High Precision, Compact Inertial Sensors for Use in Gravitational Wave Detectors

Von der QUEST-Leibniz-Forschungsschule
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades

DOKTOR DER NATURWISSENSCHAFTEN
-Dr. rer. nat.-

genehmigte Dissertation
von

M. Sci.
Jonathan Joseph Carter

geboren am 31.10.1995 in Doncaster, Vereinigtes Königreich

2024

Abstract

Operating a gravitational wave detector requires a suite of high-performance inertial and displacement sensors. These sensors are part of a complex control system which isolates optics from seismic disturbances. Despite the state-of-the-art sensors deployed in current gravitational wave detectors, noise from the control system continues to be a problem for the Laser Interferometric Gravitational Wave Observatories (LIGO). This problem will worsen for the next generation of terrestrial gravitational wave detectors and, hence, needs technological developments in all sensing and control scheme aspects. One of these needs is for better precision inertial sensors in places that currently can not be probed, such as the suspension chains of the optics. We therefore need to develop high precision, compact inertial sensors.

Compact inertial sensors have limited test masses and higher resonance frequencies than their bulkier equivalents. Both of these degrade the performance of the sensors. In order to counteract this effect, a high mechanical Quality factor (Q factor) is required. High Q factors can be achieved with fused silica oscillators but require careful design of test mass suspensions to produce useful oscillation modes. Furthermore, to make a good inertial sensor, these oscillators must be integrated with a precise method of reading out the inertial motion of the test mass.

This thesis discusses the design and testing of such compact inertial sensors. Part I of this thesis begins with a discussion of inertial sensor design. The noise terms of relevance to these sensors are evaluated for their mechanical behaviours. The effect of every design parameter on the mechanical behaviour of oscillators is calculated using a combination Finite Element Analysis and analytical modeling. The responses to these design parameter changes can then be used to optimise the mechanical behaviours for noise performance.

Part II explores the development and testing of sensors. This part covers both the oscillating optic and optical readout methods. The tools for designing inertial sensors were put into practice. Several oscillator designs were produced and tested. The resulting designs achieve mQ products of 2000 kg, high enough for the gram scale designs to reach thermal noise levels comparable to state-of-the-art sensor available today. Furthermore, methods achieving high-reflectivity on the oscillators without degrading the Q factor is shown.

A high-precision optical readout scheme is developed. The scheme uses an optical resonator to achieve sub $\mathrm{fm}(\mathrm{Hz})^{-1/2}$ performance at frequencies above 80 Hz. It is shown that a frequency reference using an ultra-low expansion glass spacer can achieve satisfactory performance such that the laser's frequency noise is no longer a limiting factor. Huddle tests are used to evaluate the performance of some initial samples. The scheme does not quite reach the thermal noise floor due to some issues controlling the sensors, but promising initial results are shown.

This thesis makes the case that compact optical inertial sensors are a viable solution for seismic isolation systems in the next generation of gravitational wave detectors. Oscillators have been produced with excellent mechanical performance that can be combined with high-precision readout methods to produce state-of-the-art inertial sensors in a compact design. These are excellent candidates for use in gravitational wave detectors due to their compact size, vacuum compatibility, and excellent performance.

Keywords: Inertial Sensor, High Q, Interferometery

Acknowledgements

I certainly choose a very chaotic time to do a PhD in physics. The last few years have been difficult for everyone and we have all relied on others for support through them. It has been a great relief to weather the storm surrounded by supportive friends, family, and colleagues. Every time I have approached someone for help, people have responded with kindness and insight I could not hope to match. Environments like this do not happen by accident. I would therefore like to thank Professor Karsten Danzmann for building the AEI into the place it is today, and all my colleagues for helping make this a wonderful place to study and work. There are some who deserve special mention, however.

I first want to thank my parents for their support. I know me leaving the country was difficult for you, and I appreciate you not trying to stop me. Zuza, your support through the last few years has meant more than you could ever know.

Several people have suffered reading through my incoherent ramblings and helped turn them into the thesis presented today, and I wish to give thanks to you all; Stefan Ast, Gerald Bergmann, Jiang Ji Ho Zhang, Sina Köhlenbeck, Malte Misfeldt, Alvise Pizzella, Jasper Venneberg, and Kevin Weber.

Despite the variety of deranged ideas I present to them, both the electric and mechanical workshops have always been able to produce exactly what I need and so I would like to thank all the staff there for their skill and understanding. Gerrit Kühn and Patrick Oppermann, despite whatever terrible thing I have done to the CDS in any given week, you have always been able to get it back to working order and allow me to continue my work. Thank you for all the help, and sorry for all the trouble.

Throughout this thesis I have received excellent support and advice from supervisors who have helped keep me on track. I would therefore like to thank Professors Gerhard Heinzel, Oliver Gerberding, Felipe Guzman, and Andreas Tünnermann for guidance throughout the project.

Sina, words cannot express my gratitude; thank you.

Contents

Contents

___

Acronyms

ADC   Analogue to Digital Converter

AEI   Max Planck Institute for Gravitational Physics (Albert Einstein Institute

AOM   Acousto Optic Modulator

ASD   Amplitude Spectral Density

CDS   Control and Data System

EOM   Electro-Optic Modulator

ET   The Einstein Telescope

FEA   Finite Element Analysis

FDT   Fluctuation Dissipation Theorem

FSR   Free Spectral Range

IBS   Ion Beam Sputtering

IOF   Fraunhofer-Institut für Angewandte Optik und Feinmechanik (Jena)

ISO   International Organization for Standardization

HoQI   Homodyne Quadrature Inteferometer

LIGO   Laser Interferometer Gravitational-Wave Observatory

LZH   Laser Zentrum Hannover

NPRO   Non-Planar Ring Oscillator

OIS   Optical Inertial Sensor

PDH   Pound Drever Hall

PEEK   PolyEther Ether Ketone

PSD   Power Spectral Density

PZT   Piezoelectric (Stack)

SAS   Spring Anti-Spring

RMS   Root Mean Squared

TED   Thermoelastic Damping

TTL   Tilt to Length Coupling

UGF   Unity Gain Frequency

ULE   Ultra Low Expansion

Chapter 1

Gravitational Wave Detectors

Over a billion years ago, two black holes completed their long, spiralling dance in a violent climax that sent ripples through the universe. These ripples radiated in all directions through space at the speed of light until they reached the earth, where incredibly sensitive devices lay in wait. This event became the first direct observation of gravitational waves and led to many things; the birth of gravitational wave astronomy as a field, the awarding of the 2017 Nobel prize for physics to Rainer Weiss, Kip Thorne, and Barry Barish, and to me moving to Germany to research the development of inertial sensors. In this chapter, I will introduce the current state of gravitational wave astronomy, explain briefly what gravitational waves are, how we detect them, and how my and several other groups' research on inertial sensors is essential for the next generation of gravitational wave detectors.

1.1 Gravitational Wave Detection

A key result of Einstein's theory of general relativity is that, in the weak field limit, the solution of the field equation takes the form of a wave equation. These waves propagate through space at the speed of light through the medium of space-time, altering the geodesic path length of a particle in free fall. One of the greatest achievements in modern physics has been the development of technology capable of measuring gravitational waves.

Any object undergoing acceleration in a non-spherically symmetric manner will act as a source of gravitational waves; however, only the most extreme astrophysical events produce gravitational wave signals that we can hope to measure. The first direct signal physicists observed was a merger of two compact bodies, in this case two black holes. Only an event as cataclysmic as this would produce a signal that we could hope to measure. Even then, the ratio of the difference in geodesic length between two points with and without the presence of a gravitational wave, called the characteristic strain, is of order 10^{-21}. This scale alone already lays bare the enormous challenge that measuring gravitational waves presents.

Despite the scale of the challenge, the Laser Interferometer Gravitational-Wave Observatory (LIGO) collaboration observed the first direct evidence of gravitational waves emitted by the merger of a binary black hole in 2015 [1], during the advanced LIGO [2] first observational run, swiftly followed by another detection [3]. The second observational run led to another ten confirmed signals [4], which included the remarkable event GW170817 [5]. Not content with simply being the first detection of a binary neutron star, the then-recent addition of the advanced Virgo detector [1, 6] to the gravitational wave observatory network allowed for a good degree of sky localisation. Ultimately, this enabled electromagnetic observatories to measure the Gamma Ray Burst corresponding to this merger [7], displaying the capabilities of gravitational wave astronomy for all to see. After some further upgrades to the detectors, the third observation run commenced, where detection candidates became almost a weekly occurrence [8]. Building these large catalogues enables astrophysicists to perform the statistical analysis necessary to understand the cauldrons from which the earliest stars formed and answer some of the most fundamental questions to all mankind, those of our universe's origins.

The LIGO and Virgo detectors, which were the first to achieve sufficient sensitivity to detect gravitational waves, are all ground-based interferometric detectors. There are now several of these detectors constructed, and under construction around the world [2, 6, 9]. These are mostly direct upgrades on the sites of the initial gravitational wave detectors and are usually classified as the second generation of gravitational wave observatories.

While the specifics of these detectors vary from site to site, the overall principle remains the same. Each is a Michelson interferometer with Fabry Perot cavities in the arms. The arms are typically kilometres long; for example, LIGO detectors are 4 kilometres long. When gravitational waves pass through the interferometer, they modulate the phase of the light detected by changing the travel time of light in the two arms. The two arms' phase difference is read out and searched for signatures of gravitational wave signals. When multiple detectors have simultaneous signals above a statistical significance threshold, a gravitational wave has been detected [1, 3, 4, 5].

All noise sources in the detectors must be suppressed to a level below the signal of gravitational waves, which are typically of order $10^{-18}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ for the band of a few Hertz to kHz. This requires extremely stable lasers for precision measurement [10], high vacuum environments [2], and sophisticated suspension systems to isolate the optics from environmental noises [11, 12, 13, 14].

1.1.1 Seismic Noise and Seismic Isolation Systems

The detection of gravitational waves requires the use of freely falling test masses. Terrestrial detectors have the problem that the ground of the building is moving at a level many orders of magnitude higher than a signal from a gravitational wave source. The sources of this motion are varied, from oceans waves to local disturbances like cars and trains. We must therefore take great caution to isolate our detector from these unwanted disturbances.

Seismic isolation systems in gravitational wave detectors must reduce the seismic noise in the detection bandwidth to below the level of the gravitational wave signal and control the optics precisely at the desired, narrow operating point. To achieve these aims, gravitational wave detectors deploy both passive and active isolation systems [15, 16, 17].

The passive suspension system suppresses the noise in the measurement bandwidth. A passive system is one with no controls to reduce the transmitted motion. Gravitational wave detectors use suspended mirrors in multistage pendula, which can swing freely. The transmission of motion through a pendulum is shown in Figure 1.1 with an intuitive demonstration shown in Figure 1.2. The behaviour of these suspensions is strongly dependent on frequency. Low-frequency displacements, those lower than the natural frequency of the pendula, are fully transmitted through the suspension to the optic and must be suppressed with other means. Above the resonance, motion is suppressed, which goes as $1/f^2$. Each additional stage adds another resonance peak and a further $1/f^2$ slope above it, as is seen by 1.1. Gravitational wave observatories usually use at least four stages to suspend their core optics to create a $1/f^8$ slope. The slope creates a seismic cliff that governs the detectors' lower frequency limit. The only ways to lower the cliff frequency on earth are to decrease the motion reaching the suspensions by choosing seismically quieter locations [9, 18, 19, 20], better motion suppression in the control systems [15, 16, 17, 21, 22], and making the suspension chains longer.

Because motion at frequencies below the resonance of the passive suspension is directly transmitted to the test masses, alternate means of motion reduction are needed. For example, at LIGO transmitted motion is reduced with a combination of feedforward and feedback controls. These controls act on a platform mounted to the ground with actuators, which can drive the platform in all degrees of freedom [15, 16, 17]. The platforms and ground around them are monitored with motion sensors to know how these actuators should react. Large high performance inertial seismometers and geophones inform most of the sub-Hertz response. The sensor outputs are combined using a sophisticated control scheme to take each of their most sensitive bandwidths to create a "Super Sensor" to drive the platform [23].

In addition to the direct noise contribution of seismic and control noise in the detector bandwidth, the extraordinarily narrow operating point of the detectors places
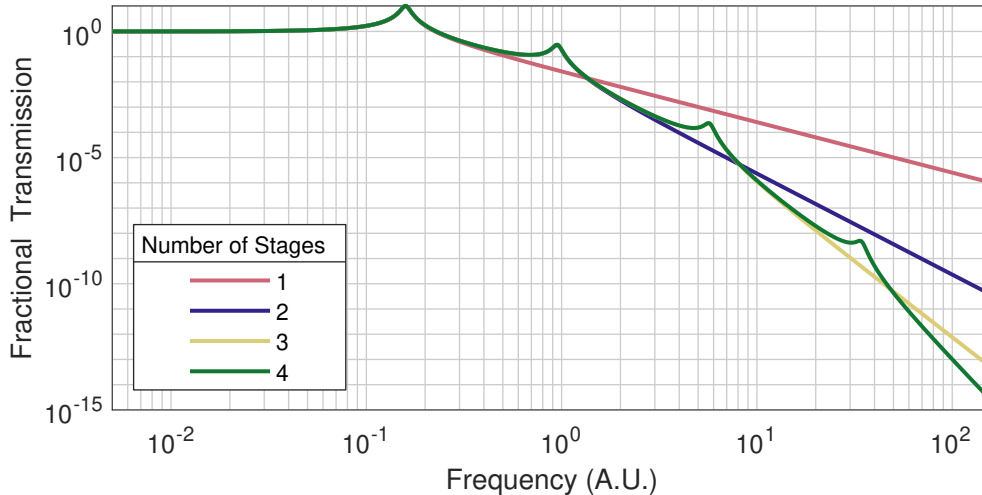
Figure 1.1: The mechanical transmission of motion through suspensions with different numbers of stages. Each stage adds $1/f^2$ slope to the overall gradient of transmission at higher frequencies, while also adding an additional resonance. The resonance values here are arbitrary. At frequencies below the first resonance, all motion is transferred through the system; an intuitive demonstration is shown in Figure 1.2.

stringent requirements on the residual vibration outside of the measurement band [12, 24, 25]. For the observatories to be operational, the suspended mirrors must be held at an ideal operating point [26, 27, 28, 29]. If the low-frequency ground motion is insufficiently suppressed, it can cause the mirrors to move too far from the operating point. Losing and re-acquiring a lock is a time-consuming process [30, 31, 25]. The duty cycle in the third observation run of the two LIGO detectors was 74.6% for the Hanford and 77% for Livingston sites [32], while Virgo achieves above 80% in science mode [31]. Due to the gravitational wave detectors' scale, the opportunity cost of downtime becomes huge. Furthermore, the need for coincidence data amplifies this effect. The network is only sensitive to extreme sources with only two online detectors. However, without at least three detectors online, triangulation with enough precision for optical counterpart measurements becomes unlikely [5]. With only a three-detector network, as Kagra is still undergoing commissioning to meet design sensitivity, and assuming no coherence between detector offline times, a three-network detector is in operation only $\sim 46\%$ of the time. The latter of these assumptions is not strictly true due to seasonal variations in seismic activity [33, 34], as well as servicing scheduled to occur simultaneously between the sites. However, it still gives a reasonable estimate of the problem.

A better informed Super Sensor would reduce detector downtime, and lead to more operational time. [23]. Much work is therefore dedicated to installing new sensors [35, 36, 28] and better using the sensors available to control the detector [22].

1.2   From Second to Third Generation

The field now sits on the cusp of 3rd generation gravitational wave detectors, which intend to push ground-based detection to its limits. The Einstein Telescope (ET)
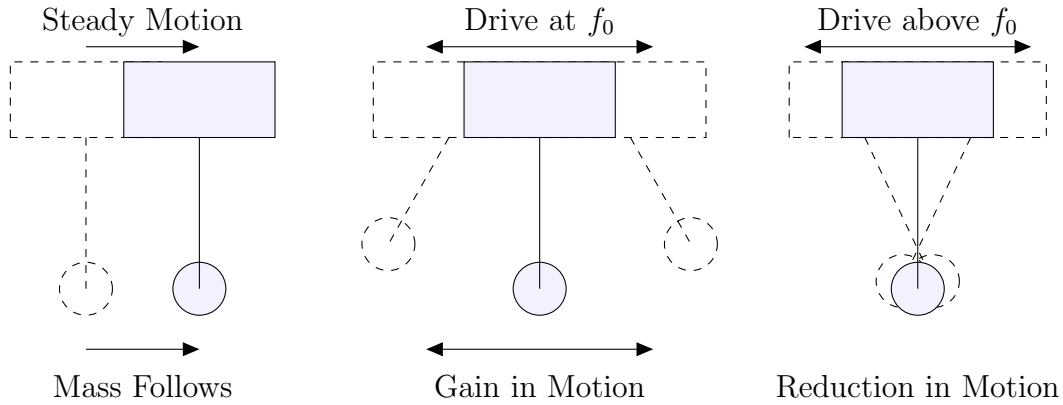
Figure 1.2: A simple experiment which can be done with a pendulum to understand passive isolation. Swinging slowly will result in the bob following behind. Driving close to the resonance will result in a huge gain in motion for the bob, while driving quickly will result in the bob staying inertially still.

detector [37] led by the European community and Cosmic Explorer [38] by the USA begin their initial developments. Both build upon the understanding gained from LIGO and Virgo. They aim to push detection capabilities to encompass nearly all black-hole and neutron star mergers in their sensitive frequency bands.

Because of the gap in sensitive bands between earth and space-based detectors, there is a strong desire to push the lower frequency sensitivity to its lowest possible limit. Doing so enables multi-band astronomy where interesting signals detected in the Laser Interferometer Space Antenna (LISA) band can be tracked as the black holes inspiral. Then the merger and ringdown can be observed by the ground-based network. [39]

1.2.1   Einstein Telescope

ET will likely be Europe's primary contribution to the third generation of gravitational wave detectors. The detector is now in the phase where key design decisions, such as optical layouts, targeted noise performance [37, 40], and locations [18, 19, 41, 42] must be decided. Now is the ideal time for developing technologies as proof of principles for further refinement in the design, which the telescope can use to achieve its performance.

As the design is still in the planning stages, all designs are liable to changes. However, several features are common to most proposals for the overall design, which make ET unique from other gravitational wave detectors:

- Underground laboratory. The detector will be buried underground to isolate better seismically, and lower Newtonian noise contributions [43].

- A triangular geometry so to form a null instrument. This is necessary if ET proves to be the only detector operating at such high precision, in order to statistically distinguish between noise and signals.

- Two combined inteferometers, optimised for different frequency bands. One as a low frequency, cryogenic inteferometer, the other a high frequency, high laser power inteferometer.
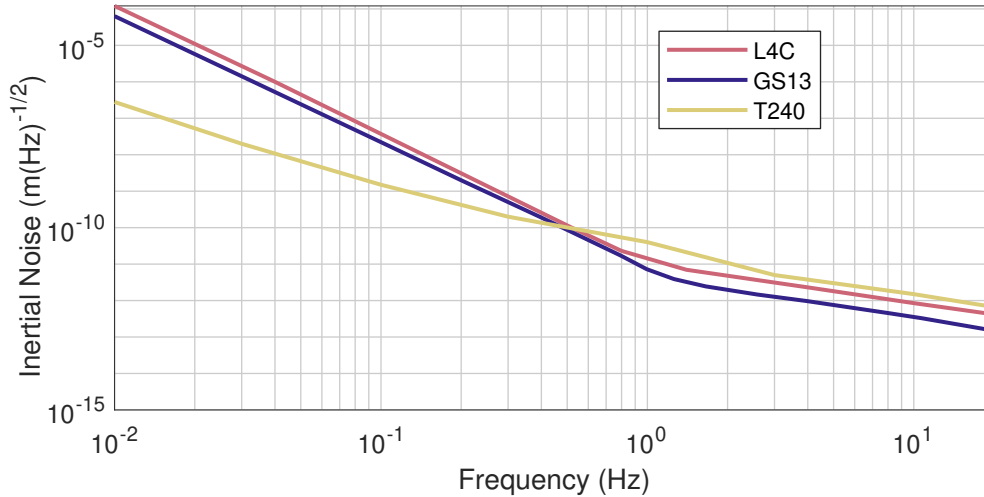
Figure 1.3: The noise performance of several inertial ground sensors used at LIGO. The sensors are the trillium 240s force feedback seismometer [51], Geotech Model S-13 [52], and the L4C geophone [53]. (acceleration equivalent (D.1))

A key aim of the detector is achieving a lower frequency sensitivity down to the Hz regime. The frequency band of a 1-10 Hz is of great interest to the community because of the vast variety of astrophysical sources proposed to lie there [44, 45, 46].

To measure this bandwidth will require pushing the control bandwidth down into the sub-Hz area. To be able to actively isolate at such low frequencies with sufficient precision a whole new suite of high sensitivity inertial sensors, tilt sensors [35, 36], displacement sensors [47], tilt decoupled sensors [48, 49], and sophisticated control schemes [15, 16, 17, 21, 22, 50] will be required. A key problem with many of the high performance classical inertial sensors is their size. Many use payloads in the kg scale, require electronics, and are not vacuum compatible. This makes them unsuitable for use in the suspension chains of gravitational wave detectors and limits how we are able to probe and control the payloads.

1.3   Commercial Inertial Sensors

In Chapter 3 I show the factors which lead to a low noise inertial sensor. These are a low natural frequency, a high mass, and a high mechanical Quality factor (Q factor), coupled with a low noise readout. Most high-performance inertial sensors in LIGO use a combination of high mass and very soft springs. Whilst this provides high-performance seismometers [51, 52, 53], the bulky nature of the sensors means mounting them on isolation stages becomes difficult. Furthermore, their large size and complicated electronics are not vacuum compatible, so they must be held in vaults or kept far away from the platforms. The distance reduces the coherence between the seismometer measurement and ground motion at the platform [15]. The performance of some of the highest performing inertial sensors used in LIGO are shown in Figure 1.3

Micro-Electro-Mechanical Systems (MEMS) accelerometers have been an area of rapid innovation over the last 30 years [54, 55, 56]. As these devices are for use in

consumer electronics, such as phones and navigation systems [57], most of the research in this field goes towards lowering costs [55] and self-calibration [58]. Many of these MEMS achieve quality factors close to the bulk loss of their material, with Q factors well into the millions [59]. Despite this, they can not achieve the same level of noise performance as large high performance seismometers simply because typical test masses are of order of μg.

To further open the possibility space for inertial sensors, we would like to take a best of both approaches, with oscillators with high Q factors from MEMS designs, but large test masses and lower resonance frequencies like conventional high performance inertial sensors. Ideally these would then still fit in compact housings with their readout methods. Doing so would produce high precision compact inertial sensors.

