise remains one of the most important areas of research to allow the realisation of gravitational wave detectors with higher sensitivity. The work in this thesis shows several promising avenues which should be further pursued: particularly identifying materials with low shear mechanical loss and the use of amorphous silicon/ silicon nitride highly reflective coatings. Further appropriate analysis of mechanical loss measurements to determine shear and bulk loss as well as accounting for thermoelastic effects upon coated silicon samples will be essential for more accurately predicting the thermal noise of new coatings.

Bibliography

[1] A. Einstein, “Die Grundlage der allgemeinen relativitatstheorie,” *Annalen der Physik,* vol. 69, pp. 769 - 822, 1916.

[2] B. Schutz, “Determining the Hubble constant from gravitational wave observations,” *Nature,* vol. 323, pp. 310 - 311, 1986.

[3] B. Abbott, R. Abbott, T. Abbott, M. Abernathy et al, “The rate of binary black hole mergers inferred from Advanced LIGO observations surrounding GW150914,” *The Astrophysical Journal,* vol. 833, no. 1, 2016.

[4] B. Abbott, R. Abbott, T. Abbott, M. Abernathy et al, “The basic physics of the binary black hole merger GW150914,” *Annalen der Physik,* vol. 529, no. 1, 2017.

[5] R. Taylor and J. Hulse, “Discovery of a pulsar in a binary system,” *The Astrophysical Journal,* vol. 195, pp. 51 - 53, 1975.

[6] B. Abbott, R. Abbott, T. Abbott, M. Abernathy et al, “Observation of gravitational waves from a binary black hole merger,” *Physical Review Letters,* vol. 116, no. 061102, 2016.

[7] B. Abbott, R. Abbott, T. Abbott, M. Abernathy et al, “GW170817: Observation of gravitational waves from binary neutron star inspiral,” *Physical Review Letters,* vol. 119, no. 161101, 2017.

[8] B. Abbott, R. Abbott, T. Abbott, M. Abernathy et al, “GW170814: A three- detector observation of gravitational waves from a binary black hole coalescence,” *Physical Review Letters,* vol. 119, no. 141101, 2017.

[9] J. Weber, “Detection and generation of gravitational waves,” *Physical Review,* vol. 117, no. 1, 1960.

[10] J. Aasi, B. Abbott, R. Abbott, T. Abbott et al, “Advanced LIGO,” *Classical and Quantum Gravity,* vol. 32, no. 7, 2015.

[11] F. Acernese, T. Adams, M. Agathos, K. Agatsuma et al, “The Advanced Virgo detector,” *Journal of Physics: Conference Series,* vol. 610, no. 1, 2015.

[12] M. Abernathy, F. Acernese, S. Aoudia, P. Amaro-Seoane et al, “Einstein gravitational wave telescope (ET) conceptual design study,” 2011. [Online]. Available: https://tds.ego-gw.it/ql/?c=7954. [Accessed 01 03 2018].

[13] A. LIGO, “The next step in gravitational wave astronomy,” Advanced LIGO, 24 October 2016. [Online]. Available: https://www.advancedligo.mit.edu/. [Accessed 02 03 2018].

[14] B. Schutz, “Gravitational wave sources,” *Classical and Quantum Gravity,* vol. 13, pp. 219 - 238, 1996.

[15] B. Abbott, R. Abbott, T. Abbott, F. Acernese et al, "A gravitational wave standard siren measurement of the Hubble constant," *Nature,* vol. 551, no. 7678, pp. 85 - 88, 2017.

[16] R. Kumar, "Aspects of suspension design for the development of advanced gravitational wave detectors," P.h.D thesis, University of Glasgow, Glasgow, 2013.

[17] NASA, "Supernovae," NASA, 1 January 2011. [Online]. Available: https://imagine.gsfc.nasa.gov/science/objects/supernovae2.html. [Accessed 02 03 2018].

[18] S. Chandrasekhar, "Solutions of two problems in the theory of gravitational radiation," *Physical Review Letters,* vol. 24, pp. 611 - 615, 1970.

[19] C. A. Observatory-CAHA, "Asymmetric supernovae: Not all stellar explosions expand spherically," Science Daily, 27 02 2011. [Online]. Available: https://www.sciencedaily.com/releases/2011/02/110224145803.htm. [Accessed 19 06 2018].

[20] B. Sathyaprakash and B. Schutz, "Astrophysics and cosmology with gravitational waves," *Living Reviews in Relativity,* vol. 12, no. 2, 2009.

[21] B. Abbot, R. Abbott, T. Abbott, M. Abernathy et al, "First search for gravitational waves from known pulsars with Advanced LIGO," *The Astrophysical Journal,* vol. 839, no. 12, 2017.

[22] Executive Comittee, "The LSC - Virgo white paper on gravitational wave searches and astrophysics executive summary 2014 - 2015 edition," in *Technical Report*, 2014, p. 16.

[23] B. Abbot, R. Abbot, R. Adhikari, J. Agresti et al, "Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: results from the second LIGO science run," *Physical Review D,* vol. 76, p. 082001, 2007.

[24] Cardiff University, "Wobbling neutron stars," Cardiff University, [Online]. Available: http://blogs.cardiff.ac.uk/physicsoutreach/gravitational-physics-tutorial/wobbling-neutron-stars/. [Accessed 19 06 2018].

[25] P. Patel, Search for gravitational waves from nearby neutron star using barycentric resampling, Pasadena: California Institute of Technology; PhD thesis, 2010.

[26] B. Owen, L. Lindblom, C. Cutler, B. Schutz et al, "Gravitational waves from hot young rapidly rotating neutron stars," *Physical Review D,* vol. 58, no. 084020, 1998.

[27] B. Abbott, R. Abbott, R. Adhikari, J. Agresti et al, "Searching for a stochastic background of gravitational waves with the laser interferometer gravitational wave observatory," *The Astrophysical Journal,* vol. 659, no. 2, 2007.

[28] X. Zhu, E. Howell, T. Regimbau, D. Blair et al, "Stochastic gravitational wave background from coalescing binary black holes," *The Astrophyscial Journal,* vol. 739, no. 2, 2011.

[29] X. Zhu, E. Howell and D. Blair, "Observational upper limits on the gravitational wave production of core collapse supernovae," *Monthly Notices of the Royal Astronomical Society,* vol. 409, pp. L132 - L136, 2010.

[30] P. Lasky, M. Bennett and A. Melatos, “Stochastic gravitational wave background from hydronamic turbulence in differentially rotating neutron stars,” *Physical Review D,* vol. 87, no. 063004, 2013.

[31] B. Allen and J. Romano, “Detecting a stochastic background of gravitational radiation: signal processing strategies and sensitivities,” *Physical Review D,* vol. 59, no. 102001, 1999.

[32] J. Firmage, Revolutions in physics: exploring the evolution and state of modern physics and the possibilities that a new paradigm holds for human civilisation, USA: Xlibris corporation, 2009.

[33] B. Abbott, R. Abbott, T. Abbott, M Abernathy et al, “GW151226: Observation of gravitational waves from a 22 solar mass binary black hole coalescence,” *Physics Review Letters,* vol. 116, p. 241103, 2016.

[34] D. Castelvecchi, “LIGO detects whispers of another black hole merger,” *Nature,* vol. 534, no. 7608, 2016.

[35] B. Abbott, R. Abbott, T. Abbott, F. Acernese et al, “GW170104: Observation of a 50 solar mass binary black hole coalescence,” *Physical Review Letters,* vol. 118, no. 22, p. 221101, 2017.

[36] B. Abbott, R. Abbott, T. Abbott, F. Acernese et al, “GW170608: Observation of a 19 solar mass binary black hole coalescence,” *The Astrophysical Journal Letters,* vol. 851, no. 2, 2017.

[37] L. V. S. collaborations, “Update on gravitational wave science from the LIGO-VIRGO scientific collaborations,” G7 ministerial meeting on science, 27 09 2017. [Online]. Available: https://www.youtube.com/watch?v=xR6d8V5oh0o. [Accessed 2 03 2018].

[38] J. Weber, “Evidence for discovery of gravitational radiation,” *Physical Review Letters,* vol. 22, pp. 1320 - 1324, 1969.

[39] J. Weber, “Anisotropy and polarisation in the gravitational radiation experiments,” *Physical Review Letters,* vol. 25, pp. 180 - 184, 1970.

[40] J. Levine and R. Garwin, “Absence of gravity wave signals in a bar at 1695Hz,” *Physical Review Letters,* vol. 31, pp. 173 - 176, 1973.

[41] J. Tyson, “Null search for bursts of gravitational radiation,” *Physical Review Letters,* vol. 31, pp. 326 - 329, 1973.

[42] V. Braginsky, A. Manukin, E. Popov, V. Rudenko et al, “Upper limit of the density of extraterrestrial gravitational radiation,” *Soviet Physics,* vol. 39, p. 387, 1974.

[43] A. Vinante, R. Mezzena, G. Prodi and S. Vitale, “Thermal noise in a high Q ultracryogenic resonator,” *Review of Scientific Instruments,* vol. 76, p. 074501, 2005.

[44] C. Costa and O. Aguiar, “Spherical gravitational wave detectors: MiniGRAIL and Mario Schenberg,” *Journal of Physics: Conference Series,* vol. 484, no. 1, 2012.

[45] A. Saha and S. Gangopadhyay, “Resonant detectors of gravitational wave as a possible probe of the noncommutative structure of space,” *Classical and Quantum Gravity,* vol. 33, p. 205006, 2016.

[46] M. Gertsenshtein and V. Pustovoit , “On the detection of low frequency gravitational waves,” *Soviet physics,* vol. 16, no. 2, pp. 433 - 435, 1962.