1.3.1   High Quality Factor, Precision Inertial Sensors

Guzman et al. were the first to show a high Q oscillator with a gram scale mass, leading to a mQ product, which acts as a figure of merit of a sensor's noise floor, comparable to a GS13 [60, 61]. They used a $100\,\mu$g scale test mass and achieved a Q factor of order 10^5 before some necessary readout tools were attached, which led to the degradation of the quality factor to around 40,000.

The work showed two key innovations essential for the future of the field. The first was using laser-assisted etching of fused silica for manufacturing the parts. Fused silica is well studied as having a low bulk mechanical loss at room temperature [62], as well as many other desirable mechanical properties leading to low Thermoelastic Damping (TED) and Akhiezer loss coefficients. The desired geometries were manufactured using a new 3D subtractive laser-assisted chemical etching technology. The technology, developed by the company FEMTOprint [63], uses a femtosecond laser to activate the areas of geometry that need removing. Then a wet chemical hydrofluoric acid is applied to remove these areas, leaving only the desired structure.

The second innovation Guzman et al.'s work showed was the geometry required to make a linear in-plane oscillation. Using a parallelogram structure, they achieved a soft device in one degree of freedom, but stiff in all others. The linear mode of oscillation makes the calibration possible for the large dynamic range required for high Q oscillations. The geometry of this motion is shown in Figure 1.4.

The work by Guzman et al. [60] was pioneering, but designed to be a high-frequency accelerometer sensitive at kHz frequencies. For the technology to be useful to gravitational wave astronomy, the geometry must be adapted to support larger test masses and drive the sensitive band down to the sub-Hz regime. Combining a high-performance oscillator with a technique capable of reading out the test mass motion to a sufficient level will lead to a revolution in the field of inertial sensing, with compact optomechanical sensors appearing in a wide range of fields.

1.4   Thesis Structure

This thesis discusses my work developing compact sensors that gravitational wave detectors can use in the active isolation of test masses. The sensors developed here aim to achieve the high performance of the best seismometers in a compact, vacuum-compatible

(a) Parallelogram Oscillation   (b) Single Flexure Oscillation
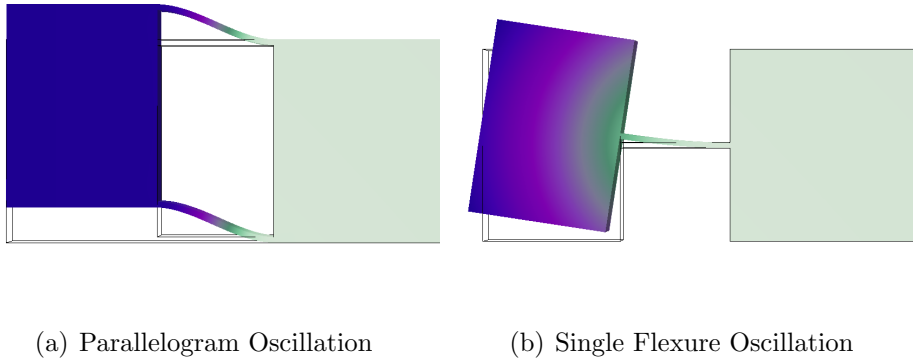
Figure 1.4: The fundamental mode of oscillation for two versions of an oscillator. The colour indicates the extent of motion of that point with arbitrary scaling relative to being at rest, shown by the black outline. The base of each oscillator is fixed on the bottom surface of the part. The parallelogram oscillation is confined to one dimension, meaning any point of the test mass is suitable for readout. Any beam hitting the single flexure geometry must be carefully positioned or can lead to calibration errors. Furthermore, the transfer function of the oscillator would require a more complicated transfer function to model.

package. The production of such sensors would be an immeasurably valuable tool to not only the seismic isolation of gravitational wave detectors, but any field with a need for precision measurement of acceleration with limited weight/space requirements. While we are a long way from the regular deployment of such sensors, it is my hope that this thesis can significantly contribute to the field.

The thesis is split into two parts. The first is a discussion dedicated to the design of low noise inertial sensors. This part is meant as a exploration of general inertial sensor design that other groups could use to develop their own designs. I first explore how inertial sensors work in Chapter 2, before discussion of how to evaluate noises in inertial sensors through Chapter 3. Chapter 4 then discusses considerations one must make when designing inertial sensors, and discusses the tools available to achieve those goals.

Part II discusses my development of inertial sensors for use in gravitational wave detectors. The part is broken down into two chapters, which correspond to the two key areas of inertial sensor design, the oscillator and the readout. The development of low thermal noise oscillator is discussed through Chapter 5. I then explore how to read out ground motion from such oscillators in Chapter 6. In Chapter 7 I present a process for designing inertial sensors and show how it was used to make the sensors presented throughout the project.

The final chapter of this thesis, Chapter 8, discusses the broader context of this work. I conclude with what I think the future of this exciting field will be.

Part I

Design of Inertial Sensors

Chapter 2

Inertial Sensors

The need to measure inertial motion occurs in almost every mechanical system. Each application has its challenges and system requirements, which warrant a new kind of sensor. Gravitational wave astronomy is no different, with a strong need for high performance and tilt-free sensors. Despite the differences in design philosophy, most inertial sensors operate on very similar principles, with two areas being the focus. The first is how to suspend a test mass and keep it inertially stable; the second is how to measure the change in test mass position. This chapter discusses both of these in the general case of inertial sensors.

2.1   Inertial Sensors

An inertial sensor is a damped harmonic oscillator in a box, with the model used for the following derivations shown in Figure 2.1. The only positional change one can measure is the distance between the edge of the apparatus and the test mass, Δx. It is possible to relate this quantity to the desired ground motion x_{g} by following how the system responds to external stimuli.

The equation of motion of a velocity damped (with the case of a displacement damped oscillator discussed in Appendix A) test mass is given by

$$m\ddot{x}_{\mathrm{m}} + b\dot{x}_{\mathrm{m}} + kx_{\mathrm{m}} = f_{\mathrm{ext}}, \tag{2.1}$$

where, k is the spring constant, b is the damping ratio, m is the suspended mass, and f_{ext} is the force of the spring compressing under motion from ground

$$f_{\mathrm{ext}} = kx_{\mathrm{g}} + b\dot{x}_{\mathrm{g}}. \tag{2.2}$$

Substituting in for this force and changing to a function of the measured displacement

$$\Delta x = x_{\mathrm{m}} - x_{\mathrm{g}} \tag{2.3}$$

gives

$$m\Delta\ddot{x} + b\Delta\dot{x} + k\Delta x = -m\ddot{x}_{\mathrm{g}}. \tag{2.4}$$

The constants are now replaced with their more natural counterparts

$$\omega_0 = \sqrt{k/m}, \tag{2.5}$$

$$Q = \frac{\omega_0 m}{b}, \tag{2.6}$$

where ω_0 is the natural angular frequency of oscillation of the system and Q is the mechanical Quality factor. The Q factor is effectively a measure of energy loss throughout the system, with high Q factors meaning the energy is well preserved. Its physical definition is

$$Q = 2\pi \frac{\text{Energy Stored}}{\text{Energy Dissipated per Oscillation Cycle}}, \tag{2.7}$$

which makes it a very intuitive property for discussing the thermal noise, as discussed in section 3.2. A Fourier transform can now be taken of equation 2.4 to give

$$\omega^2 X_{\mathrm{g}} = -\omega^2 \Delta X + \omega_0{}^2 \Delta X + \frac{i\omega_0\omega}{Q}\Delta X, \tag{2.8}$$

where the capital denotes the function is Fourier transformed and therefore a function of ω. Rearranging for the ratio of the two displacements results in the transfer function

$$\frac{\Delta X}{X_{\mathrm{g}}} = \frac{-\omega^2}{\left(\omega^2 - \omega_0{}^2 - \frac{i\omega_0\omega}{Q}\right)}. \tag{2.9}$$

This produces the transfer function shown in Figure 2.2. It can be seen from this that the sensor acts as a high-pass filter. To get back x_{g} from Δx this response must be
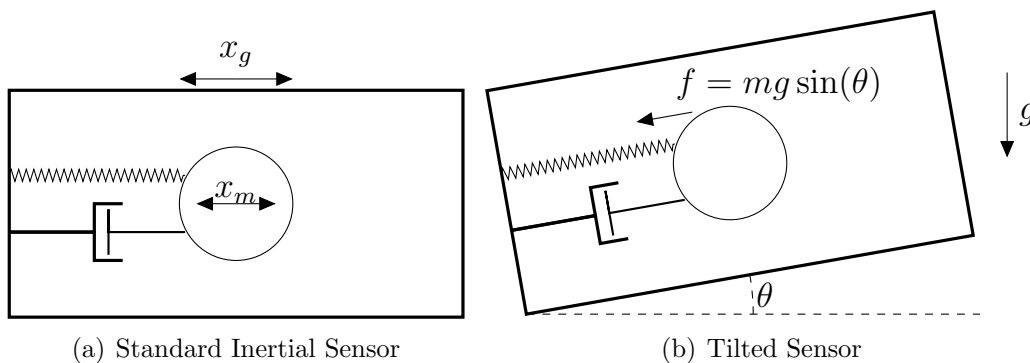
(a) Standard Inertial Sensor

(b) Tilted Sensor

Figure 2.1: (a) The basic layout of a simple seismometer. Ground motion x_g moves the outer case. By compressing the spring of spring constant, k, and damping constant, b, a force is exerted on the mass, causing a motion of x_m. (b) How gravity and tilt of a seismometer by an angle θ applies an additional force to the test mass.

inverted. Doing so produces a large amplification of low-frequency noises, so a low natural frequency is desired.

Some other considerations are necessary to bear in mind whilst working with inertial sensors.

The first is that the system will have other modes of oscillation in its design. These higher modes of oscillation will often present problems to the readout. So it often becomes beneficial that the sensor is soft in one degree of travel so that its fundamental mode is low frequency, but very stiff in all other degrees of freedom [61, 64].

The second is that the response is assumed to be linear, regardless of displacement. Depending on the system design, there will always be some point where this is false due to overextending suspensions. Unless well modelled, this limits the dynamic range of the sensor. One can either design the suspension to survive the dynamic range or use a force feedback control system to keep the test mass centred on the frame [65].

2.1.1 Tilt Injection in Inertial Sensors

Tilt coupling to translation is one of the most problematic issues for inertial sensors at low-frequency detection. It proves especially problematic for the LIGO Hanford site, as the desert environment allows strong winds to blow [28]. When these hit the detector, it acts like a large sail causing it to roll [66].

To understand how problematic this becomes, the calculation in section 2.1 is repeated, but the external force is replaced by

$$f_{\text{ext}} = kx_{\text{g}} + b\dot{x}_{\text{g}} + mg\theta, \tag{2.10}$$

with the additional term coming from the geometry as shown in Figure 2.1 (b); θ is assumed to be small. The final equation for X_g then becomes

$$X_{\text{g}} = \Delta X \left(-1 + \left(\frac{\omega_0}{\omega} \right)^2 + \frac{i\omega_0}{\omega Q} \right) - \frac{\text{g}}{\omega^2} \Theta. \tag{2.11}$$

The final term reveals the problem posed, as the $1/\omega^2$ means a massive amplification in tilt coupling at low frequencies.
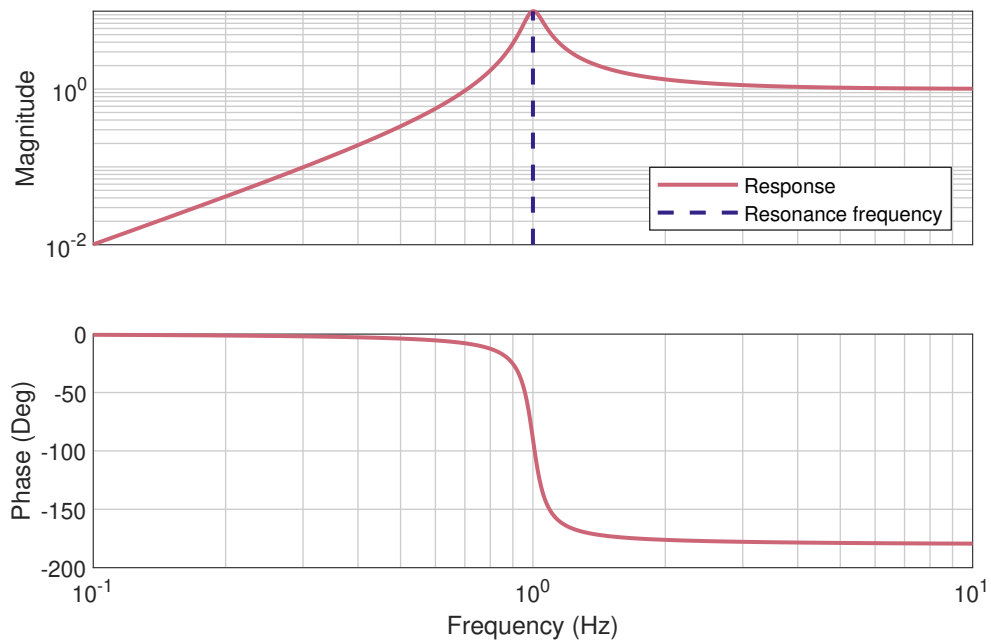
Figure 2.2: The transfer functions of the ground motion to seismometer output, for a seismometer with a Q factor of 10 and an f_0 of $1\,\text{Hz}$.

Chapter 3

Noise in Inertial Sensors

Inertial sensors must measure the displacement of a suspended test mass. When doing so we must contend with spurious noise sources. This chapter looks at these noise sources and the terms that control them. Understanding the design's limiting noise sources is essential to minimise them. This chapter discusses the origins, contributions, and fundamental limits on noise sources that change the test mass's position and limit our ability to measure its position.

3.1   The Units of Noise

Throughout this thesis, the Power Spectral Density (PSD) and its square root, the Amplitude Spectral Density (ASD), are used to evaluate the noise performance of devices. The PSD can be calculated directly from the Wiener–Khinchin theorem, which states "The power spectral density of a stationary, random process is given by the Fourier transform of the autocorrelation of the process" [67], or numerically

$$\tilde{S}^2(\omega) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S(t+\tau)S(t)e^{(i\tau\omega)}dtd\tau, \tag{3.1}$$

where S is the signal, t is the time series of the data, and τ is an arbitrary time offset between the two signals, which is integrated over. The PSD effectively measures the noise power per bandwidth as a function of frequency. A bandwidth of $1\,\mathrm{Hz}$ is used as a common reference.

Inertial sensors add a layer of complexity to this, as it is important to distinguish between the different noise signal units. This thesis uses three different noise units due to the different sources of noise coupling to different parts of a device. These are:

- Displacement noise which describes any noise terms that directly correlate to errors in the position of the test mass, mostly related to readout methods. It has units $\mathrm{m(Hz)}^{-1/2}$.

- Acceleration noise which describes any noise which leads to a direct force applied to the test mass. It has units $\mathrm{ms}^{-2}\mathrm{(Hz)}^{-1/2}$.

- Inertial equivalent noise which describes any device intended to measure ground motion's performance as it gives a measure of the noise the device will add to the measurement in terms of displacement. It has units $\mathrm{m(Hz)}^{-1/2}$.

While inertial equivalent noise and displacement noise have the same units, they are not equivalent. The displacement noise does not account for the coupling of the ground motion to the test mass motion and only describes the actual motion undergone by the test mass. As it depends on the oscillator's transfer function, it should not compare the performance of different sensors directly. Rather it is useful when assessing the readout of the test mass. Both inertial equivalent noise and acceleration noise are measurements of the performance of a device. Figure 3.1 shows how a flat displacement noise on the test mass would be projected onto the inertial or acceleration noise.

Typically, inertial equivalent units are used amongst the ground-based gravitational wave community, while the space-based community use acceleration units. This thesis presents noise budgets in inertial equivalent units throughout, but the acceleration equivalents for relevant figures are included in the appendix.

It is possible to convert directly between these units. Acceleration noise is related to displacement noise by the transfer function

$$\tilde{A}(\omega) = \tilde{X}(\omega) \cdot \mid \omega_0^2 - \omega^2 + \frac{i\omega_0\omega}{Q} \mid \tag{3.2}$$

where A is the acceleration noise, X is the displacement noise ω_0 is the angular resonance frequency and ω is the angular Fourier frequency, and Q is the Q factor. Inertial
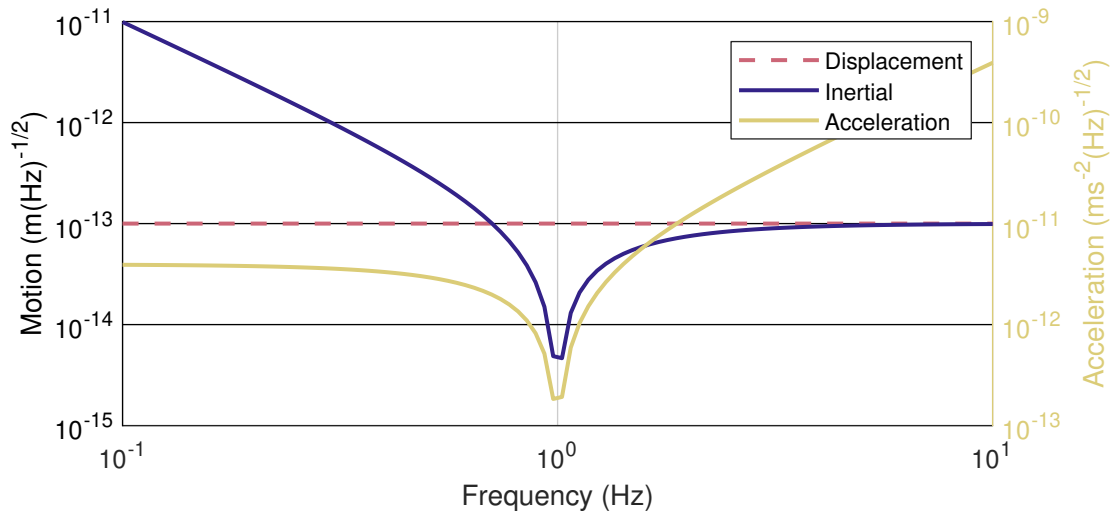
Figure 3.1: The figure shows how a flat displacement noise at $10\text{-}13\,\text{m(Hz)}^{-1/2}$ (pink, dashed) of the test mass would couple to the overall noise of a oscillator with a Q factor of 50 and a f_0 of 1 Hz. Shown in inertial noise (blue), or acceleration noise (yellow, right axis).

equivalent units are related to acceleration by

$$\widetilde{IEq}(\omega) = \tilde{A}(\omega)/\omega^2. \tag{3.3}$$

3.2 Thermal Noise

The Fluctuation Dissipation Theorem (FDT) describes how the intrinsic thermal fluctuations of a system couple to the dissipation of energy from that system's mechanical oscillations [68, 69]. The FDT is a two-way process, meaning heat dissipated by the oscillating motion excites the system energy, which acts as a heat bath. The energy in this heat bath also leads to the direct motion of the suspended mass. The latter of these is the source of what is here called thermal noise. However, because it is related to the mechanical dissipation from the system, these properties can be used to estimate the thermal noise level in an oscillator.

The strength of this coupling is described by

$$\tilde{A}^2(\omega) = \frac{4k_{\mathrm{B}}T}{m^2}\mathbb{R}(Z(\omega)) \tag{3.4}$$

where K_{B} is the Boltzmann constant, A is the acceleration of the test mass, and Z is the mechanical impedance. The mechanical impedance of any linear system is defined as the relation between external force, F_{ext}, and velocity, $\dot{x}$,

$$F_{\mathrm{ext}} = Z\dot{x}. \tag{3.5}$$

From here, it is possible to calculate how different channels of thermal noise contribute to the system from their mechanical impedance. We consider the cases of a damping ratio, b, proportional to displacement and velocity below.

Viscous Thermal Noise

A typical example of this is a gas damped (or any oscillator whose damping is proportional to velocity) simple harmonic oscillator, whose equation of motion is

$$f_{\text{ext}} = m\ddot{x} + b\dot{x} + kx \tag{3.6}$$

for a proof mass m, damping coefficient b, and spring constant k. Z can be calculated by substituting equation 3.6 into 3.5

$$Z = \frac{m\ddot{x} + b\dot{x} + kx}{\dot{x}} \tag{3.7}$$

$$\tag{3.8}$$

which we take the Fourier Transform of to express Z as a function of frequency

$$Z = m\frac{\int_{-\infty}^{\infty} \left\{ \frac{d^2}{dt^2}X(\omega)e^{i\omega t} + \frac{b}{m}\frac{d}{dt}X(\omega)e^{i\omega t} + \frac{k}{m}X(\omega)e^{i\omega t} \right\} d\omega}{\int_{-\infty}^{\infty} \frac{d}{dt}X(\omega)e^{i\omega t}d\omega}. \tag{3.9}$$

As only the real terms are needed, we are left with

$$\mathbb{R}(Z(\omega)) = b = \frac{\omega_0 m}{Q} \tag{3.10}$$

where the Q factor has been substituted using the definition from equation 2.6 from the previous chapter. Substituting this into Equation 3.4 gives the acceleration spectral density

$$\tilde{A}(\omega) = \sqrt{\frac{4k_{\text{b}}T\omega_0}{Qm}}. \tag{3.11}$$

Structural Damping Noise

For a bending flexure, the damping term is proportional not to the velocity term but to its displacement. This is because, as discussed in section 3.3, damping in these devices is often related to structural deformation. This partial derivation follows a full derivation in Agatsuma's thesis [70][chp 3]. The equation of motion of this system is described by

$$f_{\text{ext}} = m\ddot{x} + \kappa x \tag{3.12}$$

where κ is the spring constant also containing the damping information. We can then take the Fourier transform of this equation

$$F_{\text{ext}} = -m\omega^2 X + k(1 - i\phi)X. \tag{3.13}$$

Here κ has been broken down into its real part, k describing oscillatory motion and the imaginary part ϕ, which describes the fractional energy dissipated per radian of oscillation. Comparing the definition of ϕ here and that of the Q factor, it can be seen that

$$\phi = \frac{1}{Q}. \tag{3.14}$$

From this, the mechanical impedance is given by

$$Z = \frac{-m\omega^2 X + k(1 - i\frac{1}{Q})X}{i\omega X} \tag{3.15}$$

$$\tag{3.16}$$

Again removing the imaginary terms leaves

$$\mathbb{R}(\mathbf{Z}(\omega)) = \frac{k}{Q\omega} = \frac{\omega_0^2 m}{\omega Q} \tag{3.17}$$

following through as before, we find that

$$\tilde{A}(\omega) = \sqrt{\frac{4k_{\mathrm{b}}T\omega_0^2}{m\omega Q}}. \tag{3.18}$$

The additional factor ω_0/ω results in an additional noise at frequencies below resonance compared to a viscous damping model. Saulson dedicates some effort to describing this in his works [71, 72, chp 7].