[47] G. Moss, L. Miller and R. Forward, “Photon noise limited laser transducer for gravitational antenna,” *Applied Optics,* vol. 10, no. 11, pp. 2495 - 2498, 1971.

[48] R. Forward, “Wideband laser interferometer gravitational radiation experiment,” *Physical Review D,* vol. 17, pp. 379 - 390, 1978.

[49] D. Shoemaker, R. Schilling, L. Schnupp, W. Winkler et al, “Noise behavior of the Garching 30m prototype gravitational wave detector,” *Physical Review D,* vol. 38, pp. 423 - 432, 1988.

[50] A. Abramovici, W. Althouse, J. Camp, D.Durance et al, “Improved sensitivity in a gravitational wave interferometer and implications for LIGO,” *Physics Letters A,* vol. 218, pp. 157 - 163, 1996.

[51] R. Drever, J. Pugh, W. Edelstein and H. Ward, “Proceedings of the Royal Society of London,” *Proceedings of the Royal Society of London,* vol. 38, no. 2, pp. 11 - 18, 1979.

[52] H. Ward, J. Hough, G. Newton, B. Meers et al, “Laser interferometric sensing techniques for very small displacements with applications to gravitational wave detectors,” *IEEE Transactions on Instrumentation and Measurement,* vol. 34, no. 2, pp. 261 - 265, 1985.

[53] M. Beccaria, M. Bernardini, S. Braccini, C. Bradaschia et al, “Relevance of Newtonian seismic noise for the VIRGO interferometer sensitivity,” *Classical and Quantum Gravity,* vol. 15, pp. 3339-3362, 1998.

[54] R. Robie, Investigations of the mechanical properties of thin film mirror coatings for use in future gravitational wave detectors, PhD thesis, Glasgow: University of Glasgow, 2018.

[55] M. Plissi, C. Torrie, M. Husman, N. Robertson et al, “GEO600 triple lock pendulum suspension system: seismic isolation and control,” *Review of Scientific Instruments,* vol. 71, pp. 2539-2545, 2000.

[56] S. Braccini, C. Bradaschia, R. Fabbro, A. Virgilio et al, “Seismic vibrations mechanical filters for the gravitational wave detector VIRGO,” *Review of Scientific Instruments,* vol. 67, pp. 2899-2902, 1996.

[57] F. Matichard, B. Lantz, R. Mittleman, K. Mason et al, “Seismic isolation of Advanced LIGO review of strategy instrumentation and performance,” *Classical and Quantum Gravity,* vol. 32, no. 18, p. 185003, 2015.

[58] P. Saulson, “Terrestrial gravitational noise on a gravitational wave antenna,” *Physical Review D,* vol. 30, no. 4, pp. 732 - 736.

[59] S. Hughes and K. Thorne, “Seismic gravity-gradient noise in interferometric gravitational wave detectors,” *Physical Review D,* vol. 58, p. 122002, 1998.

[60] M. Beker, G. Cella, R. De Salvo, M. Doets et al, “Improving the sensitivity of future GW observatories in the 1-10Hz band: newtonian and seismic noise,” *General Relativity and Gravitation,* vol. 43, pp. 623-656, 2011.

[61] M. Bader, A. Bertolini, J. Brand and J. Heijningen, “Newtonian noise: survey of low seismic noise environment and sub ng instrustmentation,” in *4th ELITES General meeting*, Tokyo, 2015.

[62] G. Cagnoli, L. Gammaitoni, J. Hough, J. Kovalik et al, “Very high Q measurements on a fused silica monolithic pendulum for use in enhanced gravity wave detectors,” *Physical Review Letters,* vol. 85, no. 12, p. 2442, 2000.

[63] A. Buonanno and Y. Chen, “Quantum noise in second generation, signal - recycled laser interferometric gravitational wave detectors,” *Physical Review D,* vol. 64, p. 042006, 2001.

[64] C. Caves, “Quantum mechanical noise in an interferometer,” *Physical Review D,* vol. 23, pp. 1693-1708, 1981.

[65] R. Weiss, “Electromagnetically coupled broadband gravitational antenna,” Quarterly Progress Report 105, MIT, 1972.

[66] H. Luck, M. Hewitson, P. Ajith, B. Allen et al, “Status of the GEO600 detector,” *Classical and Quantum Gravity,* vol. 23, no. 8, pp. 71-78, 2006.

[67] A. Perot and C. Fabry, “On the application of interference phenomena to the solution of various problems of spectrocopy and metrology,” *The Astrophysical Journal,* vol. 9, p. 87, 1899.

[68] G. Harry, “Advanced LIGO: the next generation of gravitational wave detectors,” *Classical and Quantum Gravity,* vol. 27, p. 084006, 2010.

[69] L. laboratory, “LIGO's interferometer,” LIGO laboratory, [Online]. Available: https://www.ligo.caltech.edu/page/ligos-ifo. [Accessed 10 03 2018].

[70] B. Meers, “The frequency response of interferometric gravitational wave detectors,” *Physics Letters A,* vol. 142, pp. 465 - 470, 1989.

[71] J. Veitch, I. Mandel, B. Aylott, B.Farr et al, “Estimating parameters of coalescing compact binaries with proposed advanced detector networks,” *Physical Review D,* vol. 85, p. 104045, 2012.

[72] H. Luck, M. Hewitson, P. Ajith, B. Allen et al, “The GEO600 project,” *Classical and Quantum Gravity,* vol. 14, pp. 1471-1476, 1997.

[73] H. Luck, M. Hewitson, P. Ajith, B. Allen et al, “The upgrade of GEO600,” *Journal of Physics: Conference Series,* vol. 228, no. 1, p. 012012, 2010.

[74] K. Dooley, “Status of GEO600,” *Journal of physics: Conference Series,* vol. 610, p. 012015, 2015.

[75] LIGO, “LIGO Hanford,” LIGO, 02 05 2005. [Online]. Available: https://www.ligo.caltech.edu/image/ligo20150731f. [Accessed 12 03 2018].

[76] LIGO, “LIGO Livingston,” LIGO, 19 05 2015. [Online]. Available: https://www.ligo.caltech.edu/image/ligo20150731c. [Accessed 12 03 2018].

[77] B. Barish, “The Laser Interferometer Gravitational Wave Observatory LIGO,” *Advances in Space Research,* vol. 25, no. 6, pp. 1165 - 1169, 2000.

[78] J. Smith and F.T.L.S collaboration, “The path to the enhanced and advanced gravitational wave detectors,” *Classical and Quantum Gravity,* vol. 26, no. 11, p. 114013, 2009.

[79] A. Brooks, “The status of Enhanced LIGO,” in *AIP congress*, Adelaide, 2008.

[80] F. Acernese, P. Amico, M. Alshourbagy, F. Antonucci et al, “The VIRGO 3km interferometer for gravitational wave detection,” *Journal of Optics: Pure and Applied Optics,* vol. 10, p. 064009, 2008.

[81] R.Takahashi and the TAMA collaboration, “Status of TAMA300,” *Classical and Quantum Gravity,* vol. 21, pp. 403 - 408, 2004.

[82] D. Tatsumi, R. Takahashi, K. Arai, N. Nakagawa et al, “Current staturs of Japanese detectors,” *Classical and Quantum Gravity,* vol. 24, pp. 399 - 403, 2007.

[83] K. Somiya and the KAGRA collaboration, “Detector configuration of KAGRA - the Japanese cryogenic gravitational wave detector,” *Classical and Quantum Gravity,* vol. 29, no. 12, p. 124007, 2012.

[84] LIGO scientific collaboration, “Instrument science white paper - LIGO-T15TBI–v1,” LIGO scientific collaboration, 2008.

[85] M. Zucker, “Getting an A+: Enhancing Advanced LIGO,” in *LIGO - DAWN Workshop 2*, 2016.

[86] D. McGuigan, C. Lam, R. Gram, A. Hoffman et al, “Measurements of the mechanical Q of single-crystal silicon at low temperatures,” *Journal of Low Temperature Physics,* vol. 30, no. 5-6, pp. 621-629, 1978.

[87] L. Anderson and H. Bommel, “Ultrasonic absorption in fused silica at low temperatures and high frequencies,” *Journal of the American Ceramic Society,* vol. 30, no. 115008, pp. 125-131, 1955.

[88] Indigo, “LIGO- India,” Indigo, [Online]. Available: http://www.gw-indigo.org/tiki-index.php?page=LIGO-India. [Accessed 11 03 2018].

[89] F. Acernese, S. Aoudia, P. Amaro-Seoane, F. Barone et al, “Einstein gravitational wave telescope conceptual design study: ET-0106A-10,” *European Commission,* no. 3, p. 211743, 2011.

[90] P. Corneille, “Einstein Telescoop: gravitatiedetector van de toekomst,” Spacepage, [Online]. Available: https://www.spacepage.be/artikelen/sterrenkunde/observatoria-wereldwijd/einstein-telescoop-gravitatiedetector-van-de-toekomst. [Accessed 12 03 2018].

[91] K. Danzmann, P. Seane, H. Audley, S. Babak et al, “Laser Interferometer Space Antenna,” LISA, 2017.

[92] R. Brown, “A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen plants; and on the general existence of active molecules in organic and inorganic bodies,” *Philosophical Magazine,* vol. 4, pp. 161 - 173, 1828.

[93] A. Einstein, “On the movement of small particles suspended in a stationary liquid demanded by the molecular-kinetic theory of heat,” *Annalen der Physik,* vol. 17, pp. 549-560, 1900.

[94] H. Callen and R. Greene, “On a theorem of irreversible themodynamics,” *Physical Review,* vol. 86, pp. 702-710, 1952.

[95] H. Callen and T. Welton, “Irreversibility and generalised noise,” *Physical Review,* vol. 83, pp. 34-40, 1951.