These results show the factors needed for a low thermal noise in inertial sensors. They are a high Q factor, a high test mass, and a low ω_0. By building inertial sensors in a compact form we loss the ability to increase the mass, and have higher ω_0. To compensate for this we must make the Q factor as high as possible.

3.3   Mechanical Loss Estimation

Every non-ideal oscillator is damped. This section covers the leading terms in small mechanical oscillators such as TED, bulk loss, and surface loss.

3.3.1   Thermoelastic Damping

TED is caused by an irreversible heat flow process, as shown in Figure 3.2. When a flexure bends, one side extends and one contracts. Areas under contraction heat up, and those that are extended cool down. Local heating and cooling across the beam leads to a temperature gradient. Heat flowing across the gradient is irreversible, so the system's overall entropy must increase, i.e. energy has left the oscillating system and must couple to the heat bath. This process and effect has been well derived and described in many places [73, 74, 75], and only the relevant results are stated here.

TED in thin beams can be estimated by the Zener approximation, as is the method used by [74]. A derivation of a direct solution is presented in Appendix B. The key result from this is that the Q factor from TED is given by

$$Q_{\mathrm{TED}}^{-1} = \frac{\alpha^2 E T}{C_{\mathrm{p}}} \frac{\omega\tau}{1 + \omega^2\tau^2} \tag{3.19}$$

where α is the coefficient of thermal expansion, E is the Youngs modulus, T is the temperature, ρ the density, C_p the volumetric heat capacity at constant pressure, *omega*
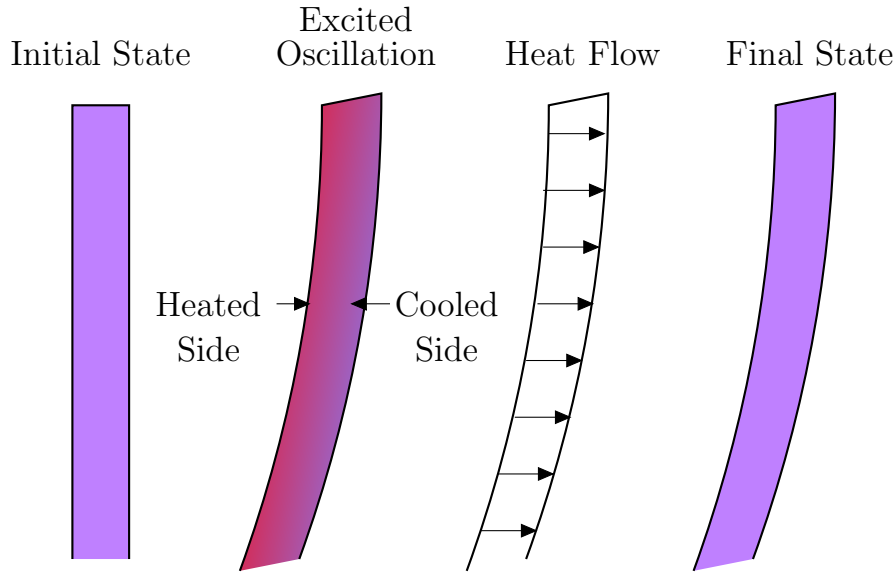
Figure 3.2: The process of TED in a flexure during oscillation. Heat flow is caused by the local heating of the compressed side and cooling of the expanded side. As the entropy of the system increases over one cycle, it is not reversible. Hence, some energy must be dissipated to the environment.

the angular Fourier frequency, and τ is the thermal relaxation time. In the case of a thin beam, this is approximated by

$$\tau = \frac{d_{\mathrm{fwidth}}^{2}}{\pi^2 \chi} \qquad (3.20)$$

with d_{fwidth} being the flexure thickness, and χ the thermal diffusivity. Other geometries have been solved, but solutions become increasingly complex [75, 76]. It is for this reason that for more complex geometries Finite Element Analysis (FEA) is often used to predict the effect of thermoelastic damping [77, 78, 79].

The Q factor has an explicit dependence on frequency. When substituted in equation 3.18, at frequencies significantly below the thermal relaxation time, the frequency dependence cancels out, and TED loss behaves like viscous damping. The second key effect is the occurrence of the TED peak. It is, therefore, advisable to tune design parameters such that the TED peak is outside of sensitive frequency ranges.

3.3.2   Material, Bulk, and Surface Loss

Another leading loss term is the intrinsic loss of the material, ϕ_{Mat}. The bulk of any material will have channels where energy can be dissipated. Typical values for fused silica are of the order $\phi_{\mathrm{Bulk}} = 10^{-7}$ at room temperature [62]. The surface of the material will have a less regular lattice and there may be contaminants, typically treated as having a bulk material with the materials intrinsic loss and a surface layer with a lower intrinsic loss, ϕ_{Surf} penetrating to some depth d_{s}. Gretarson et al. estimated a value of $\phi_{\mathrm{Surf}} = 10^{-5}$, in the case of well handled fused silica fibres [80]. The value could vary wildly based on treatment, handling, and manufacturing, but is useful as an

estimate for samples which are kept in a clean environment. We use this interpretation to estimate a material loss term of

$$\phi_{\text{Mat}} = \phi_{\text{Bulk}} + \phi_{\text{Surf}} \, d_{\text{s}} \frac{\int_S \epsilon(\vec{r}) \mathrm{d}S}{\int_V \epsilon(\vec{r}) \mathrm{d}V} \tag{3.21}$$

where ϵ is the strain energy density per unit volume. To solve the integrals in this equation, finite element analysis is used, as discussed in Section 4.4.3. In the same section, I present measurements of the surface roughness of several oscillator samples used in this project. These measurements can be used to estimate values of d_{s}.

3.3.3   Clamping Losses

Any inertial sensor must in some way be fixed to the thing that is being measured. The means by which the device is fixed down is a potential source of mechanical loss, called clamping loss.

   The nature of this loss makes defining general equations describing it impossible. Each situation is unique and would have to be derived. Unlike other factors, which are accounted for in oscillator design, clamping loss must be considered when designing a mount. It is always a leading cause of suspicion when a theoretical Q factor is not matching an expected Q factor. To test whether clamping is the limit, small adjustments can be made to, for example, clamping pressure to see if the Q changes. If the Q changes with clamping conditions, it is likely that this is the limiting effect.

   Papers on the subject of clamping loss tend to be limited to discussion of very frequently used geometries [81, 82], or to specific use cases in large scale projects, such as the suspensions in the LIGO detectors [83]. While this cannot be used to determine the expected clamping loss, several general means of reducing it are discussed. Cagnoli showed that increasing the pressure and reducing the contact point both lead to an reduction in clamping loss [83]. Further research and derivations on nanometre scale resonators by Wilson Rae [84, 85] produced result which could model phonon dispersion on these scales. While not directly applicable to larger scales, the effects are similar. They again show that clamping area should be minimised, but also suggest that the density of the clamp should be much higher than the substrate. The density gradient reflects most of the phonons back into bulk oscillating material, which will have a considerably lower bulk loss. The physical interpretation of this is that the loss comes from friction heating the contact point between the clamp and the sample. This heated point radiates energy to the environment and acts as loss. A stiffer clamp reduces this motion and so the friction. The development of clamps and mounts for oscillators is discussed in Section 5.2.2.

3.4   Optical Readout of Inertial Test Masses

Due to the remarkable precision of interferometry, many state of the art inertial sensors use interferometry for readout of test mass position [61, 86, 87, 88]. These readout techniques use a wide range of tools available in interferometry. Several different techniques are used in this thesis. Two factors are requirements for the interferometer: sufficient

range to track the full motion of oscillation and sufficient noise performance to readout inertial motion to the desired level.

The first of these is relatively easy to calculate. We simply take an estimate of the acceleration noise spectral density of the use environment and use equation 3.2 to estimate the test mass displacement. The range can be a problem because of the large gain in amplitude for high Q factor oscillators. Figure 3.3 shows an example of this. Here the ground motion measured in the laboratory is shown along with the resulting motion of 3 oscillators with different resonance frequencies. The lower the resonance frequency, the greater the motion. The larger motion does make precise readout of the test mass position easier, but it leads to two problems. The large range of motion may saturate some interferometers and cause the oscillator to break due to extreme motion. As seen in Figure 3.3, an oscillator with $1\,\mathrm{Hz}$ resonance frequency and a Q of 100,000 will undergo $1\,\mathrm{mm}$ of motion on resonance and, as discussed in Section 4.3.1 designing an oscillator to survive such conditions is not trivial.

Once the test mass motion expected from an oscillator has been calculated, a relevant optical readout can be chosen (or conversely, if a readout scheme is already chosen, an oscillator can be designed to meet performance). All other design considerations being equal, the simplest readout scheme that meets the required performance would be the best. Here I discuss some standard and state-of-the-art techniques for test mass sensing with interferometry.

Homodyne Interferometery

The simplest interferometric approach uses a homodyne interferometer in a standard configuration such as a Michelson or Mach-Zehnder. Light from one source is split into two paths and then recombined. The interference pattern is then simply a result of path length differences. Typically one would then control the arm length such that the interferometer is held on the mid-fringe so that the response is linear. Often this is done mechanically with a piezo actuator and servo. The range of this control loop can be problematic when used in high-activity environments where these can saturate.

Optical Resonators and the Pound Drever Hall Technique

Two or more highly reflective mirrors can create an optical resonator. Here the optical cavity will only transmit light that matches its resonance condition [89]. The Pound Drever Hall (PDH) technique can then be used to measure length changes in the optical cavity by modulating and demodulating signals. The PDH response scales with the finesse of the optical cavity, and so a sharper response to length changes is achieved when compared to a single pass interferometer. The implementation of such a system is discussed in Chapter 6.

The system must be precisely held at its operating point to enable readout. The control can be done by mechanically changing the cavity length with piezo actuators or by frequency shifts. Frequency control typically allows for very fast control but limited dynamic range. The frequency shifts needed to compensate for the motion depend on cavity length as

$$\frac{c}{2L} = \nu_{\mathrm{FSR}}, \tag{3.22}$$

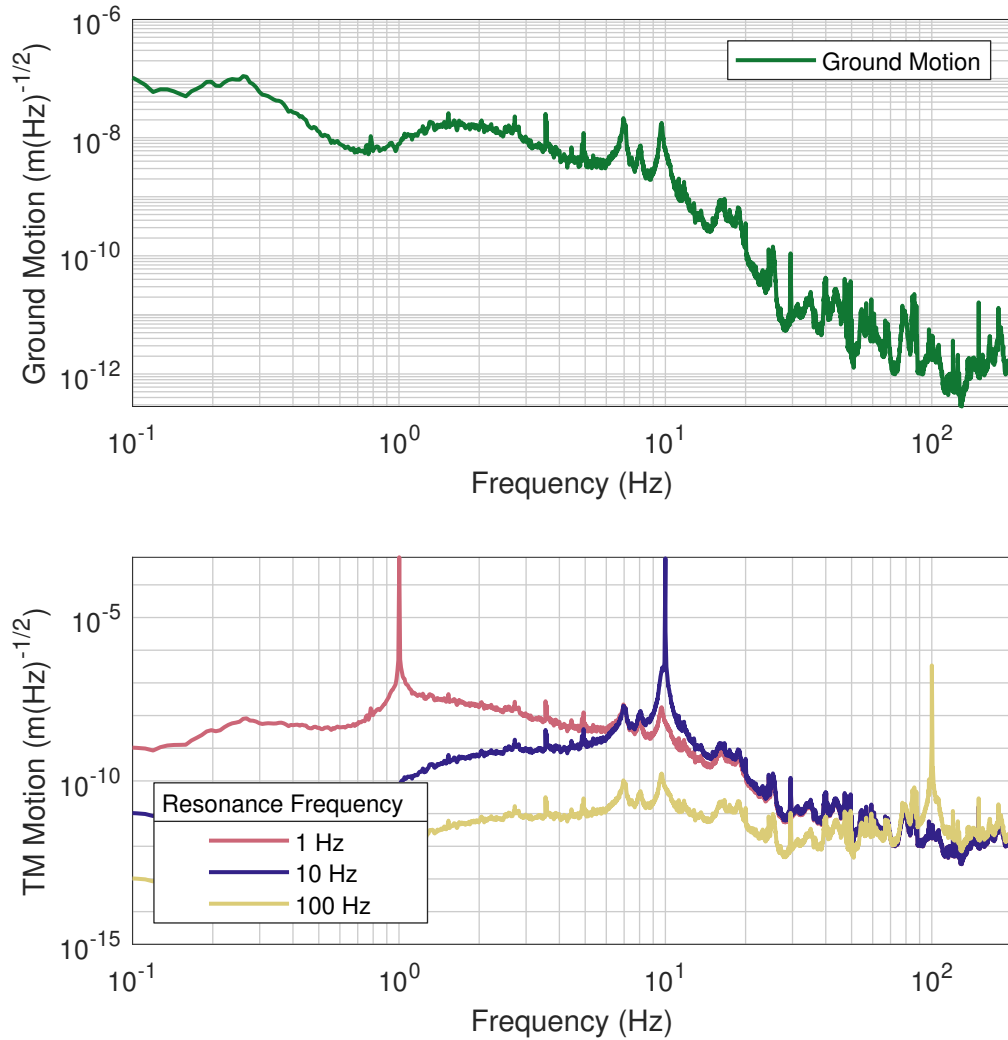Figure 3.3: (Top) Ground motion in the lab measured using an L-4C Geophone. (Bottom) The test mass motion for different frequency oscillators. Each oscillator is assumed to have a Q factor of 100,000. The lower the resonance frequency, the greater the motion. This is a double-edged sword as larger motion allows for a less sensitive readout, but it also means that the interferometer has to track and survive a more extreme stresses.

where c is the speed of light, L is the cavity length and ν_{FSR} is the distance between fringes in the cavity. Hence, a longer cavity gives a better range of control. However, as well as space considerations, longer cavities lead to a direct increase in frequency noise.

The frequency noise introduced in the readout is given by Equation 6.9, and this leaves only two options to reduce its effect: change the cavity length or reduce the incoming frequency fluctuations. The easier of these to control is cavity length, so it should be tuned to the limit of what can be controlled. If further suppression is needed, a frequency reference must be installed, which can be a rigid interferometer or an absolute time measure, such as an atomic transition. A frequency reference was needed to assist the prototype design, and its results are discussed in Section 6.2.

The combination of precision measurement and fast control but limited dynamic range makes it suitable for monitoring seismically isolated vaults when combined with a higher frequency oscillator. As these environments are the exact regions where sensors in gravitational wave observatories are used, such sensors were investigated in this thesis as the prototype devices.

Heterodyne Interferometery

Heterodyne interferometers use beams of two slightly different frequencies to generate a beat note. As the path length of the arms varies, this pattern shifts the phase of the beat note. A phasemeter can then track the beat note phase and recover the path length variations.

No control loop is needed if the modulation frequency is constant and the phase can be read out. A very large phase range can be achieved through an appropriate choice of hetrodyne frequency. The large dynamic range makes heterodyne interferometry an ideal tool for use in more active environments or when coupled with a low-frequency oscillator, as discussed in Chapter 5.5. It was also why a heterodyne interferometer was used to track the large oscillations induced during ringdown measurements discussed in Chapter 5. A common and easily available phasemeter is the Moku:Lab from Liquid Instruments. Setting up the device and using for readout is very simple. The noise performance of the Moku:Lab phasemeter was measured in the lab and is shown in Figure 3.4, where it is assumed that a 1064nm laser is used. These measurements agree well with the advertised 100 ncyc precision defined in the specifications.

Optical Heads and Compact Readout Schemes

One key advantage of a compact sensor is its size. It must be integrated with a small readout to make the full sensor in a housing. Optical fibres can inject laser beams into the sensor when using interferometric readouts. While the development of such schemes is beyond the scope of this thesis, it is worthwhile to take a quick tour of currently developed schemes and identify which ones could easily be integrated with these oscillators to enable compact, high-performance devices.

The Homodyne Quadrature Inteferometer (HoQI) is a recent development by the University of Birmingham [86]. The scheme uses two co-propagating beams with orthogonal polarisation states. One beam is sent to the test mass, and one to a reference mirror. The device reaches a noise floor of $10^{-13}\,\text{m(Hz)}^{-1/2}$ from 0.1-10 Hz. This device fits in a housing 17 cm by 10 cm, making it more of a compact interferometer

than an optical head. The setup has been used on a GS13 sensor to meet its thermal noise level. Its noise performance is shown in Figure 3.4. As the device does not require control electronics, the operating range should be large, though no specific values have been given.

Deep Frequency Modulation (DFM) is a compact interferometric technique developed at the Max Planck Institute for Gravitational Physics (Albert Einstein Institute (AEI)[90]. The technique imprints a signal on a laser. The laser beam is then split up and recombined with the path length changes modulated on this signal. The signal can then be reconstructed to give both absolute and relative phase difference measurements. The technique has the advantage that the complexity is almost entirely confined to the modulation and demodulation scheme. The simple optical requirements allowed for a demonstration of a prototype using just one optical prism [91]. Recent performance has been measured at a sub-picometer level [92] and devices developed at the AEI) are expected to also meet this performance level [93]. The measured noise performance of [91] is also shown in Figure 3.4. The technique is designed to track test masses moving over one fringe. Its large dynamic range makes the technique an ideal candidate for integration with a low frequency, high Q factor oscillator.

Previous optomechanical inertial sensors integrated a low finesse fibre cavity readout [90]. The scheme was based on a demonstration by Smith et al. [94] and achieved impressive noise performance for a very simple design, again shown in Figure 3.4. The limits of this technique are its operational range, which is approximately 25 nm. This limit is only suitable for oscillators with kHz resonances or those operating in exceptionally quiet seismic vaults.

Other groups working on similar devices are currently exploring heterodyne fibre readouts. The latest attempts use a dual-wavelength interferometer [95]. As optical fibres introduce a lot of extra phase noise from various thermal and polarisation effects, they attempted to use one wavelength as a reference and one as a test interferometer. The heterodyne interferometer then allows for a very large dynamic range. The results of their method are also shown in Figure 3.4.

Discussed here are just a small number of the many compact optical readout techniques. One key observation from Figure 3.4 is that many techniques get limited at the $1 \times 10^{-13}\text{m(Hz)}^{-1/2}$ level. While this is no fundamental limitation, this level can be reliably reached on this scale with standard optical equipment. Noise floors lower than this can be reached but require dedicated work and research efforts. These "rules of thumb" are worth considering when choosing resonance frequencies.
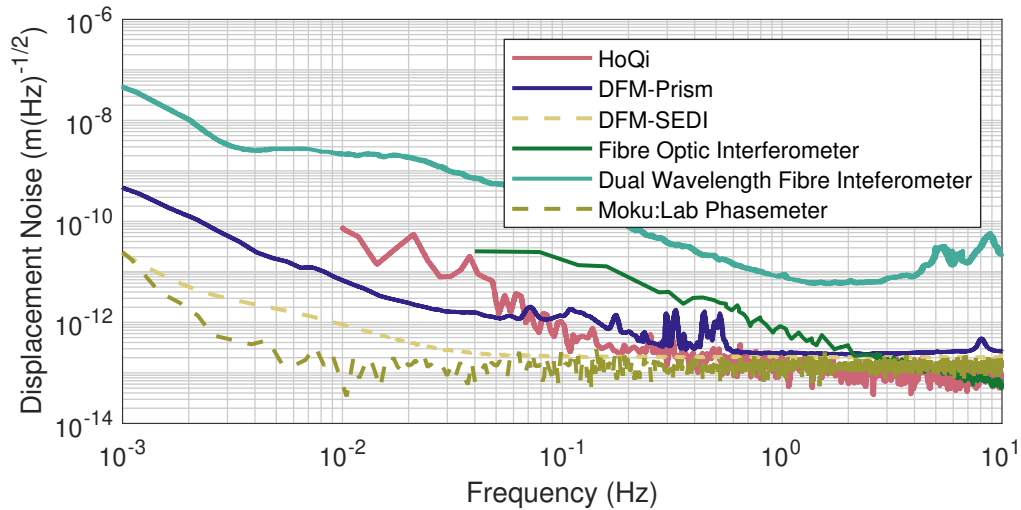
Figure 3.4: A comparison of the noise performance of several compact displacement sensors with sufficiently large dynamic ranges. The data from the HoQi is taken from [86], for the DFM-Prism from [91], DFM-SEDI from [93], the Fibre Optic Interferometer from [94], The Dual Wavelength Fibre Interferometer from [95]. The Moku:Lab Noise floor was measured in our lab under optimal instrument settings and does not show an actual displacement sensor, but rather what an ideal optical heterodyne interferometer with this readout could achieve. The DFM-SEDI line is a theoretical projection, and so it is dashed. All of these readout techniques lie an order of magnitude above what is needed for a precision readout of a 50 Hz resonator, which is the focus of Chapter 6, and so warranted the installation of an optical resonator readout. Many of these techniques achieve sufficient performance to match some of the state-of-the-art inertial sensors available today when combined with a 15 Hz oscillator, like one envisioned in Section 5.5.

Chapter 4

Designing Inertial Sensors

This chapter approaches the field of designing compact inertial sensors with a broad perspective. First, I present the design space of interest for inertial sensors in gravitational wave detectors. We use this to define what sensors we wish to make. After this, I broadened the scope to look at how each parameter in the design can be tuned to achieve the desired mechanical properties. A general approach is taken on a simplified geometry so that the effects of design parameters on the system are understood. Both Finite Element Analysis (FEA) and analytic modelling are used in this analysis. The justification for parameters chosen in the design of inertial sensors throughout this project follows naturally form this analysis. Although the primary aim here is the design of inertial sensors for gravitational wave astronomy, the studies can easily be adapted for other fields or applications.
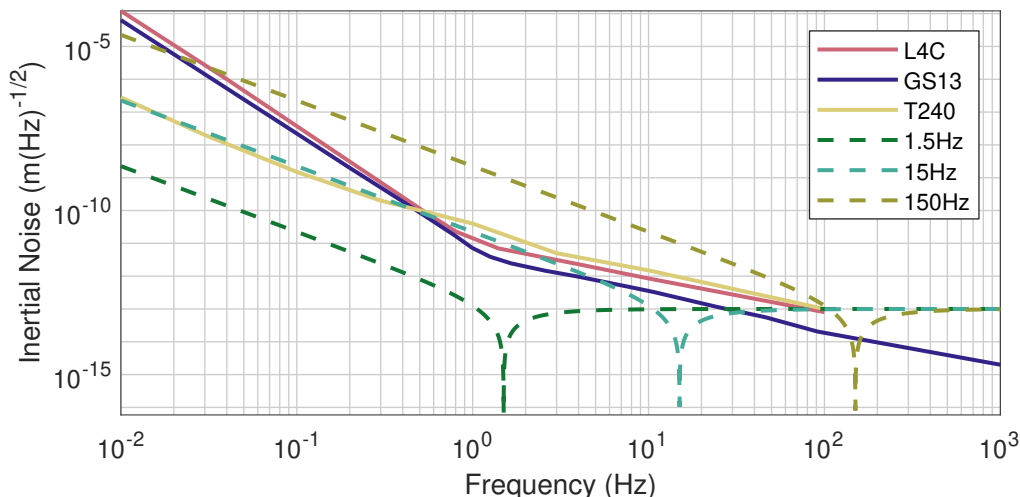
Figure 4.1: A comparison of the inertial noise floors of oscillators with differing resonance when a Moku:Lab Phasemeter (or a 16 bit Control and Data System (CDS) phasemeter) dark noise is the limiting noise source. We see that to achieve comparable noise performance to other high-performance sensors, we would need a resonance frequency below 15 Hz. (Acceleration Equivalent (D.3))

4.1   Designing High Precision, Compact Inertial Sensors

Now that we have covered the basic principle of inertial sensors, we ask: what sensors are useful for gravitational wave astronomy? For the sensors to play a part in contributing to the control scheme, they must achieve a sensitivity close to the high-precision sensors already installed there, shown in Figure 1.3.