[96] R. Greene and H. Callen, “On the formalism of thermodynamic fluctuation theory,” *Physical Review,* vol. 83, pp. 1231-1235, 1951.

[97] A. Nowick and B. Berry, Anelastic relaxation in crystalline solids, New York: Academic Press, 1972.

[98] S. Rowan, J. Hough and D. Crooks, “Thermal noise and material issues for gravitational wave detectors,” *Physics Letters A,* vol. 347, no. 1, pp. 25 - 32, 2005.

[99] T. Hong, Brownian thermal noise in interferometric gravitational wave detectors and single photon optomechanics, PhD thesis: California Institute of Technology, 2013.

[100] P. Saulson, “Thermal noise in mechancial experiments,” *Physical Review D,* vol. 42, pp. 2437-2445, 1990.

[101] A. Gillespie and F. Raab, “Thermally excited vibrations of the mirrors of laser interferometer gravitational wave detectors,” *Physical Review D,* vol. 52, no. 2, pp. 577-585, 1995.

[102] Y. Levin, “Internal thermal noise in the LIGO test masses: A direct approach,” *Physical Review D,* vol. 57, no. 2, pp. 659-663, 1998.

[103] K. Yamamoto, Study of the thermal noise caused by inhomogeneously distributed loss, PhD thesis: University of Tokyo, 2000.

[104] F. Bondu, P. Hello and J. Vinet, “Thermal noise in mirrors of interferometric gravitational wave antennas,” *Physical Letters A,* vol. 246, no. 3, pp. 227-236, 1998.

[105] N. Nakagawa, A. Gretarsson, E. Gustafson and M. Fejer, "Thermal noise in half finite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror," *Physical Review D,* vol. 65, p. 102001, 2002.

[106] G. Harry, A. Gretarsson, R. Saulson, S. Kittelberger et al, "Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings," *Classical and Quantum Gravity,* vol. 19, pp. 897-917, 2002.

[107] T. Hong, H. Yang, E. Gustafson, R. Adhikari et al, "Brownian Thermal Noise in Multilayer Coated Mirrors," *Physical Review D,* vol. 87, no. 8, p. 082001, 2012.

[108] V. Braginsky and S. Vyatchanin, "Thermodynamical fluctuations in optical mirror coatings," *Physics Letters A,* vol. 312, no. 3, pp. 244-255, 2003.

[109] M. Fejer, S. Rowan, G. Cagnoli, D. Crooks et al, "Thermoelastic dissipation in inhomogeneous media: loss measurements and displacement noise in coated test masses for interferometric gravitational wave detectors," *Physical Review D,* vol. 70, p. 082003, 2004.

[110] V. Braginsky, M. Gorodetsky and S. Vyatchanin, "Thermo-refractive noise in gravitational wave antennae," *Physics Letters A,* vol. 271, no. 5, pp. 303-307, 2000.

[111] M. Evans, S. Ballmer, M. Fejer, P. Fritchel et al, "Thermo-optic noise in coated mirrors for high-precision optical measurements," *Physical Review D,* vol. 78, p. 102003, 2008.

[112] J. Steinlechner, C. Kruger, IW. Martin, A. Bell et al, "Optical absorption of silicon nitride membranes at 1064nm and 1550nm," *Physical Review D,* vol. 96, p. 022007, 2017.

[113] P. Murray, I.W. Martin, K.Craig, J. Hough et al, "Ion beam sputtered amorphous silicon films for cryogenic precision measurement systems," *Physical Review D,* vol. 92, p. 062001, 2015.

[114] W. Yam , S. Gras and M. Evans, "Multi-material coatings with reduced thermal noise," *Physical Review D,* vol. 91, p. 042002, 2015.

[115] J. Steinlechner, I. Martin, J. Hough, C. Krüger et al, "Thermal noise reduction and absorption optimization via multimaterial coatings," *Physical review D,* vol. 91, p. 042001, 2015.

[116] O. Anderson and H. Bommel, "Ultrasonic absorption in fused silica at low temperatures and high frequencies," *Journal of the American Ceramic Society,* vol. 38, no. 4, pp. 125-131, 1955.

[117] B. Berry and W. Pritchet, "Vibrating reed internal friction apparatus for films and foils," *IBM Journal of Research and Development,* vol. 19, no. 4, pp. 334 - 343, 1975.

[118] S. Penn, P. Sneddon, H. Armandula, J. Betzwieser et al, "Mechanical loss in tantala/silica dieletric mirror coatings," *Classical and Quantum Gravity,* vol. 20, pp. 2917-2928, 2003.

[119] T. Metcalf, X. Liu, B. Houston, J. Baldwin et al, "Low temperature internal friction in nanocrystalline diamond films," *Applied Physics Letters,* vol. 86, no. 8, p. 081910, 2005.

[120] J. Ferry, Viscoelastic Properties of Polymers, Toronto: John Wiley and Sons, 1980, pp. 48-49.

[121] I. Martin, R. Nawrodt, K. Craig, C. Schwarz et al , "Low temperature mechanical dissipation of an ion-beam sputtered silica film," *Classical and Quantum Gravity,* vol. 31, no. 3, p. 035019, 2014.

[122] I. Martin, H. Armandula, C. Comtet, M. Fejer et al, "Measurements of a low temperature mechanical dissipation peak in a single layer of Tantala doped with Titania," *Classical Quantum Gravity,* vol. 25, no. 8, p. 055005, 2008.

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

[124] T. Quinn, C. Speake, R. Davis and W. Tew, "Stress dependent damping in Cu_Be torsion and flexure suspensions at stressed up to 1.1GPa," *Physics Letters A,* vol. 197, pp. 197 - 208, 1995.

[125] N. Mclachlan, Theory of vibrations, New York: Dover, 1952.

[126] D. McLachlan and L. Chamberlain, "Atomic vibrations and melting point in metals," *Acta Metallurgica,* vol. 12, pp. 571 - 576, 1964.

[127] C. Van Eysden and J. Sader, "Resonant frequencies of a rectangular cantilever beam immersed in a fluid," *Journal of Applied Physics,* vol. 100, p. 114916, 2006.

[128] D. Crooks, Mechanical loss and its significance in the test mass mirrors of gravitational wave detectors, PhD thesis, Glasgow: University of Glasgow, 2003.

[129] B. White and R. Pohl, "Internal friction of subnanometer amorphous silica films," *Physical Review Letters,* vol. 75, no. 24, pp. 4437- 4439, 1995.

[130] ANSYS, ANSYS Workbench User Guide, Canonsburg: ANSYS, 2009.

[131] COMSOL, COMSOL Structural Mechanics Module, Los Angeles: COMSOL, 2017.

[132] E. Ghanbari, I. Trigor and T. Nguyen, "A broad-beam electron cyclotron resonance ion source for sputtering etching and deposition of material," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films,* vol. 7, p. 918, 1989.

[133] D. Vine, Private Communication, 2015.

[134] M. Abernathy, Mechanical properties of coating materials for use in the mirrors of interferometric gravitational wave detectors, PhD thesis, Glasgow: Glasgow University, 2012.

[135] R. Robie, Private Communication, 2016.

[136] R. Liftshitz and M. Roukes, "Thermoelastic damping in micro and nanomechanical systems," *Physical Review B,* vol. 61, no. 8, pp. 5600 - 5609, 2000.

[137] M. Abernathy, G. Harry, J. Newport, H. Fair et al, “Bulk and shear mechancial loss of titania-doped tantala,” *Physics Letters A,* vol. In press, 2017.

[138] S. Gras, H. Yu, W. Yam, D. Martynov et al, “Audio-band coating thermal noise measurement for advanced ligo with a multimode optical resonator,” *Physical Review D,* vol. 95, no. 2, p. 022001, 2017.

[139] Tafelmaier Daunschicht-Technik GmbH, “Rosenheim,” http://www.tafelmaier.de Germany. [Online].

[140] G. Schlichtherle, G. Strauss, H. Tafelmaier, D. Huber et al, “Reactive low volatage ion plating (RLVIP),” *Vakuum in Forschung und Praxis,* vol. 17, no. 4, pp. 210-217, 2005.

[141] A. Cumming, K. Craig, I. Martin, R. Bassiri et al , “Measurement of the mechanical loss of prototype GaP/AlGaP crystalline coatings for future gravitational wave detectors,” *Classical and Quantum Gravity,* vol. 32, no. 3, p. 035002, 2015.

[142] E. Hirose, K. Craig, H. Ischitsuka, I. Martin et al, “Mechanical loss of a multilayer tantala/silica coating on a sapphire disk at cryogenic temperatures: toward the KAGRA gravitational wave detector,” *Physical Review D,* vol. 90, no. 10, p. 102004, 2014.

[143] SIOS, “Laser interferometer,” SIOS, [Online]. Available: http://www.sios-de.com/products/length-measurement/laser-interferometer/. [Accessed 21 12 2017].

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

[145] D. Crooks, G. Cagnoli, M. Fejer, G. Harry et al , “Experimental measurements of mechanical dissipation associated with dieletric coatings formed using silica, tantala and aluminium oxide,” *Classical and Quantum Gravity,* vol. 23, no. 15, pp. 4953-4965, 2006.

[146] T. Hsu, MEMS and microsystems: design, manufacture and nanoscale engineering, New Jersey: John Wiley and Sons, 2008, pp. 172 - 173.

[147] COMSOL, “Private Communication,” 2017.

[148] M. Fejer, *Private Communication,* 2018.

[149] S.Kumar and M.A. Haque, “Reduction of thermo-elastic damping with a secondary elastic field,” *Journal of Sound and Vibration,* vol. 318, no. 3, pp. 423 - 427, 2008.

[150] H. Pan, Y. Juang, S. Huang, L. Kuo et al, “Stress effect on mechanical loss of SiNX and amorphous silicon film deposited by PECVD method on silicon cantilever,” in *LSC meeting*, Pasadena, 2015.