There are two ways that better sensors can improve the performance of the seismic isolation system: sensors that have significantly higher performance than the current ones or achieve similar performance in a compact form. Higher sensitivity seismometers are an approach best achieved by bulkier sensors. While the high Q factor sensors can achieve this, the size and scale of the oscillators would make manufacturing difficult. We, therefore, focus here on achieving high-performance sensors in a compact form. There are many opportunities for such sensors in the LIGO isolation platform schemes. Particular examples are that there is no inertial sensing in the lower suspension stages [15, 16, 17], largely because of space requirements.

Furthermore, a ring of compact vertical sensors around the platform would allow a measure of inertial pitch and yaw on the platforms, which could contribute to better tilt isolation. Currently tilt sensors are installed on the ground next to the platform [35]. However, an on-platform tilt measure is a critically missing piece of information at the moment [36].

We derive the design principles for inertial sensors designed and tested in this work from these requirements. The key features that must be shown then are that performance comparable to a sensor such as a L4C or GS13 seismometer can be achieved, that this can be done on a centimetre scale sensor, and that the sensors can work in any orientation and survive handling. The sensor head (the part that must be deployed in

the region we wish to study) can be separated from its readout using an optical fibre, but the sensor head must be compact. We define this as having no dimension that exceeds 10 cm, and, ideally, being smaller in every direction. Working sensors must be shown before any specific one can be deployed on-site. So, the designs here are proof of principle sensors, meant to show the capability of the technologies and not necessarily fit in a specific role.

4.1.1   Design Goals

We now need to define the parameters needed to achieve this goal. We start by considering the readout methods discussed in Section 3.4, and notice most single pass interferometer instruments have noise floors around the 10^{-13}m(Hz)$^{-1/2}$ level. We can then use this to find the resonance frequency of the oscillator needed with one of these schemes to meet our noise requirements. The oscillator's transfer function derived in Equation 2.9 can convert a displacement noise floor into inertial equivalent noise. The projection of sensors with a displacement noise floor limited by a Moku:Lab in inertial noise is shown in Figure 4.1. Various resonance frequencies of oscillators are compared. From the figure, we immediately see that the frequency region of 15 Hz and below is necessary to achieve comparable performance with a GS13.

We were now presented with a dilemma. As we discover in Section 4.3.1, we can already see that approximately 40 mm long flexures would be needed. At the start of the project, this length was an order of magnitude longer than what had been produced, using the subtractive manufacturing techniques used to make previous inertial sensors [60, 63], and even now such long flexures are prohibitively expensive for iterative design and testing. The difficulty in producing such lengths with subtractive methods is somewhat fixed, as the design cost is linearly related to the amount of material that must be removed. Therefore, we investigate alternative manufacturing methods in Section 5.5. Rushing straight for this goal, however, would be unwise. We instead developed a road map to this endpoint. The first goal was to show that these size oscillators could achieve high Q factor oscillations. We choose to make these resonance frequencies above 100 Hz to examine and gain experience with our handling and manufacturing steps for things that degrade Q factors or damage flexures before moving to lower frequency (and more expensive) designs. When we have oscillators that behave as expected, we can move onto an optimised design, which reaches useful noise performance to the gravitational wave community but has complicated readout tools. In parallel, we develop techniques to scale up the design to make low-frequency oscillators, which could be used with compact readout schemes. The processes of iterative design and testing of drum oscillators to achieve each of these goals are followed throughout Chapter 5.

4.1.2   Design Geometries

One factor decided very early on was folding the flexures around into a circularly symmetric drum geometry, shown in Figure 4.2. The folding allowed for multiple flexure parallelograms on one oscillator. We decided to fold the flexures into the complicated drum geometry for two reasons. The first was pushing higher-order modes to significantly higher frequencies, as the springboard and torsional modes would otherwise
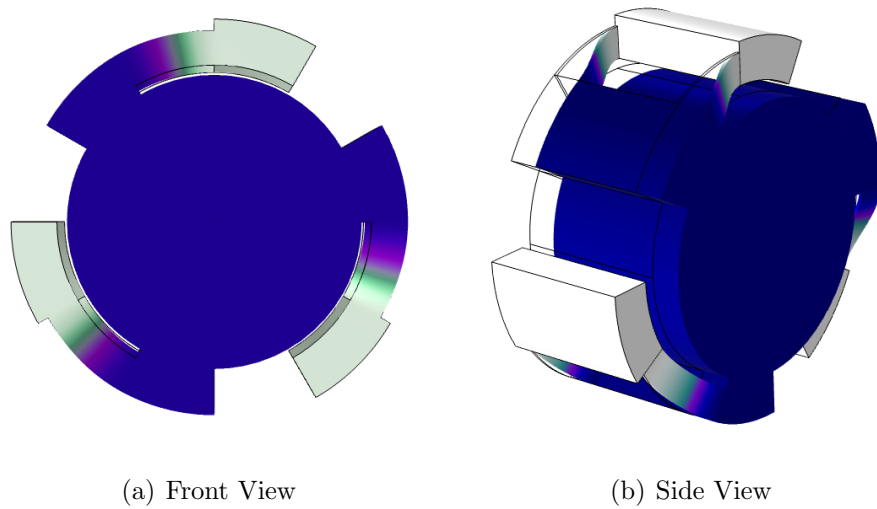
(a) Front View   (b) Side View

Figure 4.2: The internal structure of a "3D Drum" oscillator structure with the frame suppressed. The locations of the flexure pairs can be seen. The geometry allows an arbitrary scaling of the test mass without concerns of the springboard mode of oscillation. As with the parallelogram oscillator, the motion is entirely in a linear direction.

Three sets of 100µm thick parallelogram flexures to support test mass

Silicate Bond to hold assemble together

3 gram test mass coated with highly reflective coating

5cm long ULE spacer to form high finesse optical resonator

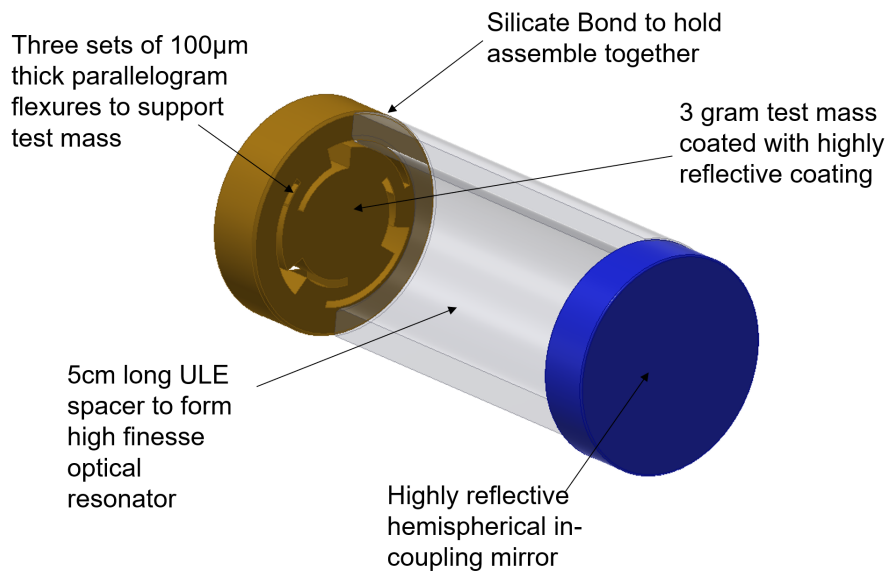Highly reflective hemispherical in-coupling mirror

Figure 4.3: The overall design of the sensor. The part consists of 3 pieces of glass, each of which has a different colour in the illustration. In orange is the oscillator that encodes the movement of the ground. The clear part is the rigid spacer. In blue is the in-coupling mirror; this can be highly reflective to allow for a high-finesse cavity for precise readout.

be problematic. Greater mode separation allowed for the support of much larger test masses in a more compact form, substantially reducing thermal noise. By having a large disk as a test mass, the distance between flexures and worked regions can be increased, minimising the chances of damage. The worked regions in this design would be the outer frame, which was to be bonded to the spacer, and the area to be coated in the middle. Finally, making the oscillators 1" diameter optics will fit in most standard mounts and tools for optics, making handling much easier.

The oscillators were typically etched from wafers several mm thick. As discussed in Section 4.4, the flexure separation (effectively wafer thickness here) can be cancelled out in the optimum case, where we can make longer flexure, lower resonant frequency oscillators to compensate. However, most oscillators in the project were designed with specific fundamental resonate frequencies, f_0 in mind. Here, stiffer flexures and large masses would lead to an improvement in oscillator performance. The deepest depth that FEMTOprint could etch with their technology was 7 mm, so typically we aimed to have designs this thick. As this project was ongoing against the backdrop of the chaos of the global COVID19 pandemic, often glass wafers of specific thicknesses were not in stock, and lead times tended to be measured in years. Therefore, we took the approach of just using what was available. Doing so often meant using the standard wafer thickness of 6.35 mm. These thicknesses were then compensated for in the design.

For oscillators with a natural frequency above 15 Hz, we would require state-of-the-art readout techniques and high Q factors to match the performance of classical sensors. A readout noise level of order $\mathrm{fm(Hz)}^{-1/2}$ would be required. This level performs orders of magnitude better than most of the techniques listed in Section 3.4, so a different approach was needed. High-finesse optical resonators [1] respond sharply to length changes, allowing for excellent noise performance. In order to keep the resonator stable and aligned, a spacer can be used. Making the spacer out of a Ultra Low Expansion (ULE) glass will help keep the cavity resonant with light fields even in the presence of thermal drifts. A high reflectivity on the oscillator would be needed to make a high-finesse optical resonator. Previous experiments had difficulties attaching mirrors to the samples [60]. We could use standard optical coating techniques by making the sample out of a 1-inch diameter round wafer. Hence, these samples would be coated with an Ion Beam Sputtering (IBS) coating and the effect on the Q factor measured. These considerations produced the overall design sketch shown in Figure 4.3. The optical readout method was developed in parallel with the oscillators and informed design decisions for oscillators throughout Chapter 5. The development of the readout scheme is discussed in Chapter 6.

4.2   Manufacturing Gram Scale Oscillators

In order to narrow down the infinite possibility space of inertial sensor design to a region worth studying, we must look at what can be manufactured. As we aim to use high Q factor to achieve good performance in a compact form, we need materials

---

[1] As we now have both mechanical and optical aspects that are resonating, descriptions can get confusing very quickly. I will establish a convention throughout the thesis of using oscillation to describe a mechanical movement and resonance to describe a light field inside an optical cavity.

that show a low bulk loss. Materials often used for high Q factor microresonators are fused silica [96, 97, 98] and silicon [54, 99]. Often, these layers are mixed or doped with other materials to achieve other specific properties. Most of the work on upscaling MEMs style resonators has used fused silica oscillators [60, 61, 100, 101] over silicon because of its comparatively higher shear strength. However, some are silicon inertial sensors [102] for its capabilities at cryogenic temperatures. We focus on fused silica designs throughout this thesis for a few reasons, partly because of the shear strength, partly because we knew designs could achieve high Q factor, and partly because it was a material with which companies were willing to work.

Glass structures like those studied here can be produced additively [103], where the material is added in the required shape, or subtractively, where the material is removed from an initial wafer, leaving only the desired geometry. The former very quickly becomes impractical at this scale. The method used for most oscillators produced in this study was with the techniques developed by the company FEMTOprint [63], which is the method used by other groups making similar sensors [60, 101]. FEMTOprint uses a two-step process where a femtosecond pulsed laser targets the material to be removed. The targeted areas are much more reactive, so a wet chemical HF acid is used to etch out the regions, leaving the remaining structure. This process has been shown to produce high Q factor oscillators but does have some limitations. First, the surface roughness left by the method is not smooth, measured in Section 4.4.1, which presents a limit on Q factor from surface losses. The second is that the laser's penetration depth defines the scale of one dimension that can be etched. The largest depth we could get cut was 7 mm deep. Finally, because the main cost is given effectively in running the laser, there is a limit to what is financially feasible regarding material removal. This scaling meant that oscillators become increasingly unfeasible when pushing resonance to a lower frequency via this production method. For this reason, we look at a method of bonding various subassemblies together to make larger geometries in section 5.5.

4.3   Designing Low Frequency Oscillators

The designs we present are of a complex geometry, but we still wish to get a handle on how the system behaves according to our design parameters. Therefore, we wish to use a simplified geometry to study the tools we have at our disposal and good design techniques. We do this with a parallelogram linear oscillator as shown in Figure 1.4 a). Each parallelogram junction acts similarly in the more complex geometries, although the difference must be analysed for those cases. While the effect of mismatches is discussed in Section 4.5, the assumption is that a symmetrical design is desired, i.e., all flexures are identical by design. The geometry studied and the definition of each parameter we will use throughout this chapter are shown in Figure 4.4.

4.3.1   Low Frequency Oscillation

From Chapter 2.1, a need for low-frequency oscillations has been demonstrated. Equation 2.5 gives an oscillator's natural frequency, which leaves two choices for lowering natural frequencies: reducing the spring constant, k and increasing the suspended Test Mass. If the material is fixed in choice, this only leaves geometric considerations, which
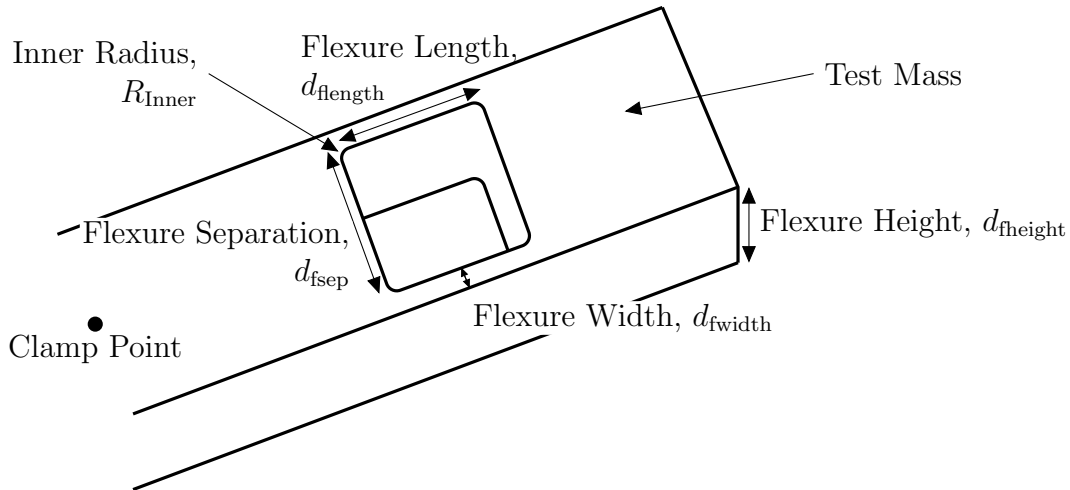
Figure 4.4: The geometric definition of the terms discussed in flexure design. Although the overall geometry of the device may change, the definitions of flexure length, separation, width and height stay the same. In this case one flexure junction is shown, in others, multiple flexure junctions may be used. The test mass is the suspended mass that is not rigidly attached to the ground, it is always on the opposite side of the flexure junction to the clamping point. The clamp point is the point (or interface) at which the monolithic piece is rigidly attached to the ground.

means reducing d_{fwidth} and d_{fheight}, while increasing d_{flength} and increasing the suspended mass. Any of these changes make samples more prone to brittle fracture. We must, therefore, study how far we can push the designs and still have them survive handling.

A simple failure model can be used if motion and stress are confined or largely dominated by one direction. In this case, the stress induced by a given load can be simulated. Points in the geometry of high stress can be compared to the failure condition. The maximum shear stress criterion can be used to predict failure in oscillators undergoing parallel motion, where most stress is confined to the direction of travel. We can then define a safety factor for all points in the sample

$$\frac{F_{\mathrm{U}}}{\sigma} = \text{Safety Factor}, \qquad (4.1)$$

where F_{U} is the ultimate shear stress, and σ is the shear stress at the given point in the oscillator. A safety factor of less than 1 indicates the material will fail at that point; however, it is advisable to always leave some margin of error for extra unexpected loads. Furthermore, ultimate failure points in glass typically have large uncertainties as they strongly depend on local structure defects and contaminants. The highest load we expect the oscillator to face is when it is parallel to gravity. For this reason, along with the relatively high cost of making a sample, a large safety factor of at least five whilst under $1\,\mathrm{g}$ of load was used for all designs presented here. In addition, due to fused silica being brittle and, hence, having a linear response until the point of failure, a large safety margin allows the oscillators to be used as vertical seismometers. This behaviour also means that three sensors can be mounted orthogonally to create a 3D seismometer with an arbitrary set of operating axes.
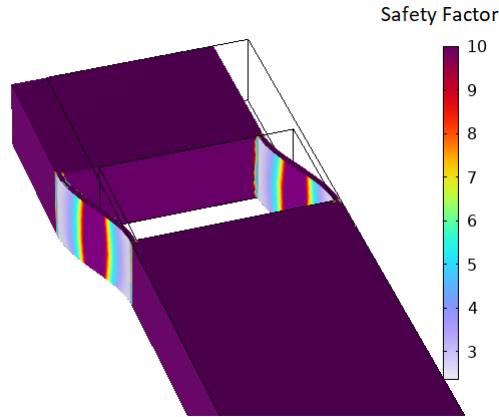
Figure 4.5: A heat map of the safety factor across various points in an example sample. The modelling of stress is done using a stationary solver in COMSOL Multiphysics. The oscillator is fixed at its clamping point, while the test mass is moved by a prescribed displacement of 25 μm.

FEA can be used to estimate the stresses across the material; an example is shown in Figure 4.5. As the stress is almost entirely in one direction, the shear strength of fused silica is used as F_{U} during 1mm of parallelogram motion. In this thesis, COMSOL Multiphysics was used to solve the stress distribution. The solid mechanics and heat transfer in solids packages were used to study the stress effects under prescribed loads. The smallest possible mesh sizes were used. Whenever a simulation is conducted, it is important to check that it has converged on a solution. The PARDISO solver was used to calculate solutions. For most simple geometries, the solution converged to a relative tolerance between iterations of 10^{-16} after ten iterations. Some of the more complex geometries discussed in 5 required closer to 100 iterations to achieve this level of convergence.

Sag, Gravity, and Low-Frequency Oscillators

Any oscillating system not in free fall will undergo sag when its direction of oscillation is parallel to that of gravity. The magnitude of this displacement d_{sag} is given by the simple equation

$$d_{\mathrm{sag}} = \frac{g}{\omega_{\mathrm{vert}}^2}, \tag{4.2}$$

where g is the local gravity, and ω_{vert} is the natural frequency in the vertical direction. As the only parameter affecting this is the resonance frequency, which scales as the inverse of the square, this becomes a serious issue for low-frequency oscillators. This scaling warranted considerable attention for the design of the next generation of devices with resonances of 15 Hz, discussed in Section 5.5.

There is only one term in the equation governing displacement, so there are limited means of avoiding it. Vertical inertial sensors often use prestrained suspensions, such that when in a vertical position, the test mass' weight pulls the piece into its desired operating condition [104]. Horizontal sensors can not use such a scheme for standard use. In principle, assuming that they are sufficiently vertically stiff, this would not

affect the performance of the horizontal sensor. However, practical use in the lab and manufacturing concerns means that the device must survive in any orientation. Here, we have two options: have the device survive naturally or use motion stops to limit the range of motion the device can undergo. The former is the simpler solution when complex multistage manufacturing processes are required. Motion stops also prevent the device from being used as a vertical inertial sensor. This project focused on oscillators that could survive under this sag.

If motion stops are to be used, their spacing must be chosen carefully. There must be at least enough distance for the oscillator to accommodate its dynamic range, discussed in Section 3.4, while being narrow enough to stop at its breaking point. The motion stops would then have to be incorporated so that the load is exceeded at no point in the device's manufacturing until installed. With the motion stops installed, the device can be modelled with a resonance frequency corresponding to the sag allowed by the motion stops when testing breaking points. This changes the scaling and trade-offs discussed in Section 4.4.2.

Several parameters can be tuned to increase the survivability of the oscillators. One that may at first seem counter-intuitive is decreasing the flexure width. Naturally, one thinks of making the support thicker, and while this would lower the stress, it also increases the natural frequency. If one increases the load to compensate for this, then the test mass' displacement remains the same, so the flexure must bend more while being stiffer, increasing material stress. Therefore, thinning the flexure and supporting less weight will increase survivability. Practical considerations limit how thin the flexures can be tuned, so we must look to other terms to improve survivability. The effects of flexure width on maximal beam stress are shown in Figure 4.6.

Alternatively, the flexure length can be increased. Doing so reduces the curvature per unit length and the stresses induced. A study was conducted using COMSOL multiphysics on a simple beam geometry undergoing prescribed displacement of one end of the beam under flexural motion. Figure 4.7 shows the result. Varying the beam's length, we can see the point at which a safety factor of five is reached. The Figure gives a good estimate of the flexure required for a specific resonance frequency. A clear indication here is that, in this regime, the length requirement increases inversely proportional to the flexure length.

For flexures of universal height much greater than their width, the distribution of stress is even across the whole height, ignoring boundary conditions at the edge of the geometry. As this stress is, therefore, only dependent on the curvature (and, in turn, the natural frequency), it would initially seem that height is a free parameter that can be tuned to achieve the right mass. However, Figure 4.5 shows that the stress is not evenly distributed across the flexure length. By varying the height of the flexure along the length, the stress can be more homogeneously distributed and reduce the critical stress at the peak points.

Figure 4.5 also shows in this mode of oscillation that stress is localised towards the end of the flexures. Therefore, we would like to make the ends of the flexure thicker and taller, and the centre thinner. This can be achieved by adding an inner radius to the flexures.

The effects of inner radius on stress distribution were studied using FEA. The geometry shown in Figure 4.8 was used for this test. A single flexure was tested, where
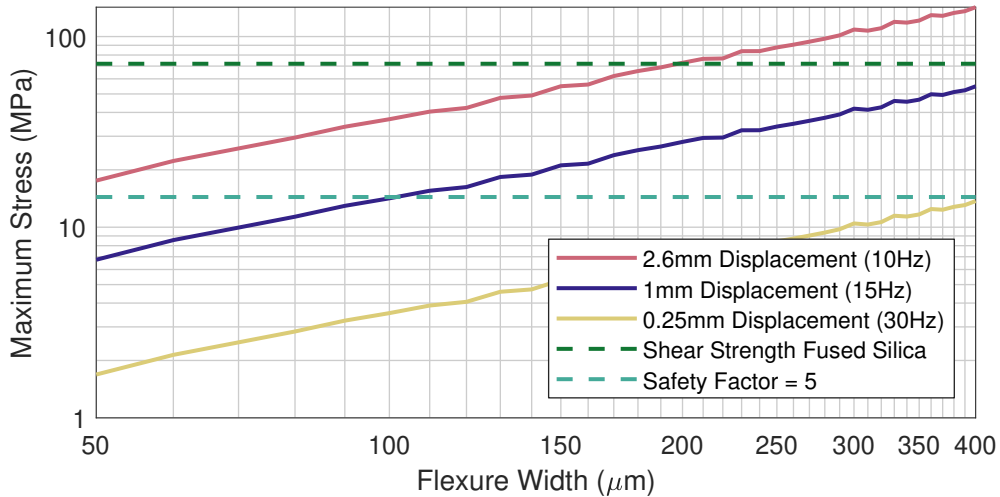
Figure 4.6: The point of maximal stress in a 40 mm thick and 8 mm high fused silica beam when undergoing flexural motion in the direction of the width, as a function of beam width. The stress was simulated using COMSOL multiphysics. The solid lines correspond to different fixed displacements, which correspond to the sag related to an oscillator of different natural frequencies under 1 g load, shown in brackets. The dashed lines correspond to the shear strength of fused silica, and this divided by five; any oscillator under this line has a safety factor of five when tilted against gravity.
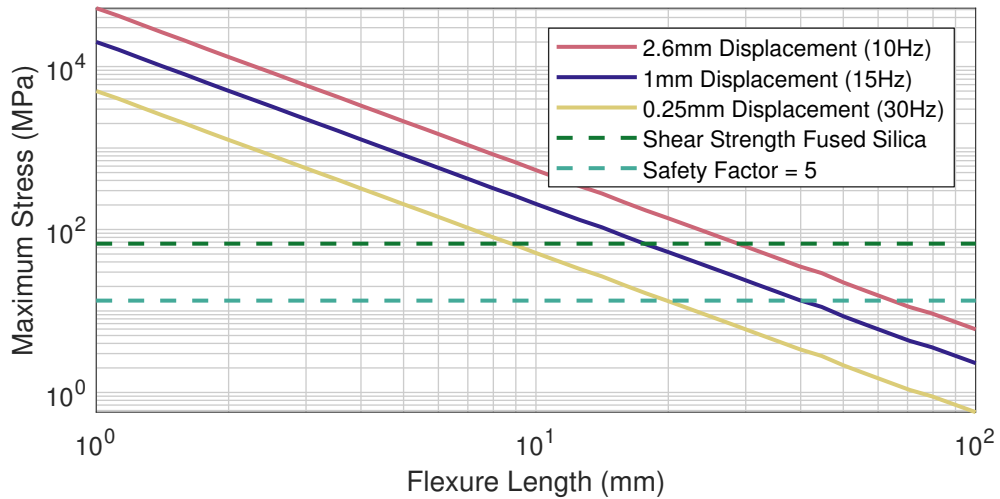

Figure 4.7: The point of maximal stress in a 100 μm thick, and 8 mm high fused silica beam when undergoing flexural motion in the direction of the width, as a function of beam length. The stress was simulated using COMSOL multiphysics. The solid lines correspond to different fixed displacements, which correspond to the sag related to an oscillator of different natural frequencies under 1 g load, shown in brackets. The dashed lines correspond to the shear strength of fused silica, and this divided by five; any oscillator under this line has a safety factor of five when tilted against gravity.
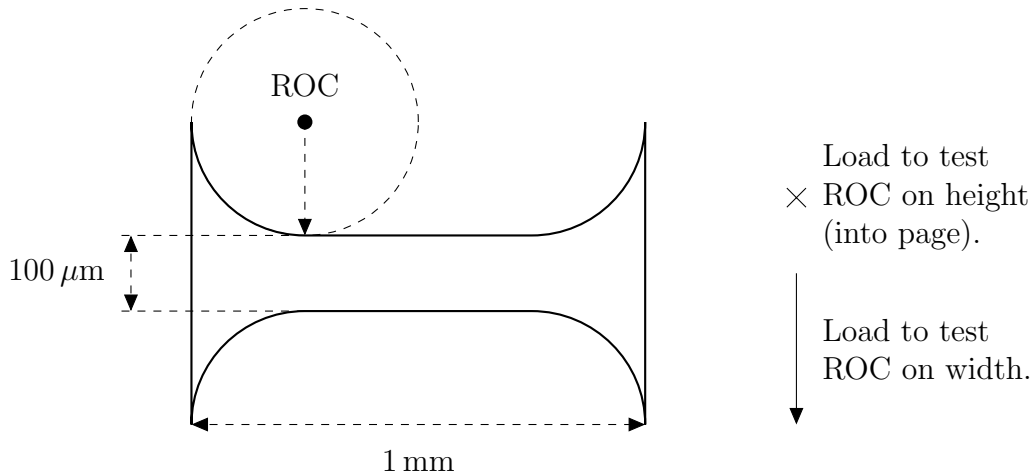
35

Figure 4.8: A diagram showing the flexure geometry used for the radius of curvature test. The design is $100\,\mu$m deep into the page. A load can be applied normal to the geometry to study the effects of flexure height. To simulate a radius of curvature in the width the load can applied downwards as shown by the arrow.

a prescribed displacement of $25\,\mu$m was applied to one end and the other fixed in place, with no motion. The flexure was 1 mm long and had a minimum width and height of $100\,\mu$m. Then, the maximum stress along the flexures for different radii of curvatures was compared.

Figure 4.9 shows the results of this simulation using COMSOL multiphysics. A radius of curvature was applied in two directions. When the radius of curvature was applied to height, it was found that a larger radius distributes the stress better over the surface. Therefore, creating an hourglass reduces the maximum stress sustained in the flexure. Changing the width along the length of the flexure creates a minimum stress with a relatively small radius of curvature, which then increases again. Both these results have a simple physical interpretation. The width response is result of a better stress distribution conflicting with thicker parts having to bend more. There is not effect of more material bending form extra height radius of curvature, so we only see the gains from better distribution. The small rise in stress as the flexure inner radius approaches half the flexure length is likely a result of stress from both ends of the flexure being pushed into the centre of the flexure. The inner radius on the height should therefore be designed just slightly short of the flexure length, when manufacturing methods allow. These results are combined with a study on the effects on Q factor in Section 4.4.3 to define what inner radii for flexure width should be used.

4.3.2   Single Mode Operation

As discussed in Section 2.1, the sensor's bandwidth is limited by the higher-order modes coupling into the oscillator's motion. Therefore, these modes must be separated from the fundamental mode by as much margin as possible. In a geometry such as Figure 4.4, two modes are of concern, which we will call the springboard motion and the torsional motion, shown in Figure 4.10.

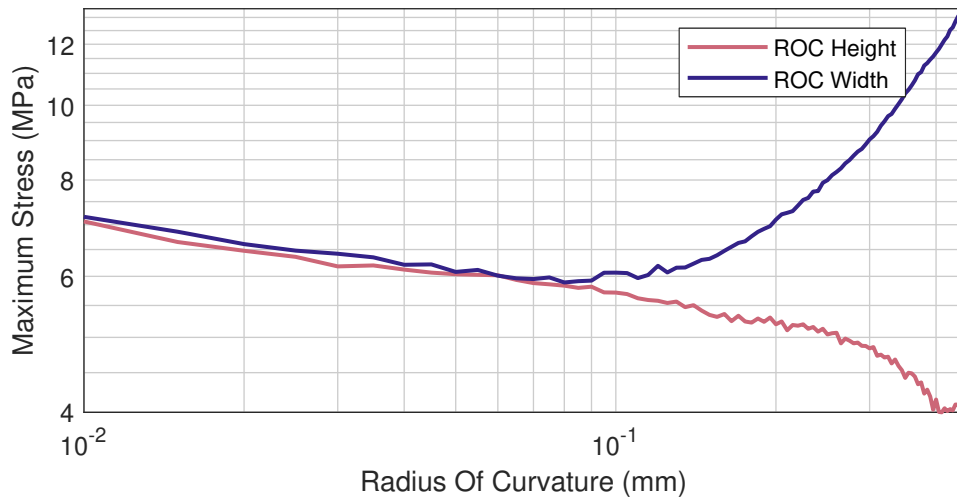The effect of varying each of the flexure parameters with linear geometry on the

Figure 4.9: The Radius Of Curvature (ROC) vs. the maximal stress in a 1 mm long sample undergoing flexural motion with a maximal end displacement of 25 µm. The radius of curvature was applied in two directions, one in the direction of flexure width and one in the direction of flexure height. The minimum flexure width and height was 100 µm when the curvature was applied in that direction. When the curvature was not applied in that direction, they were flat at 100 µm. The simulation used COMSOL's stationary solver with its solid mechanics module.

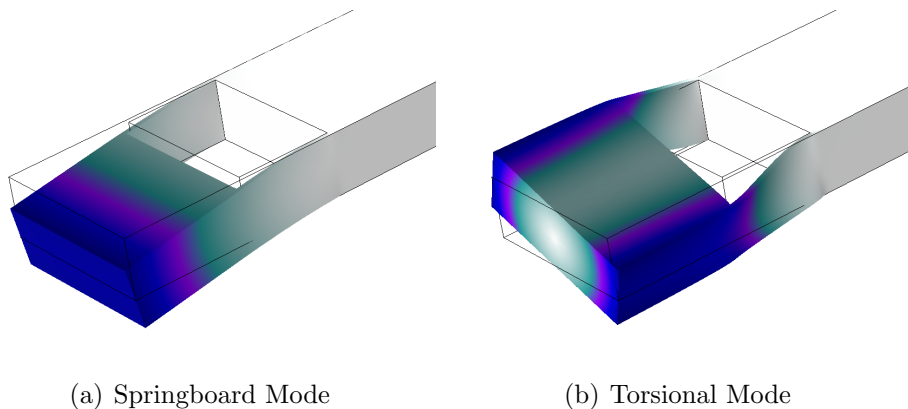(a) Springboard Mode   (b) Torsional Mode

Figure 4.10: The two modes of concern for a linear parallelogram oscillator. The springboard motion acts up-down, perpendicular to the direction of fundamental oscillation, but acts as an additional energy loss from the system. The Torsional mode will couple directly to measurements of test mass displacement and so must be suppressed using design. The heat map shows the local displacement at the extrema of the oscillation with arbitrary scaling.

eigenfrequencies of the first three modes of oscillation is shown in Figure 4.11.

The flexure width is already revealed as one of the most dominant factors. It scales with the springboard and torsional modes as power 1/2. This scaling is consistent with a linear relation between spring constant and flexure width to these modes, as expected by adding springs in parallel. Below 20 µm width, internal modes on the flexure become dominating in this geometry. However, the response of the fundamental mode to flexure width scales as power 3/2. In the case of an isotropic, homogeneous plate, the flexural rigidity, the force required to induce one unit of curvature (making it an equivalent of a bending spring constant), is given by

$$D = \frac{EH^3}{12\left(1 - \nu^2\right)}, \tag{4.3}$$

where E is the Young's Modulus, ν is the Poisson ratio, and H is the plate width, which follows a cubic relation [105, chpt 3]. It stands to reason that similar factors and relations should hold true for the geometry of these oscillators, even though some of the assumptions used in this plate theory are not appropriate; namely a homogeneous distribution of material and aspects of the curvature [106]. From Figure 4.11, we can see that increasing the material's width increases the resonance eigenfrequency of a mode parallel to that increase by a power of 3/2 and modes perpendicular to that by power 1/2.

The fundamental mode falls as power 3/2 to flexure length. The other modes do not change as a standard power law, likely because the ratio between the flexure region and test mass region is changing, each contributing differently to the mode's eigenfrequency. This changing ratio changes the dynamics of the response and so its gradient.

The flat response of the first mode to flexure height can be understood simply as adding springs in parallel, leading to a linear increase in spring constant being cancelled out by a linear increase in test mass. For similar reasons to the response of the fundamental mode to flexure width, the response of the springboard motion is to the power 3/2. The response of the torsional mode is linear for low flexure heights. However, it peaks at about 3.5mm, suggesting it becomes maximally stiff to rotation from the added material at some point. Then, the mass increase brings the mode's eigenfrequency down in frequency again.

Both the fundamental and springboard mode frequencies fall as power 1/2 to flexure separation, consistent with the linear increase in mass. The torsional mode follows a similar shape to flexure height for similar reasons.

The fundamental mode falls as power 1/2 with respect to the test mass length, showing that the linear increase in mass is the only relevant factor. The other modes, again, do not follow a straight power law; as with flexure length, the changing ratio between flexure and test mass length is the reason for this. The springboard motion tends towards a quadratic fall-off. This is assumed to be due to 1/2 power from mass increase and 3/2 power from the length of oscillation motion, as seen in the flexure length plot. The torsional mode falls as power 1/2 from linear mass increases at high test mass lengths until a drastic change at 18mm, where the axis of rotation of this mode changes to be parallel to flexure height.

Another problem one must be aware of is the internal modes in the flexures themselves. When too thin or long, the test mass and clamp side act as anchors and the
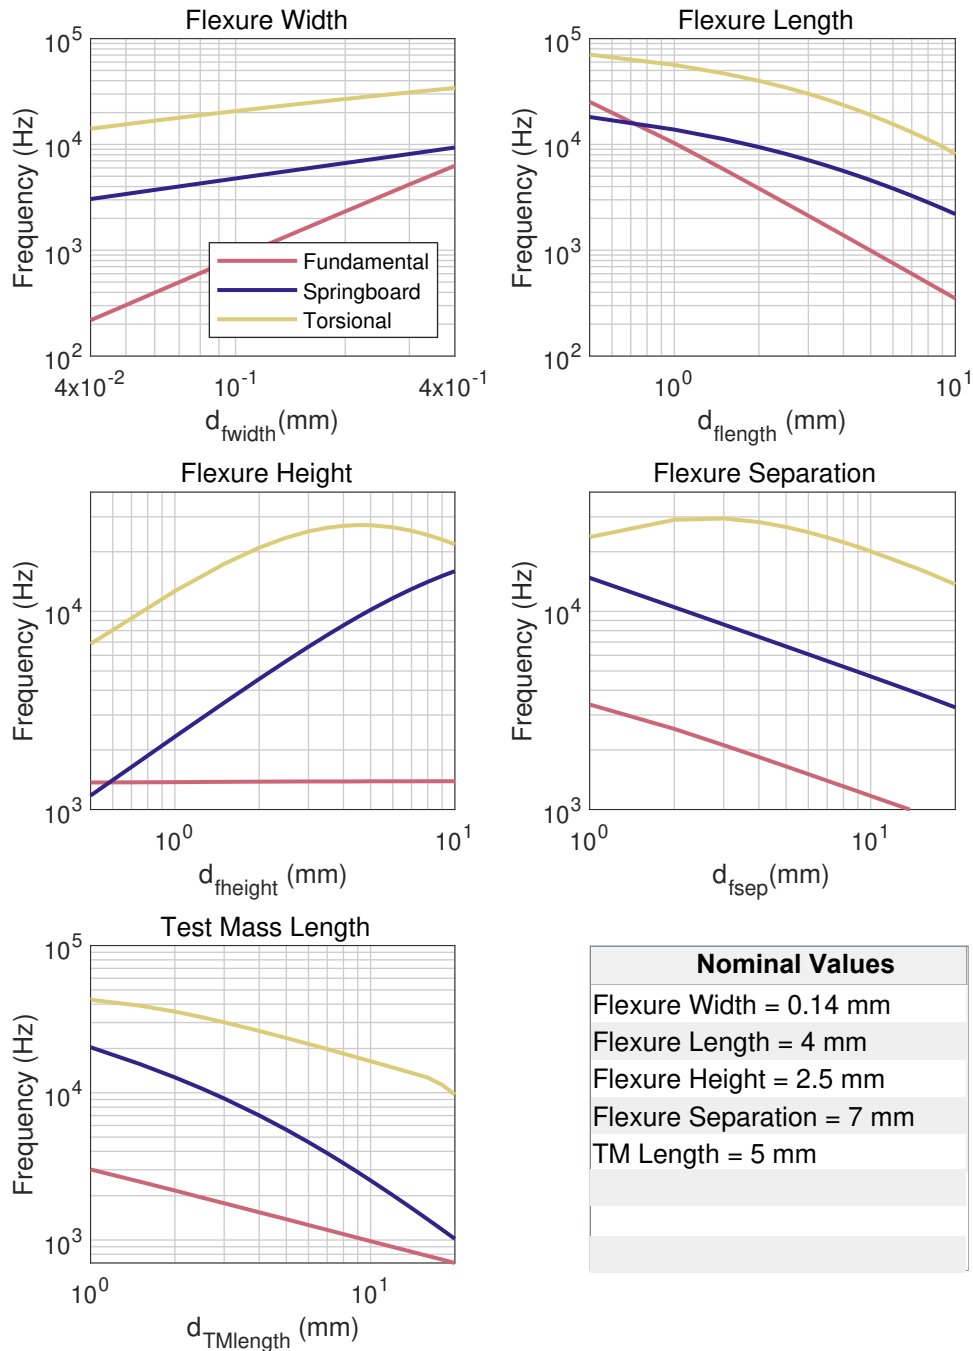
Figure 4.11: Graphs showing how the eigenfrequencies of a linear parallelogram oscillator's first three modes of oscillation vary with changing parameters from Figure 4.4. The table of nominal values shows the parameter's value when it is not under investigation. The eigenfrequencies were predicted using COMSOL Multiphysics FEA.

system becomes akin to a system where a thin beam is clamped on both ends. This behaviour is especially problematic as the motion of the fundamental mode will directly couple to these modes and lead to a large dissipation of energy from the system.

Flexure height and separation are also limited by manufacturing ability. The etching tools used to make these oscillators typically have maximum cutting depths of order millimetre; the maximum that any company would offer to cut in terms of depth was 7 mm, and the largest samples that could be fit were 4" wafers. While these companies' technology progresses quickly, these limits remain in the same ballpark figure.

The springboard mode of oscillation becomes increasingly problematic with both flexure length and test mass length increasing. Unfortunately, these are needed to make glass inertial sensors comparable in noise performance with other state-of-the-art inertial sensors. To counter this, alternatives to the setup by Guzman et al. were thought of. A novel idea here is to use an inertial disk sensor. By making the test mass a disk with mounting points at multiple locations, the device pushes any undesired modes to very high frequencies. In turn, a disk geometry allows for the suspension of much larger test masses. The test mass is inversely proportional to the square root of the fundamental mode's natural frequency, so increasing the test mass is linearly correlated to overall noise performance, assuming the limit comes from the thermal noise.

Spring Anti-Spring Resonators

Spring Anti-Spring (SAS) systems can tune resonance frequencies arbitrarily. Two springs (or inverted pendula) oppose one another to reduce the restoring force and, therefore, the system's resonance. Such systems have found use in the Virgo super attenuator [21], and in inertial sensors [107, 108]. A naive assumption from equations 3.11 and 3.18 is that tuning the resonance frequency lower can achieve a better thermal noise floor. Harms and Mow-Lowry showed this is not the case [109]. Additional thermal noise contributions counteract the improvements in thermal noise from lower resonance frequencies. Therefore, the system's thermal noise is controlled more by its size, stiffness, and Q factor.

These limitations do not mean that SAS systems have no place in inertial sensing. Instead, they change the transfer function of the test mass, increasing the coupling of test mass motion to ground motion. Any noise flat in displacement gains significantly in signal-to-noise at low frequencies. Typical noise floors which are flat in displacement noise are those related to readout. Hence, SAS oscillators can provide a boon to any sensor limited by its readout, but not concerned with thermal noise [108].

4.4   Maximising Q Factors and Minimising Thermal Noise

As it is one of the fundamental limits of the inertial noise in a sensor, thermal noise must be minimised. It is, therefore, imperative that the highest possible Q factor can be obtained in these oscillators. The contributions to thermal noise through mechanical loss channels are discussed in Section 3.2. These loss terms follow different frequency responses to each other, which are displayed in Figure 4.12. The Figure uses the example
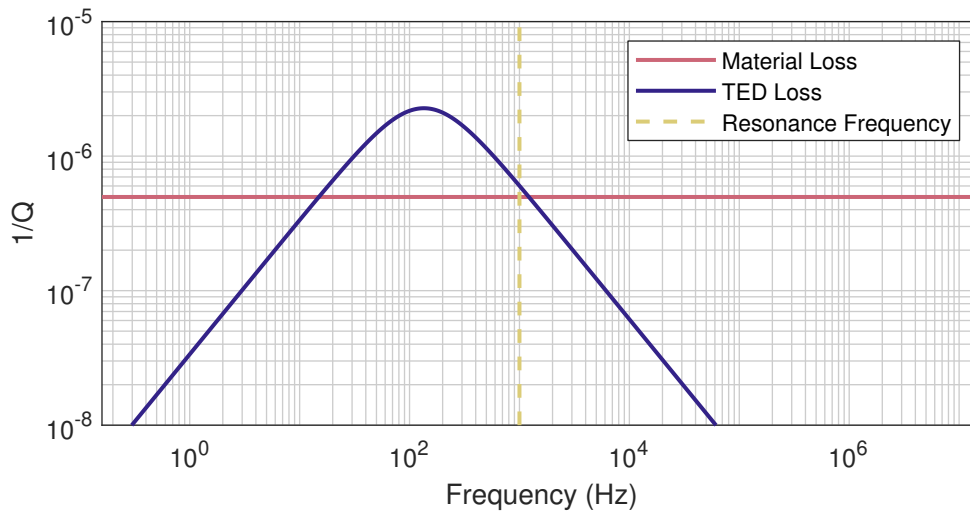
Figure 4.12: The frequency dependence of different Q terms for an arbitrary oscillator. It nominally has a flexure width of 300 μm and a resonance frequency of 1 kHz.

of a 100 μm thick cantilever. The material loss is a constant across all frequencies, whereas TED forms a peak. The height of this peak is purely material-dependent, but its location depends on flexure width. The flexure width should be tuned to move this peak out of any regions of interest, usually the measurement band and fundamental natural frequency.