[151] C. Zener, “Internal friction in solids 1.Theory of internal friction in reeds,” *Physical Review,* vol. 52, pp. 230 - 235, 1937.

[152] C. Zener, “Internal friction in solids 2. General theory of thermoelastic internal friction,” *Physical Review,* vol. 53, pp. 90 - 99, 1938.

[153] ANSYS, “Coupled field analysis guide,” ANSYS, Canonsburg, 2009.

[154] COMSOL, “Thermoelastic damping in a MEMS resonator,” 2017. [Online]. Available: https://www.comsol.com/model/thermoelastic-damping-in-a-mems-resonator-1439.

[155] M. A. Hopcroft, “What is the Young's modulus of Silicon,” *Journal of Microelectromechanical Systems,* vol. 19, no. 2, pp. 229 - 238, 2010.

[156] S. Lepage, “Stochastic finite element method for the modeling of thermoelastic damping in micro resonators,” Cranfield university, 1 December 2006. [Online]. Available: http://www.ltas-vis.ulg.ac.be/cmsms/uploads/File/Lepage_PhD.pdf. [Accessed 21 2 2018].

[157] X. Liu, D.R Queen, T.H Metcalf, J.E Karel et al, “Hydrogen-free amorphous silicon with no tunneling states,” *Physical Review Letters,* vol. 113, p. 025503, 2014.

[158] G. Stoney, “The tension of metallic films deposited by electrolysis,” *Mathematical, Physical and Engineering Sciences,* vol. 82, pp. 172-175, 1909.

[159] M. Lorenzini, “Modelling of the mode dependent mechanical losses in disc substrates,” in *GWADW*, Elba, 2016.

[160] B. Abbott, R. Abbott, T. Abbott, M Abernathy et al, “The Advanced LIGO detectors in the era of first discoveries,” *Physical Review Letters,* vol. 116, p. 131103, 2016.

[161] M. Granata, K. Craig, G. Cagnoli, C. Carcy et al, “Cryogenic measurements of mechanical loss of high reflectivity coating and estimation of thermal noise,” *Optics Letters,* vol. 38, pp. 5268 - 5271, 2013.

[162] S. Chao, H. Pan, L. Kuo, S. Huang et al, “Silicon nitride films deposited by PECVD method on silicon substrate for next generation laser interference gravitational wave detector,” in *Optical interference coatings 2016, OSA technical digest (Optical Society of America)*, 2016.

[163] L. Kuo, H. Pan, C. Lin and S. Chao, “Cryogenic losses for titania, silica, silicon nitride films and silicon substrates,” LIGO technical document LIGO-G1601703, 2016.

[164] X. Liu, T. Metcalf, Q. Wang and D. Photiadis, “Elastic properties of several silicon nitride films,” in *Proceedings of the materials research society symposium*, San Francisco, 2007.

[165] D. Southworth, R. Barton, S. Verbridge, B. Ilic et al, “Stress and silicon nitride: A crack in the universal dissipation of glasses,” *Physical Review Letters,* vol. 102, p. 225503, 2009.

[166] M. Fletcher, S. Tait, J. Steinlechner, I. Martin et al, “Effect of stress and temperature on the optical properties of silicon nitride membranes at 1550 nm,” *Frontiers in Materials,* vol. 5, no. 1, 2018.

[167] Norcada, “Norcada,” 2015. [Online]. Available: https://www.norcada.com/. [Accessed 15 01 2018].

[168] A. Alexandrovski, M. Fejer, A. Markosian and R. Route, “Photothermal common-path interferometry (PCI): new developments,” *International Society of Optics and Photonics,* vol. 7193, p. 101117, 2009.

[169] F. systems, “Thermovision A40M,” FLIR systems, 2005. [Online]. Available: http://www.movimed.com/files/Flir/A40M%20Datasheet.pdf. [Accessed 15 01 2018].

[170] COMSOL support, *Private communication,* COMSOL, 2018.

[171] P. Zwol, D. Vles, W. Voorthuijzen, M. Peter et al, “Emissivity of freestanding membranes with thin metal coatings,” *Journal of Applied Physics,* vol. 118, p. 213107, 2015.

[172] D. Heldman, “Convection,” in *Encyclopedia of agricultural, food and biological engineering*, New York, Marcel Dekker, 2003, p. 154.

[173] COMSOL, “COMSOL multiphysics user's guide v4.3,” COMSOL, 2012. [Online]. Available: http://people.ee.ethz.ch/~fieldcom/pps-comsol/documents/User%20Guide/COMSOLMultiphysicsUsersGuide.pdf. [Accessed 16 01 2018].

[174] K. Luke, Y. Okawachi, M. Lamont, A. Greta et al, “Broadband mid infrared frequency comb generation in a silicon nitride microresonator,” *Optics Letters,* vol. 40, pp. 4823-4826, 2015.