4.4.1   Surface Roughness of the Manufacturing Process

Surface roughness along the flexure is critical to the device's performance. Surface roughness contributes significantly to material loss in regions of high strain energy density. Measurements of the surface roughness can be used to estimate d_s. In contrast, in worked regions, the surface quality determines whether a specific bonding technique can work or whether coating layers will stick. Due to the fragile nature of the flexure regions, abrasive polishing techniques will ultimately destroy the flexure, so the resultant flexure roughness from manufacturing is final. FEMTOprint could not verify the surface quality of the underside of the flexures, so measurements had to be taken at the AEI.

Surface profiles can be described in many ways, but the International Organization for Standardization (ISO) has specific definitions for measuring the surface roughness of a sample, given by EN ISO 21920-2:2021. The profile is a measured height variation along a representative 1-dimensional cut of the sample. To comply with the ISO standards, a specific profile length must be measured depending on the required roughness specification.

The Roughness, or R-profile, is one of three commonly used profiles for investigating surfaces, along with the Waviness (W) and Primary (P) profiles. The distinction between these is the spatial wavelength filtering applied to the measured profile. The P profile is the full range of spatial frequencies, although optionally, the lowest and highest spatial frequencies can be filtered out if not necessary. A cut-off frequency links the W and R profiles, λ_c. The value of λ_c is again contingent on the scale studied.

The wavelengths shorter than λ_c give the R-profile and longer give the W-profile. Each of these profiles can then be assessed with a variety of metrics, including the Root Mean Squared (RMS) deviation from the mean (R_a, W_a, P_a), Maximal Peak-Valley distance (R_t, W_t, P_t), and the mean Peak-Valley distance of 5 chunks of the sample region (R_z, W_z, P_z); periodic samples can use the mean peak width, but the measurements here were not periodic.

The R-profile is more relevant to the flexure region, and the W or P-profile better assess the bonding interface. As shown in Section 4.5, the longer-term variations do not matter in the flexure region. Long term variations of flexure geometry do not lead to significant changes in angular coupling or mode shape. The short-term variations can, however, dramatically increase the region's surface area, increasing surface loss.

The measurements were taken by a laser scanning microscope produced by Keyence. This laser can use several different techniques to measure surface roughness. However, for these measurements, best pictures were seen using its confocal laser microscopy setting with good explanations of the process found here [110, 111]. The drawback of using this method is the limited field of view when using the higher magnification settings. This means the insufficient area was viewed to comply fully with the ISO standards. However, as we are not trying to achieve a specific tolerance and measure the profile, this is not too problematic of a limitation. One other problem was the dynamic range of measurements. The different roughness values measured warranted different values of λ_c. Therefore, the primary profile was used to estimate the roughness, and then the appropriate value of λ_c was chosen from this using the ISO standard.

A surface can have four types of interaction with the manufacturing process, as discussed in Section 4.2. First is a surface with no laser exposure, only the chemical etch. This should have only minimal effect on the surface roughness. Next is a surface with a laser beam passing through it, but it does not etch away any material from this side. This occurs on the upper flexure side. In principle, this should have minimal effect on the surface as the laser is not focused sufficiently to etch away material, but scattered light or other effects could occur. Opposite is a material with a laser focused from directly above onto it. As this was the only way to make the drum geometry, the surface roughness left from this needed to be measured. This makes the under side of the flexure geometry. The final surface is one where the resulting surface is parallel to the direction of the laser. This was measured using a linear oscillator style geometry produced for the bonding test experiment discussed in [112].

Damaged Batch 1 Sample   The first sample discussed is one from the first batch of oscillators. This sample had a flexure broken during manufacture but was otherwise unused, meaning any damage was part of the manufacturing process. The piece was first examined with the optical microscope. Across the whole surface, small amounts of debris were found. These looked like small contaminants from the lens tissue the piece was stored in. It had been stored for two years by the point of measurement as the microscope was made available after some time. The debris was removed with isopropanol and gentle rubbing with another lens tissue.

Three regions were then examined: a region on the test mass, where no material had been etched, then the upper unetched side of the flexure, and the underside of the flexure where the material had been removed. These regions were examined on both

Region	R_a (nm)	R_z (nm)	W_a (nm)	W_z (nm)
Upper Flexure Surface	1.8±0.3	12±2	14±3	56±14
Upper Flexure Surface: Side 2	3.0±0.6	21±5	21±8	100±45
Test Mass Surface	2.4±0.4	15±3	16±6	66±25
Test Mass Surface: Side 2	3.4±0.5	21±3	13±4	56±16
Under Flexure Surface	1700±100	10000±700	170±120	700±500
Under Flexure Surface: Side 2	810±340	6000±2000	270±190	1100±700
Bonding Test Oscillator	980±140	7000±1000	160±80	650±300

Table 4.1: The measured surface profile values for seven different faces of an oscillator. The region designates where on the oscillator the measurement was taken. All but the parallel incidence side from the Bonding test oscillator had both sides measured. As this sample was produced symmetrically, the sides do not matter on this sample. The quoted value is the mean from 6 independent sample lines in a hash pattern with two perpendicular sets of three parallel lines, an example of a cut is shown in Figure 4.13. The uncertainty is the standard deviation of these values. Using a 2σ criterion, we see no difference between the two sides of the oscillator, nor any deterioration of the upper flexure surface when compared to the test mass surface. The Under Flexure etched side 2 has a very large uncertainty compared to side 1. This is because some regions of that flexure were comparatively good with R_a <300 nm, and others were much rougher with a R_a comparable to side 1. This kind of structuring was not seen on side 1, and the reason for this is unknown. Only a limited difference is seen between parallel and normal incidence surface roughness. Hence, the parallel incidence of attack should be used where possible, but does not yield a big improvement in roughness.

sides of the oscillator to asses if the manufacturing process was not symmetric. The roughness parameters measured are shown in Table 4.1. No significant difference was found between the unetched regions (the test mass and upper flexure surface).

The next region studied was the underside of the flexure. For the laser to be focused on the underside, a lower magnification lens than the previous measurement was used (50x), such that its focal length was greater than the flexure width. This reduced the precision of the method, but as the flexures were much rougher here, the loss was not an issue. The R and W profile of these measurements is shown in Figure 4.15. The underside of the flexure had considerably worse surface quality than the upper surface. This scale is shown in Figure 4.16. The surface resembles a very heavily cratered scene with a clear spatial frequency structure. This structure likely corresponds to where the laser was focused on the material. The value of R_z is 10 μm for the underside. This structuring, in turn, suggests a very large region of the exposed surface where contaminants can dissolve. In future calculations, a value of 10 μm is used as an estimate of contamination depth, d_s.

Bonding Test Oscillator A linear oscillator of the geometry in Figure 4.4 was produced as part of an experiment on bonding losses. Although the results are not discussed in this thesis a summary of the project can be found at [112]. The laser hit the material parallel to the height of the flexure, allowing for an alternative angle of etching on
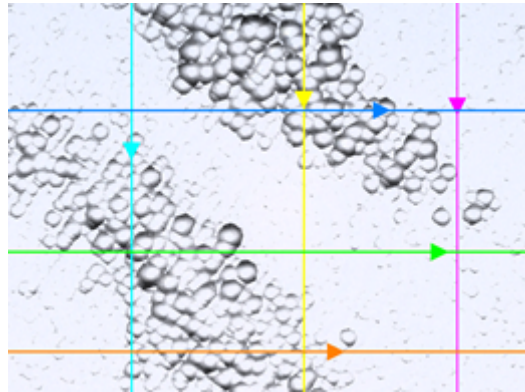
Figure 4.13: The under flexure region of side 2 studied for these measurements. The coloured lines show along which lines the profile were taken. Clearly visible are two distinct regions, one with a very rough surface compared to the other. The reason for this is unknown.
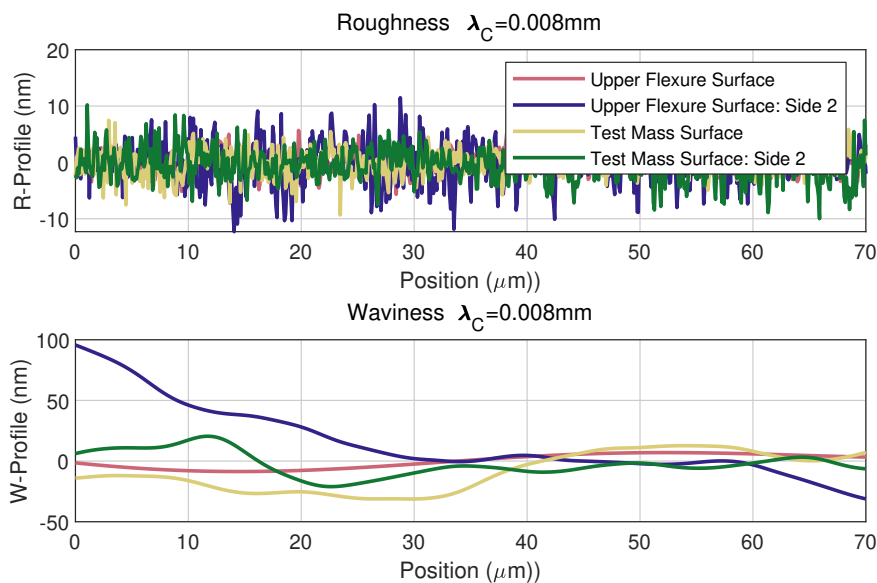

Figure 4.14: The profiles of two areas of the surface of an oscillator. One is the central test mass, and another the upper half of the flexure. Neither region was directly laser activated by the manufacturing process, but the laser passed through the surface of the flexure to focus on its underside, the effect of this on the surface roughness was unknown. Strangely the roughness and waviness profiles of the test mass region were both higher, which was somewhat unexpected. The best explanation for this is that the laser had no effect on the surface roughness, and that the region that became the test mass just had worse surface quality by chance.
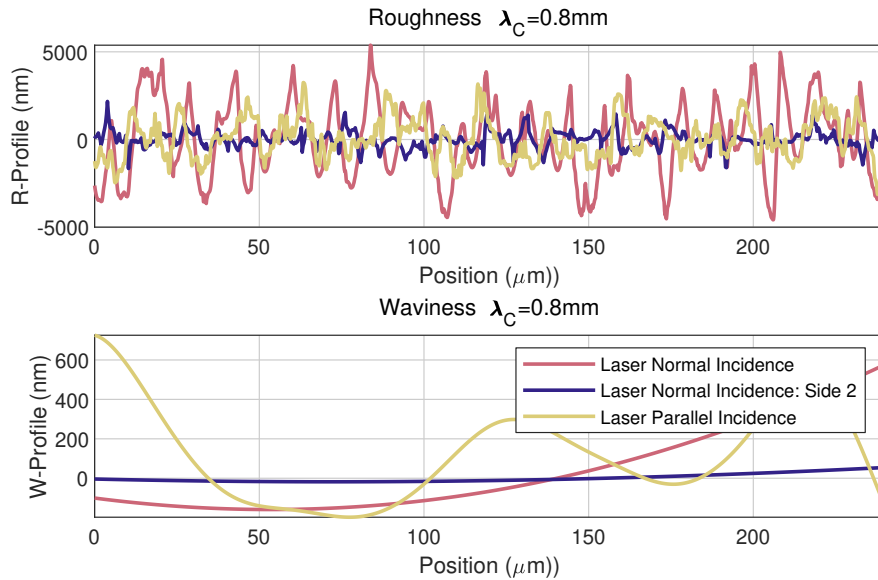
Figure 4.15: The profiles of the under flexure roughness for two different flexures. One was produced by the laser coming from directly above the surface with a normal angle of incidence and the other from along the height of the flexure, parallel to the remaining surface. The surface roughness of the normal incidence is about a factor of 2 worse depending on the metric used, but both are of order micrometers. This is a factor of ten larger than the waviness profile, suggesting the structuring along the surface has a large high spatial frequency component.
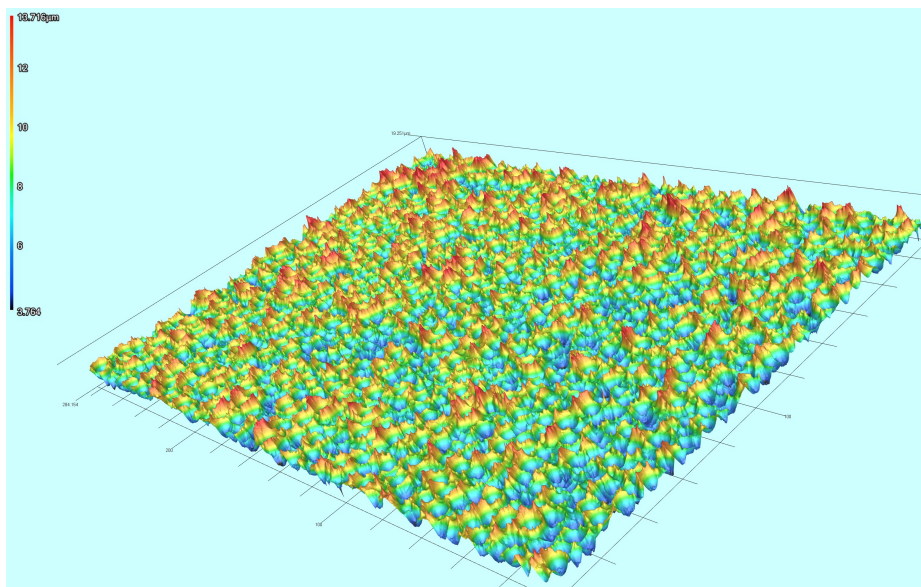

Figure 4.16: The topology of the under flexure surface after etching. The area is a square of side length 300 μm. The crater structure is very likely a result of focus points of the laser as it hits the material, creating an even region of chemical etching. The peak-valley distance of the whole area is about 10 μm.
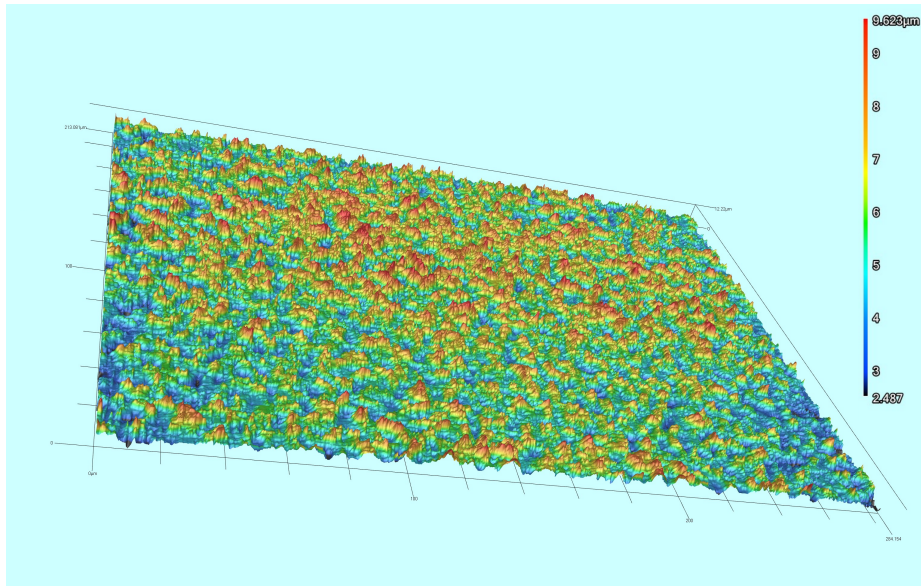
Figure 4.17: The topology of the under flexure surface after etching with a parallel incidence. The larger structures from Figure 4.16 have been removed; however, the surface still has a roughness value much higher than an untreated area.

the plane. This allowed comparison between the surface left after interaction with the laser from different angles. The results of this are listed as the parallel incidence shown in Table 4.1. The surface roughness is measurably lower than the normal incidence.. The surface topology is shown in Figure 4.17, where the previous structuring has been reduced but still leaves a rough surface. This, therefore, shows the limit of what can be achieved with these manufacturing techniques, but designs should incorporate geometry that allows this angle of attack when possible.

4.4.2   Designing Flexures with High Q Factors

We now must consider how to design flexures to minimise thermal noise. Here, we consider this in two cases, one an approach with a free design, that will minimise thermal noise, the second where a specific resonance frequency is needed.

Minimising Thermal Noise

In the first case, the readout is either arbitrarily good or not needed. The latter is rarely an option for inertial sensors. However, oscillators of this kind can also be used in a variety of applications that do not need to be read out directly, such as radiation pressure sensors [113] or when used in ponderomotive squeezing [114]. Often, these methods still rely on soft suspensions to maximise coupling between displacement and force, but as we see, this aligns with the goal of minimising thermal noise. We must consider Equation 3.18 to optimise structural damping. We will assume material constraints have defined a maximum allowable stress. Each term in this equation is defined by the flexure parameters defined in Figure 4.4. The effects of each parameter were calculated and described in Section 4.3.1, and so the resonance frequency is related

to the geometric factors by

$$f_0 \propto \sqrt{\frac{d_{\text{fwidth}}^3 d_{\text{fheight}}}{d_{\text{flength}}^3 m}}. \tag{4.4}$$

Meanwhile, the maximum allowable stress in the flexure scales was shown in Figures 4.6 and 4.7. This scales as

$$\sigma_{\text{Max}} \propto \frac{d_{\text{fwidth}}}{d_{\text{flength}}^2 f_0^2}. \tag{4.5}$$

Combining these, the maximum possible mass for a given stress will evolve as

$$m \propto \frac{d_{\text{fwidth}}^2 \sigma_{\text{Max}} d_{\text{fheight}}}{d_{\text{flength}}}. \tag{4.6}$$

Substituting this into Equation 3.18, we get that the acceleration thermal noise scales as

$$\tilde{A}_{\text{TN}} \propto \sqrt{\frac{1}{d_{\text{flength}}} \frac{1}{Q(d_{\text{fwidth}}) d_{\text{fwidth}}} \frac{T}{\sigma_{\text{Max}}^2}}. \tag{4.7}$$

The ideal flexure width depends upon the target Fourier frequency in relation to the TED peak and its magnitude with respect to surface losses. The evolution of the Q factor with respect to d_{fwidth} is

$$Q = const \qquad \text{(Bulk Loss Limited)} \tag{4.8}$$

$$Q \propto d_{\text{fwidth}} \qquad \text{(Surface Loss Limited)} \tag{4.9}$$

$$Q \propto \frac{1}{d_{\text{fwidth}}^2} \qquad \text{(TED Limited (Below Peak))} \tag{4.10}$$

$$Q \propto d_{\text{fwidth}}^2 \qquad \text{(TED Limited (Above Peak))}. \tag{4.11}$$

These responses create two scenarios; the flexures should be tuned to be as thick as possible when above the TED peak. Below the TED peak, they should be tuned such that the term

$$\frac{1}{d_{\text{fwidth}}} \left(d_s \phi_{\text{Surf}} \frac{\int_S \epsilon(\vec{r}) \mathrm{d}S}{\int_V \epsilon(\vec{r}) dV} + \phi_{\text{Bulk}} + \frac{\alpha^2 E T}{C_{\text{p}} \pi^2 \chi} \omega d_{\text{fwidth}}^2 \right), \tag{4.12}$$

is minimised across the desired frequency range. Where we have taken the equations for Q_{Mat} and Q_{TED} from Section 3.2. Typically, for gravitational wave astronomy, we aim for a control bandwidth of 0.1-100 Hz. The flexures must be centimetres thick for low-frequency seismic isolation to push the TED peak below this band. Such thick flexures are not a feasible design strategy with these manufacturing techniques and go into ranges where assumptions we have made break down. Hence, we should focus on tuning the peak to the above detection bands. We consider how best to do this for a single flexure in Figure 4.18. We can use this graph to choose a flexure width for a sensor targeting a specific bandwidth. Broadband sensors, however, have to be more careful and look across the whole frequency region. In this case, it is better to look at the effect of flexure width on the Equation 4.12, across the whole bandwidth. Figure 4.19 shows this for four different flexure widths. Although a thicker flexure improves the noise floor, it does come at the cost of high frequency performance. There are a few
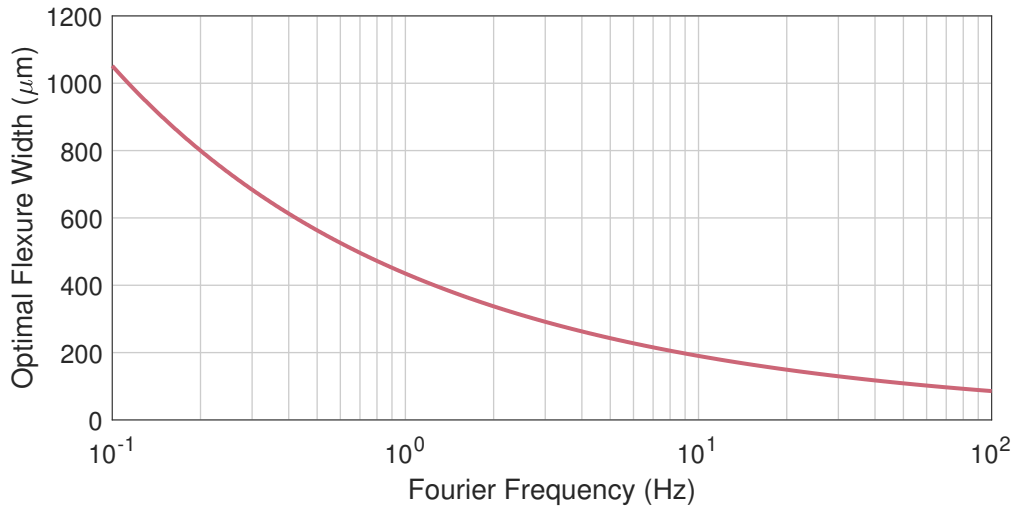
Figure 4.18: The flexure widths needed to achieve minimum thermal noise at a specific Fourier Frequency. Here the Term 4.12 was minimised for flexure width for each Fourier Frequency, and the value plotted. For this Figure, the approximation that strain density is uniformly distributed over flexure width is made (see Section 4.4.3), a material of Corning 7890-0F fused silica is used, with a surface loss of 1×10^{-5} [80], with a surface depth of $10\,\mu$m from measurements in Section 4.4.1.

ways of deciding on flexure width from this Figure. A maximum acceptable noise floor across all frequencies could be targeted (particularly if readout noise is already known), and a flexure width that does not violate this at any frequency chosen. Alternatively, some form of (possibly weighted) average taken over the desired range could be used as a minimisation criterion. Often, this decision is also simply defined by what can be made. Ultimately, optimum flexure width is a very project specific definition.