[175] M. Keevers and M. Green, “Optical properties of intrinsic silicon at 300K,” *Progress in Photovoltaics,* vol. 3, pp. 189 - 192, 1995.

[176] A. Jain and K. Goodson, “Measurement of the thermal conductivity and heat capacity of freestanding shape memory thin films using the 3w method,” *Journal of Heat Transfer,* vol. 130, p. 102402, 2008.

[177] MIT, “LPCVD silicon nitride material property database,” MIT, 19 02 2004. [Online]. Available: http://www.mit.edu/~6.777/matprops/lpcvd_sin.htm. [Accessed 23 07 2018].

[178] D. Jabaraj and M. Jaadar, “Theoretical calculation of resonant frequencies of the human alveolar wall and its implications in ultrasound induced ling hemorrage,” *International Journal of Bioscience,* vol. 3, p. 107763, 2013.

[179] Azo materials, “Silicon nitride properties and applications,” AZO, [Online]. Available: https://www.azom.com/properties.aspx?ArticleID=53. [Accessed 16 01 2018].

[180] A. Bell, “Private Communication,” 2018.

[181] J. Steinlechner, I. Martin, A. Bell, J. Hough et al, “Silicon based optical mirror coatings for ultrahigh precision metrology and sensing,” *Physical Review Letters,* vol. 120, p. 263602, 2018.

[182] M. Granata, L. Balzarini, J. Degallaix, R. Flaminio et al, “Internal frictionand youngs modulus measurements on SiO2 and Ta2O5 done with an ultra high Q silicon wafer suspension,” *Archives of metallurgy and materials,* vol. 60, no. 1, pp. 365 - 370, 2015.

[183] C. Speake, T. Quinn, R. Davis and S. Richman, “Experiment and theory in anelasticity,” *Measurement Science and Technology,* vol. 10, no. 6, pp. 430 - 434, 1999.

[184] J. Wu and C. C. Yu , “How stress can reduce dissipation in glasses,” *Physical Review B,* vol. 84, p. 174109, 2011.

[185] Norcada, "Crystalline silicon membranes," Norcada, 2015. [Online]. Available: https://www.norcada.com/products/silicon-membranes/. [Accessed 05 04 2018].

[186] J. Steinlechner, I. Martin, R. Bassiri, A. Bell et al, "Optical absoption of ion beam sputttered aSi coatings," *Physical Review D,* vol. 93, no. 6, p. 062005, 2016.

[187] D. Krick, P. Lenahan and J. Kanicki, "Electrically activve point defects in amorphous silicon nitride: An illumination and charge injection study," *Journal of Applied Physics,* vol. 64, p. 3558, 1988.

[188] G. Harry, H. Armandula, E. Black, D. Crooks et al, "Thermal noise from optical coatings in gravitational wave detectors," *Applied Optics,* vol. 45, no. 7, pp. 1569 - 1574, 2006.

[189] M. Abernathy, F. Acernese, P. Ajith, B. Allen et al, "Einstein gravitational wave telescope conceptual design study," *Technical report (ET- 0106C-10),* p. 451, 2011.

[190] O. Heavens, Thin Film Physics, London: Methuen and company, 1970.

[191] Norcada, "X-Ray microscopy windows specification sheet," 01 August 2013. [Online]. Available: https://www.norcada.com/wp-content/uploads/2013/10/Xray-Window-Specsheet.pdf. [Accessed 20 09 2018].

[192] P. Martin, A. Bendavid, M. Swain, R. Netterfield et al, "Mechanical and optical properties of the films of tantalum oxide deposited by ion assisted deposition," *MRS proceedings,* vol. 208, p. 583, 1993.

[193] S. Musikant, Optical materials, New York: Marcel Dekker Inc, 1985.

[194] L. Freund and S. Suresh, Thin film materials, Cambridge: Cambridge university press, 2003.

[195] Accuratus, "Silicon nitride ceramic properties," Accuratus, 01 January 2013. [Online]. Available: https://www.accuratus.com/silinit.html. [Accessed 20 September 2018].

[196] I. Martin, R. Bassiri, R. Nawrodt, M. Fejer et al, "Effect of heat treatment on mechanical dissipation in tantala coatings," *Classical and Quantum Gravity,* vol. 27, no. 22, 2010.

[197] A. Rua, F. Fernandex, M. Hines and N. Sepulveda, "Study of resonant frequnecies of silicon microcantilevers coated with vanadium dioxide films during the insulator-to-metal transition," *Journal of Applied Physics,* vol. 107, p. 053528, 2010.

Publications

This work resulted in the following publications:

M. Fletcher, S. Tait, J. Steinlechner, I. Martin et al, "Effect of stress and temperature on the optical properties of silicon nitride membranes at 1550 nm," *Frontiers in Materials,* vol. 5, no. 1, 2018.

J. Steinlechner, I. Martin, A. Bell, J. Hough et al, "Silicon based optical mirror coatings for ultrahigh precision metrology and sensing," *Physical Review Letters,* vol. 120, p. 263602, 2018.