In the case of this project, prototype high frequency sensors discussed in section 5.1 were designed with very thin flexures, $50\,\mu$m thick, so that they could achieve the best Q factor at their resonances above 100 Hz. The 50 Hz oscillators in Section 5.4 had a thickness of $180\,\mu$m as it offered close to the optimum point on resonance and a good performance across the relevant frequency band. The flexure width of the 15 Hz oscillators discussed in section 5.5 was again decreased to $100\,\mu$m primarily to enable a survivable oscillator at 15 Hz while still fitting in a 50 mm design. As the manufacturing process of the 15 Hz oscillator was different, they likely would have lower surface losses, so the minimum calculated here would be considerably different. Measurements of the resulting surface must still be taken to estimate its contribution.

The flexure height and separation are not related to the thermal noise. Therefore, these are the free parameters to tune the mode frequencies and the mass. Factors such as radius of curvatures were discussed in section 4.3.1, but already have clear optimums without the need for trade-offs. A study following the steps discussed there can find these points for any flexure, and the only reasons not to include them are for manufacturing simplicity.

With two caveats, the flexure length should be tuned as long as possible. Lengthening the flexures is the most difficult parameter change to achieve with current design methods in terms of cost and complexity. Methods to make the manufacturing process
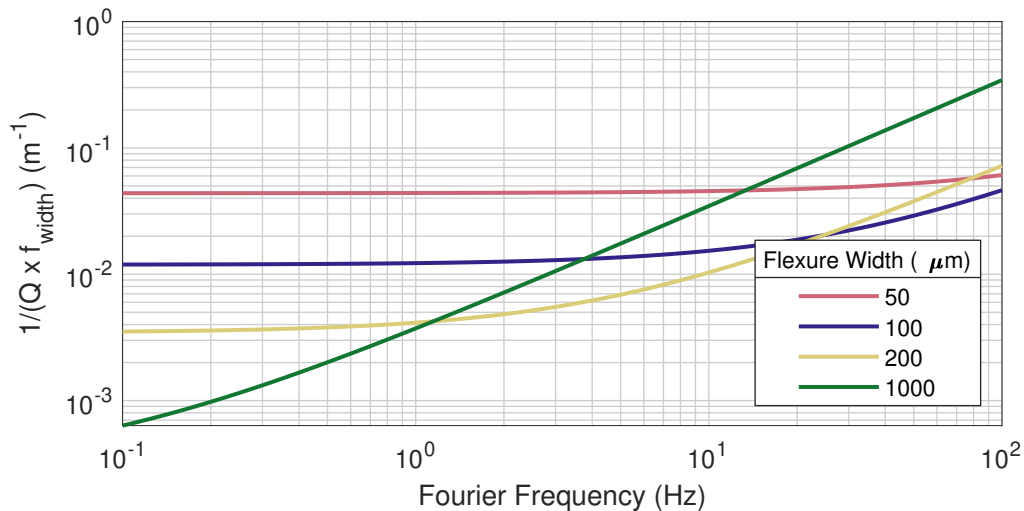
Figure 4.19: The results of the Term 4.12 for flexures of 4 different widths across the Fourier frequency range of interest to the Seismic Isolation community. The strange units are effectively a proxy for thermal noise. The same assumptions, values, and material used in Figure 4.18 are used here.

more scalable with length are being researched and are discussed in Section 5.5. Even if this can be achieved, there are still fundamental limits to how far flexure length can be taken. As the flexures get longer, the test mass gets smaller, increasing the contribution of residual air pressure damping. In extreme regimes, viscous damping becomes dominant. Viscous damping responds differently to flexure length. Furthermore, as the flexures get longer with a fixed width, a point will be reached where internal flexure modes of oscillation will approach the mode desired and become problematic. Finally, with such a soft spring, the radiation pressure noise becomes problematic for readout, so such designs are used in experiments to measure and suppress the effect [113].

Fused silica is a standard choice for a room temperature high Q oscillator [60]. Fused silica is used due to its low bulk loss, and comparatively high shear stress to other low-loss glasses [62]. Numata studied the specific bulk losses for several types of Fused Silcia [62], good candidates of specific materials were Corning 7890-0F, and Heraeus Suprasil-312. The latter of these has a slightly lower bulk loss, but the former was used due to the availability of wafers in all but the first batch of oscillators.

The other parameter that can be adjusted is the operating temperature. A problem with doing so, other than the infrastructure requirements, is that the material properties of fused silica made its TED peak much more prominent at these temperatures [115]. The TED peak means other materials must be used for cryogenic, low-loss applications. Tests are ongoing to adapt the manufacturing techniques to work with crystalline silicon, which has low bulk and TED losses at cryogenic temperature. However, the ultimate shear stress of crystalline silicon is a factor of five smaller than fused silica. Looking back to Equation 4.7, we would only see a gain in thermal noise between the designs at sub 12 K. As the low-frequency component of ET will be cooled to 10 K, such sensors would have to be developed for this use, but the sensors at room temperature are a good estimate of the performance that will be reached. Other materials, such as sapphire, have good loss characteristics at cryogenic temperatures and high shear strengths, but

are very difficult to structure. When approached, companies have said they cannot do such fine structuring with sapphire.

Resonance Frequency Fixed

Often, limits on readout and application define f_0. With f_0 defined, the question is how to optimise around this restriction for thermal noise. Again, we can consider the same factors as the previous, but now with the additional criteria that f_0 is constant. The limitation places a restriction

$$d_{\text{flength}} \propto \frac{d_{\text{fwidth}} d_{\text{height}}^{1/3}}{m^{1/3}} \tag{4.13}$$

in order to maintain the correct resonance frequency. The restriction also applies to the maximum suspendable test mass

$$m \propto \sigma_{\text{Max}}^{3/2} d_{\text{fwidth}}^{3/2} d_{\text{height}}. \tag{4.14}$$

With the resonance frequency fixed, thermal noise is controlled by the mQ product. Substituting limitations into Equation 3.18, we get the scaling

$$\tilde{A}_{\text{TN}} \propto \sqrt{\frac{T}{d_{\text{fheight}} Q(d_{\text{fwidth}}) d_{\text{fwidth}}^{3/2} \sigma_{\text{Max}}^{3/2}}}. \tag{4.15}$$

The condition on width is slightly adapted from the free design case

$$\frac{1}{d_{\text{fwidth}}^{3/2}} \left(d_s Q_{\text{Surf}}^{-1} \frac{\int_S \epsilon(\vec{r}) \mathrm{d}S}{\int_V \epsilon(\vec{r}) dV} + Q_{\text{Bulk}}^{-1} + \frac{\alpha^2 E T}{C_{\text{p}} \pi^2 \chi} \omega d_{\text{fwidth}}^2 \right) \tag{4.16}$$

which must again be minimised with equivalent Graphs to Figure 4.18 and 4.19. This again leads us to a similar optimisation study as Equation 4.12 for flexure width.

A key difference here is the scaling with flexure height. Increasing the flexure height allows a greater mass to be suspended, making the flexures stiffer. This scaling will allow for a heavier test mass without compromising maximal stress. The limits here are again similar to the length scaling in the previous case. The increased height will lead to a larger volume etched, increasing the manufacturing cost and complexity of the devices. Furthermore, figure 4.11 shows that the torsional mode will eventually become problematic as the height is increased. The length of the flexures is defined by what is needed to survive under load. The supported mass can then be tuned to give the correct f_0.

As energy costs of etching the material drive the cost of the manufacture, it is unlikely that a dramatic cost reduction will occur soon. Therefore, there is a strong desire to investigate manufacturing techniques that scale better with flexure lengths; such an approach is discussed in Section 5.5

4.4.3   Strains and Corners

A deeper look at the material loss term is needed as it is a dominating loss term for most of the frequency ranges of interest. For this, a linear oscillator of the geometry
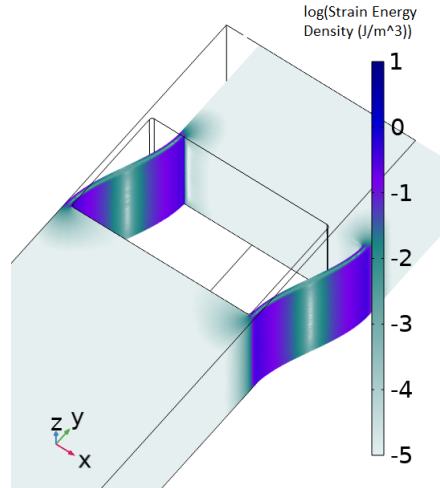
Figure 4.20: A logarithmic heat map of the strain energy distribution in a linear oscillator when the test mass undergoes a 1 μm displacement. The strain energy is heavily localised to the corners of the oscillator.

of Figure 4.4 is again considered. The first thing to study is the distribution of strain density, or ϵ, from Equation 3.21. So far, we considered the energy to be homogeneously distributed through the flexure region, but this is not the case.

The theoretical strain energy distribution was solved for an ideal plate in Appendix B. At the point of maximum displacement during oscillation, all energy should be potential energy. The strain energy distribution is given by

$$\mathcal{E}_{\text{PE}} = d_{\text{height}} \frac{\langle z, z \rangle}{2} \frac{1}{1 - \nu^2} \left(E \kappa_{yy}^2 \right), \tag{4.17}$$

where ν is the Poisson ratio, E is the Young's Modulus, and κ_{yy} is the element of the curvature tensor with respect to the direction parallel to flexure length. Except for edge effects, the strain is only dependent on edge curvature, and uniform across the width and height. Therefore, we only consider effects across flexure length.

FEA with COMSOL Multiphysics was used to solve the strain distribution. A stationary analysis model with a prescribed displacement applied to the test mass was used. Shown in Figure 4.20 is a heat map of the strain density, which is localised to the corners of the flexures in the case of a 1 μm displacement.

Since the material loss depends on the surface-to-volume ratio, adding more material to the flexures in high-energy density regions makes sense, allowing better energy distribution along the flexure. The width can be altered by adding an inner radius to the flexure corner. A test was performed whereby different inner radii were tested to find an optimum point for strain distribution. The radius was added as indicated by Figure 4.4. Comsol was again used to simulate these effects, using a combination of static and eigenfrequency study. The effect of different inner radii on strain distribution is shown in Figure 4.21. Adding the inner radius pushes the peak of the strain distribution inwards before decaying into the bulk.

The effect of this on surface loss is then shown in Figure 4.22. To calculate the surface loss, Equation 3.21 was used, with the distribution of strain densities shown in Figure 4.21 as the estiamte for ϵ. The result of d_{s}=10 μm is taken from 4.4.1 and a
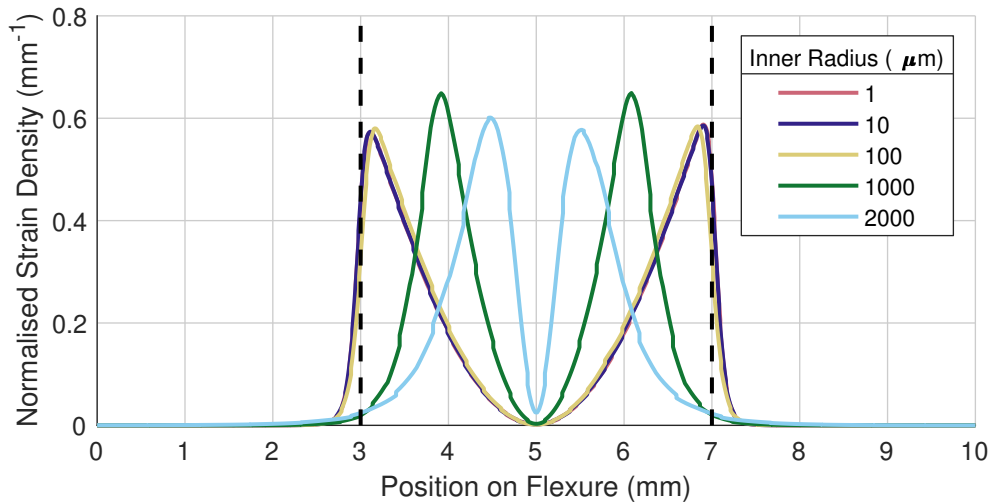
Figure 4.21: The normalised distribution of strain energy density for a cross-section across the flexures at positions along the flexure length. The energy is normalised, such that it is a ratio of total energy. This is shown for different inner radii of curvatures at the flexure ends. The black dashed lines show the ends of the flexure region and the return to the bulk material. The geometry studied is the same as defined in Figure 4.11

.

surface loss value of 1×10^{-5} from [80]. We see in Figure 4.22 that a large inner radius leads to a lower surface loss. However, this only has a noticeable effect when the inner radii are above about $10\,\%$ of the flexure length. The corresponding study on extra stress in Figure 4.9 shows that stress becomes more localised and would need longer flexures to compensate. The extra stress scales roughly as power $1/2$, while the gain in surface loss scales roughly linearly. Considering the two cases discussed in Section 4.4.2, we see that the gains of surface loss are cancelled out for the case f_0 is a free parameter above the optimum for any radius of curvature above its optimimum stress point, Instead, if f_0 is fixed by external constraints, inner radii scales as power $1/4$ with thermal noise. This is only true while surfaces losses are dominating. As the gains from increasing the inner radius are small, they should only be done when it does not add to manufacturing complexity.

When f_0 can be chosen, this means adding an inner radius $10\,\%$ of flexure length. Alternatively, when a specific f_0 is required, the optimum radius of curvature is a little under half the length, but we will only see gains while surface loss limited.

4.4.4   Worked Regions

Several regions of the oscillator require additional manufacturing steps to be made into a working inertial sensor. These regions will be called the worked regions. Some examples are the surfaces that have coatings applied, regions where readouts are attached, and where the device is mounted. Tampering with these regions can have unexpected results on the Q factor. The distance between these and flexures regions should be maximised to minimise the loss on the Q factor.

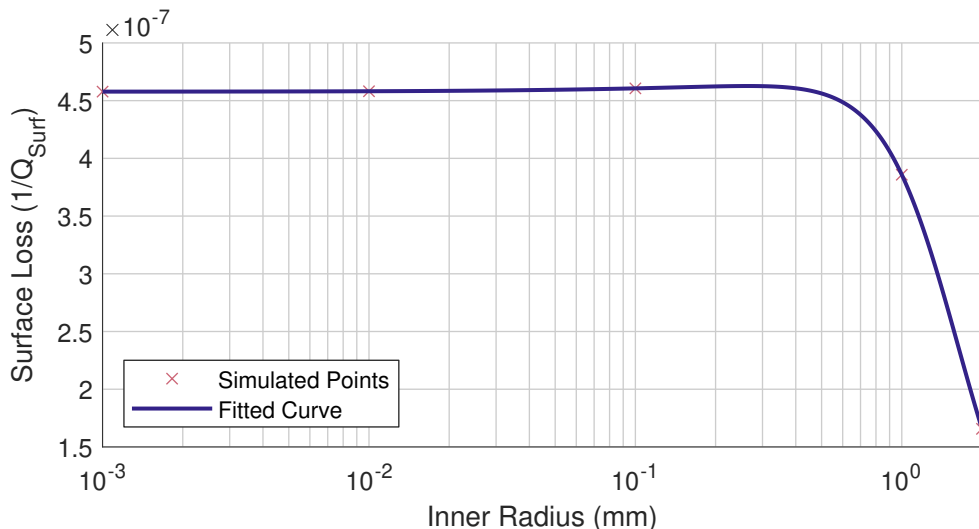Similar studies to that which produced Figure 4.21 can be done, but where the

Figure 4.22: The surface loss of an oscillator with geometry used in Figure 4.11 when different inner radii are used. The simulation was done for the five inner radii in Figure 4.21, and fitted with a cubic interpolation routine. A larger inner radius leads to a lower surface loss.

worked region is modelled as a separate block with a much higher intrinsic bulk loss. When doing so, we find that very little strain energy is distributed in the worked region, even when it is just a small distance away from the flexures. In principle, this means that the distance of worked regions can be quite close, typically just a few millimetres. However, in practice, bonded areas apply stress to the surface. These stresses can cause exceptionally high losses, as shown in Section 5.2.4, with effects difficult to simulate or predict. Therefore, it is advisable to test these worked regions and methods on lighter or simpler samples first, as is the approach in Section 5.3.

4.5   Tolerances and Mismatches

The current discussion has so far focused on idealised cases. When made, every defined parameter will have some deviation from the idealised case. The effect of these aberrations and how best to define them is discussed in this section. As the flexure region defines the oscillatory behaviour, and small deformities greatly impact performance, tolerances on the flexures are the most important for oscillator design.

The surface profiles of flexures produced using the methods discussed in this thesis are presented in Section 3.3.2. The surface roughness is a measure of the very high spatial frequency effects. As such, its effects on the large-scale flexure geometry are limited and so only matter with regard to material loss calculations. The waviness profiles measured in this section are orders of magnitude lower than the quoted geometric deviations. The tolerances of concern are those that disturb the larger flexure geometry, especially in an uneven manner.

We will, for this study, confine ourselves to long-term deviations from a specified flexure geometry. This would still be an infinite possibility space, but when consulting with companies and manufacturers, three common deformations come to light that

must be defined. Controlling the average width of the flexures across the structure is very difficult, with tolerances typically of the order of +/-10%. It is also possible that the flexure width of one or both varies along the length of the flexure. For example, if the wafer or etching machine is not perfectly flat, the width would change linearly over its length, creating a trapezoid geometry. The trapezoid geometry affects either or both flexures and could happen in the same direction or opposite to each other. Finally, the surfaces of the test mass and base facing each other may not be flush. As the distance between the two would vary, one flexure would have to be longer to compensate for this. We investigate the impact of each of these on the design throughout this section.

After defining the parameters we wish to explore, we can study this again for the case of a linear oscillator to understand the general effect of these mismatches. Similar studies would be needed on specific geometries that are to be manufactured. As an example, we use the same linear oscillator model from Figure 4.11 and the same FEA model.

The effect of varying average flexure width on the modes is shown in Figure 4.23. Here, two cases are studied, one where the average width is split evenly between the two flexures and one where one flexure is held at 140 μm width and the other varied. We see limited effects on any mode from the distribution, but the mode frequency depends on the total flexure mass. The only relevant mode frequency effect is if the device can survive operation with any flexure width within quoted manufacturing tolerances. This is best done by repeating FEA simulations with extrema of the tolerances and checking safety factors in these cases. Although the mode frequency effect was minimal in this case, the mode shape itself could be distorted. Imbalances between the flexures could lead to mode shape distortion. In practice, this distortion would mean the test mass no longer being perpendicular to the mode of oscillation. The effects of differing imbalance on test mass angle, θ, relative to the load applied, were studied using the same linear oscillator as the study in Figure 4.11 and the results are shown in Figure 4.24. These graphs show how θ evolves with mismatches between different parameters.

The first plot shows how under 1 g load, θ varies with the width of one of two flexures changing. The simulation shows the expected result, when the flexures are the same width, there is no angular offset, but the angle increases as the flexures mismatch changes. θ scales linearly with respect to load, as shown in the fourth plot on this Figure. This is problematic as it is a noise source that scales with signal, making it very difficult to model. The angle of the test mass will couple into readout measurement noise through Tilt to Length Coupling (TTL). The effects of TTL are subtle and numerous. They are thoroughly detailed by Hartig et al. [116]. The coupling to noise will vary enormously depending on the detector configuration and the readout scheme. With estimates of the angular noise introduced by the mismatch, the effect on the readout can be predicted using the models in the paper.

The effect of trapezoid-shaped flexures is shown in the second plot in Figure 4.24. Two cases are shown here, with opposite ends of the flexures being the ones that are mismatched. Both sides show similar results but with the rotation in opposite directions. When the base is varied, a slightly greater coupling to angle is seen, but both show a greater susceptibility to tilt than the purely average width change or length change cases. Tolerances should therefore be defined to minimise this criterion best. This is best done by defining a maximum and minimum value of the flexure width
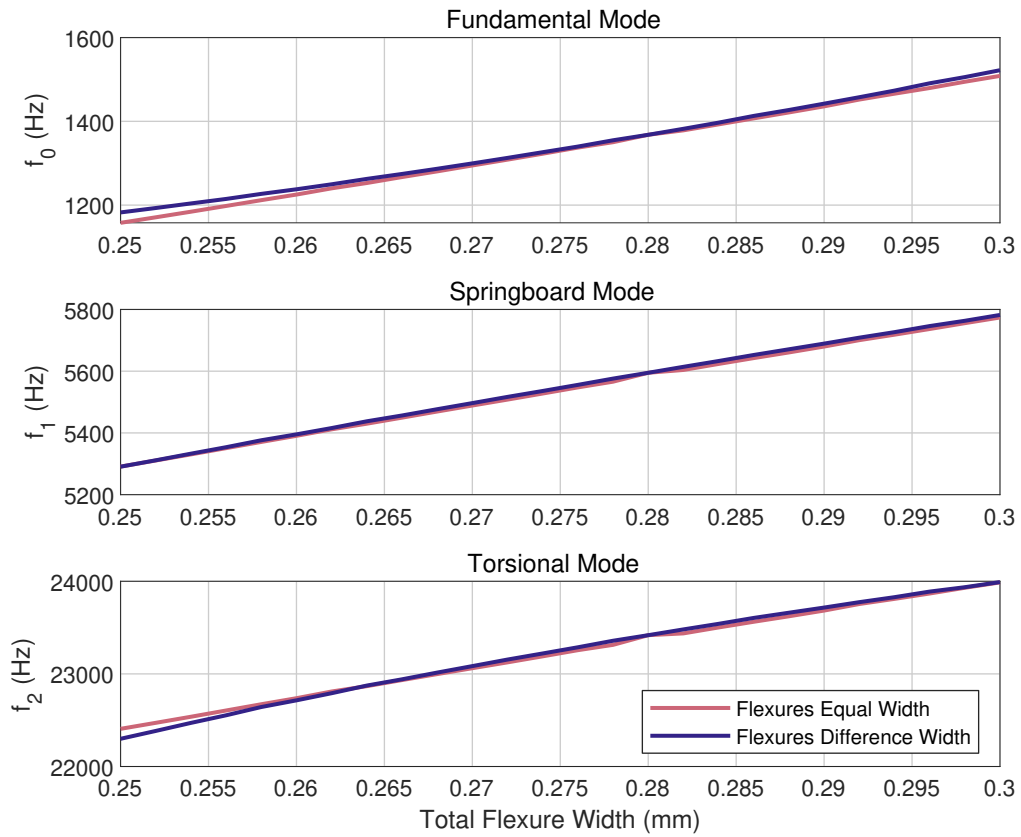
Figure 4.23: Dependency of the frequencies of the first three modes of a linear oscillator, with parameters described in 4.11, with respect to flexure width mismatches. The graph shows two cases. The red curve describes the case when the total flexure width is evenly divided between the two flexures. On the other hand, the blue curve is when one flexure of the pair is fixed at $0.14\,\mathrm{mm}$ thick while the other changes width. By comparing both, we see that the mode frequency is only slightly dependent on the ratio between the two flexures and is mostly dependent on the total width of the two flexures combined.
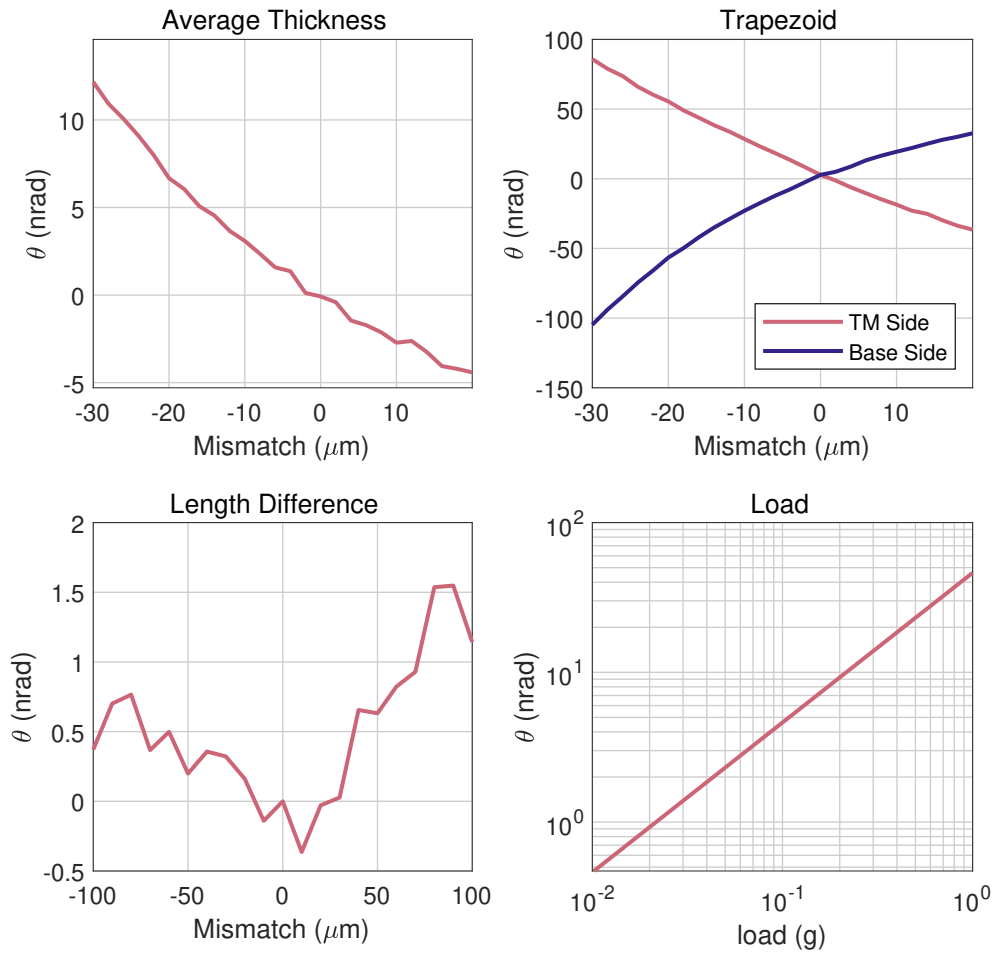
Figure 4.24: The angular dependence of a linear oscillator with regards to a variety of tolerance mismatches between flexure pairs. The tests were conducted with a linear oscillator as described in Figure 4.11. A stationary FEA analysis was performed in COMSOL. The tests were done with 1 g of load on the oscillator in the direction of the fundamental mode. The angle of the test mass relative to the load is shown for several cases. The first plots shows the case of two flexures, which are perfect cuboids with a width difference. The second graph shows when the flexure is a trapezoid, with the width changing over its length. The case where the base side has the correct width and the test mass side is mismatched is shown in red, while the blue line shows the opposite orientation. The length difference shows the case of one flexure being longer or shorter than the other. The fourth plot shows how a fixed mismatch in average flexure width of 20 µm scales with θ with differing loads.

that cannot be exceeded at any point along the length. Typically, this criterion has an achievable +/-10 µm tolerance.

The length of the two flexures was changed so that one was longer than the other. The result on θ is shown in the third plot of Figure 4.24. The effect of the length change was so small that uncertainties from the mesh elements seemed to dominate the test. Even with different lengths, each flexure can act independently and linearly without changing the mode shapes. Hence, this seems not to be a critical criterion.

4.6   Chapter Summary

This chapter aimed to build the toolbox to design optical inertial sensors for use in gravitational wave detectors. A key element here is the coupling of mechanical motion and optical readout. The interaction between these two determines our ability to choose readout techniques and the scale of difficulty in mechanical design.

A full description of the parameters available for oscillator tuning has been introduced. How these parameters tune the oscillatory behaviour and survivability of the designs has been thoroughly evaluated. As the oscillators are often mechanically simple, these geometric factors are the only meaningful way to tune our systems to behave as we want; understanding them is crucial for those looking to design such sensors or similar optomechanics.

The effects of tolerances are also investigated here. If not properly accounted for, they can take a near-perfect design and leave it a pile of shattered glass. There are several ways of defining tolerances with such a beam, but an absolute maximum and minimum width along the flexures is shown to constrain the design best.

Part II

Experimental Validation of Inertial Sensor Designs

Chapter 5

Designing and Characterising Drum Oscillators

The process of designing inertial sensors discussed in Part was developed alongside the design and manufacture of actual samples. This chapter presents results on producing oscillators. We also test the properties of these oscillators to show how well the prediction match produced samples. Testing samples for their mechanical properties is best done using a ringdown method. In this chapter, how these measurements are taken and what lead to the best results are discussed.

5.1   Prototype Drum Oscillators

A prototype drum oscillator was designed to allow us to get experience with handling samples. The designs were produced towards the start of the work, before much of the analysis in Chapter 4 was conducted. Hence, the designs were not optimised, but these prototypes had a few key science goals. They aimed to show that high Q factors could be maintained with more complicated geometries than a linear oscillator. Furthermore, the AEI had limited experience handling these samples; we aimed to build up our tools and ability to handle these samples. To do this, we aimed for a higher frequency oscillator design, which would be easier to handle.

These samples were originally designed with short flexures, which was acceptable due to the high resonance frequencies. They were able to survive 10 g of load. The original flexures were designed to have thicknesses of 50 μm, but the companies approached could not manufacture this. To accommodate them, we increased this thickness to 80 μm, which pushed the expected drum mode frequencies up to about 300 Hz. The higher frequencies would further limit the sensor's performance at low frequencies, but this was a small concern as these were mostly initial test samples.

The oscillators were made by FEMTOprint, a company which uses a 2-step manufacturing process for glass structuring [63] whose methods are discussed in 4.2. The first batch was a batch of six. Three different designs were produced, as shown in Figure 5.1. Design 2 had 4 flexure pairs. Originally, this was planned so that the pairs could be aligned with gravity to see what effects this had on Q factor. It was unknown if they would compensate for one another favourably. The other design change was the 7mm thick design for design 3. Thicker designs would lead to more mass and a lower f_0, and so a better thermal noise. FEMTOprint was unsure of their ability to cut to that depth, so to get pieces in hand, only two were ordered with this thickness. Two versions of each design were ordered.

5.2   Ringdown Experiment

A ringdown experiment is a common method of measuring an oscillator's Q factor. To perform this experiment, the oscillator must be excited and its decaying oscillations measured. The time taken for the energy of the oscillation to fall to $e^{-2\pi}$ is recorded

(a) Design 1   (b) Design 2   (c) Design 3
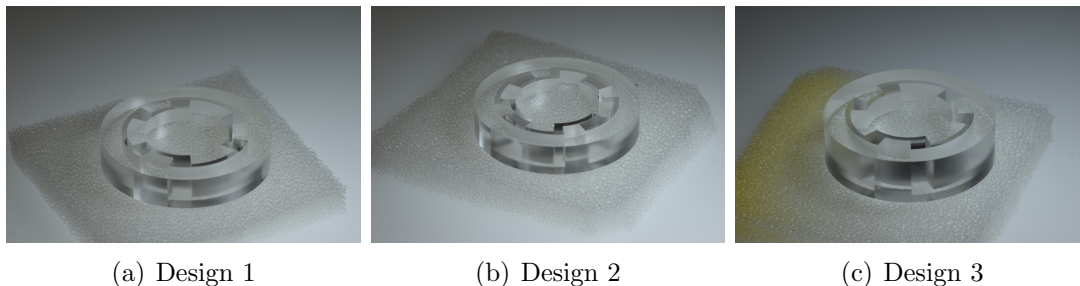
Figure 5.1: Photos of the three oscillator designs in the first testing batch. Design 1 has three flexure pairs and is 5 mm thick. Design 2 has four flexure pairs and is 5 mm thick. Design 3 is identical to design 1, but is 7 mm thick.
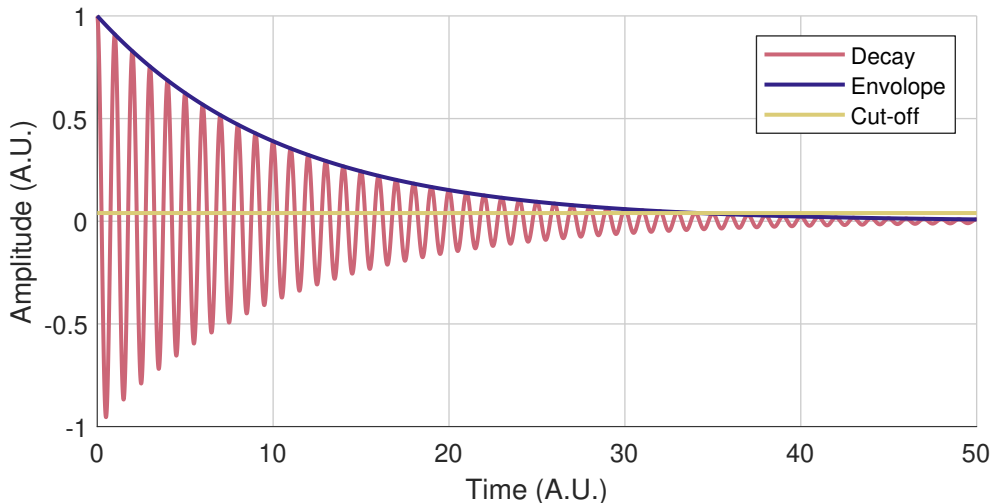
Figure 5.2: An example decay of oscillation of an oscillator with a time period with 1A.U. The Cut-off line is when the amplitude of oscillation equals $e^{-\pi}$ of the initial amplitude. The decay falls below the cut-off point at time=33 showing a Q factor of 33.

(remembering that energy is proportional to displacement squared). Multiplying by the oscillation frequency gives an approximate value for the Q factor as long as it is sufficiently large. Figure 5.2 shows a typical excitation and decay of oscillation. For this experiment, we take the peak amplitude of each oscillation and fit these points to the function

$$A = A_0 e^{\left(-\frac{\pi f_0 t}{Q}\right)}, \tag{5.1}$$

where A is the amplitude of oscillation, f_0 is the natural frequency, and t is the time since excitation. The fit of this gives the envelopes shown in many figures in this section. The quality of the fit of the envelope to the data can be used to assess the measure of the Q factor. When there is a discrepancy, another mode of oscillation is likely causing a disturbance, or the signal decayed below measurement noise. In order to perform a ringdown measurement of a mechanical oscillator, we need two things: a means of exciting motion and a means of measuring motion.

We achieved the former by mounting a large Piezoelectric (Stack) (PZT) behind the oscillator holder. A step voltage would be applied across the PZT, causing it to expand. The stack would expand into the holder of the oscillator driving it. A kick voltage of 20 V was sufficient to excite motion at about 1×10^{-7} m; enough to be easily seen above the interferometers noise floor.

5.2.1   Heterodyne Interferometer

Laser interferometry was the most appropriate readout method, as the test pieces were optical in nature. The system needed a large dynamic range as we expected large Q factors. Therefore, we used heterodyne interferometry.

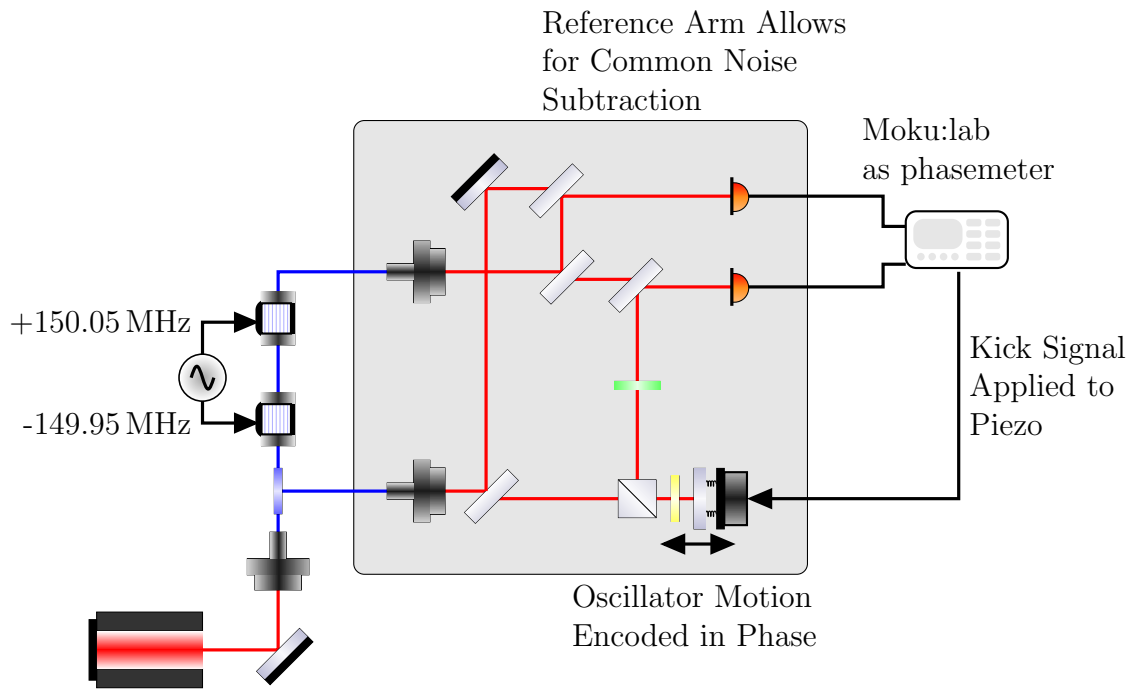Figure 5.3 shows the schematics of the interferometer used for this. The beam

Figure 5.3: The schematics of the ringdown experiment. The setup has two interferometers, both heterodyne, read out with a phasemeter. The one which does not contain the oscillator is used as a reference. The other is the measurement interferometer. All free beam beamsplitters are 50/50. Contrast can be maximised by matching the power in the two arms of the measurement interferometer using the appropriate fibre-splitting ratio, which must account for the sample's reflectivity. The Acousto Optic Modulators (AOMs) have a combined efficiency of $\sim 25\%$. Fused silica has an uncoated reflectivity of $\sim 4\%$ at 1064 nm.

undergoes preparation in fibre outside the vacuum with the Acousto Optic Modulators (AOMs). In the vacuum chamber, it is converted to free-beam via fibre collimators. We use two interferometers. The reference interferometer consists of only fixed optics. Meanwhile, one arm of the measurement interferometer includes the oscillator, which encodes its motion in the phase of the beat note. A Moku:Lab acting as a phasemeter reads out both arms, using IQ demodulation. The difference between the two interferometer phases subtracted the common motion to reduce the noise floor further. The kick from the PZT can excite other modes of oscillation in the mount and oscillator. In principle, these modes can be filtered out. However, waiting for these modes to decay before fitting the main oscillation is still advisable. Countermotion between oscillator modes and mount modes may lead to excess friction losses.

The oscillator must operate in a vacuum to prevent air pressure damping. A Q factor of 7000 was the highest measured in air in the prototype batch, while coating testers discussed in Section 5.3 only could achieve a Q factor of 1000 in air due to their low mass. The free beam optics are breadboard mounted to allow for quick switching between experiments in the vacuum and are shown in Figure 5.4. Mounted on a breadboard outside the vacuum were the fibre optics. The breadboard sits upon four VIB100 isolation feet from the Newport Corporation to isolate from seismic disturbances. The experiment sits in a vacuum tank; the pre-pump was outside the lab to avoid additional disturbances. However, the turbo pump is mounted below. The turbo pump can simply not be activated to reduce vibrations in the experiment so that we can choose between lower air pressure and additional disturbances. The turbo pump spins at $800\,\text{Hz}$, so it should be sufficiently far away from most oscillators' resonance frequencies. We reached a pressure of $1 \times 10^{-2}\,\text{mbar}$ with pre-pump alone and a pressure of $1 \times 10^{-7}\,\text{mbar}$ after 24 hours with the turbo pumping. We found that there was no measurable difference in Q factor at either of these pressures, so $1 \times 10^{-2}\,\text{mbar}$ was sufficient for these measurements. We measured the noise floor of the setup at several points during the iterative design process. Figure 5.5 shows the final noise floor achieved. The Figure shows two different lasers used during the thesis work. The diode laser proved to be the leading noise source across the entire frequency band of interest (10-1000 kHz). The laser diode was the only laser available during the measurements before coating. After coating, a Non-Planar Ring Oscillator (NPRO) laser was set up for use in the cavity readout; we used this for all subsequent measurements.

5.2.2 Clamping Loss

We did not achieve a Q factor of more than 7000 measured during the first measurements, even though a limit of 400000 was expected. All initial samples measured a Q of around 7000, suggesting that there was a common element limiting the performance. The main suspect for this was the mounting technique for the oscillators. The original was a standard mirror mount with a side retention arm for clamping, shown in figure 5.6 (a).

This mounting concept was a concern as the side mount effectively held the oscillator in place with friction. The test mass moving forward and back would cause the frame to move. Rubbing the frame against the screw would lead to local heating, and then the heat would be radiated away from the locally hotter spot, acting as a loss channel.
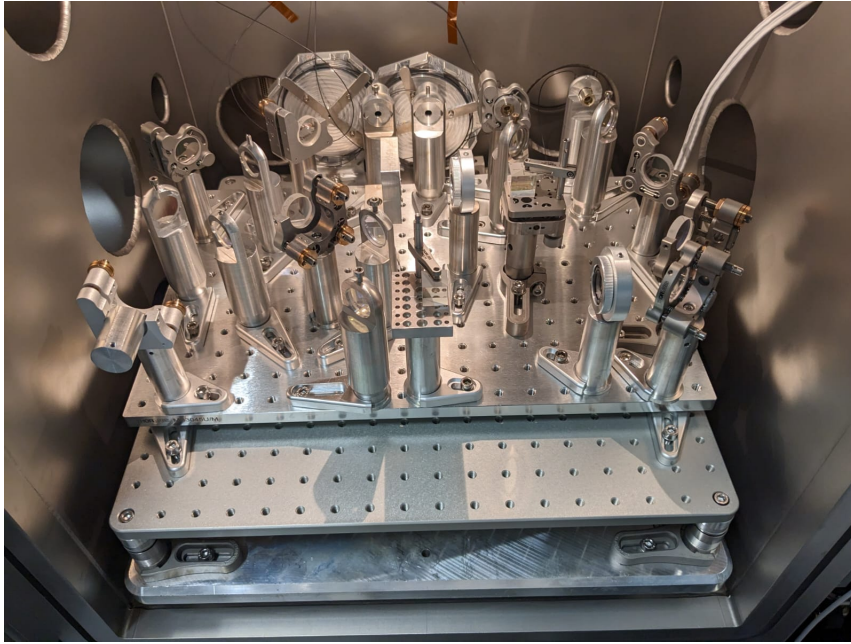
Figure 5.4: A photo of the optical breadboard used in the ringdown measurement. A large base breadboard is attached to four vibrational isolation feet to keep the experiment isolated from excitation from the ground or pumps.
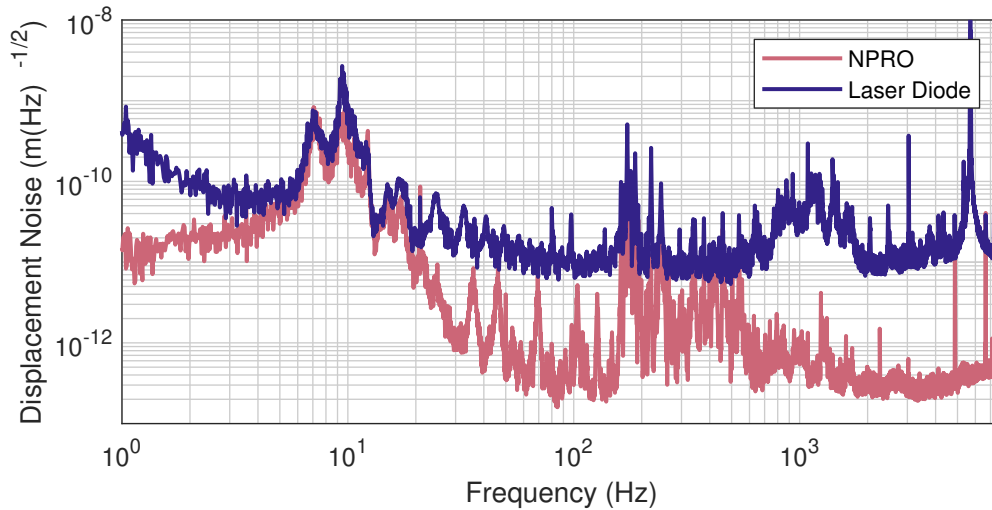

Figure 5.5: The noise floor of the ringdown experiment when using the two different lasers available in the work. The laser diode was used first until the NPRO was installed. The common noise peak at 7Hz is the expected resonance of the Newport isolation feet. The other peaks are likely mirror mount resonances. These measurements were taken at 1×10^{-2} mbar. The optical power in both arms was ∼5 mW for both lasers, although the NPRO was operated at maximum power (500 mW) and simply tapped with a fibre splitter. A contrast of 80 % was measured.

(a) Side Screw

(b) Rigid PEEK Disk/Front Loaded Spring

(c) Front Loaded Spring with Alignment Stage and Adaptor Plate

(d) Front Loaded Spring with Rigid Block
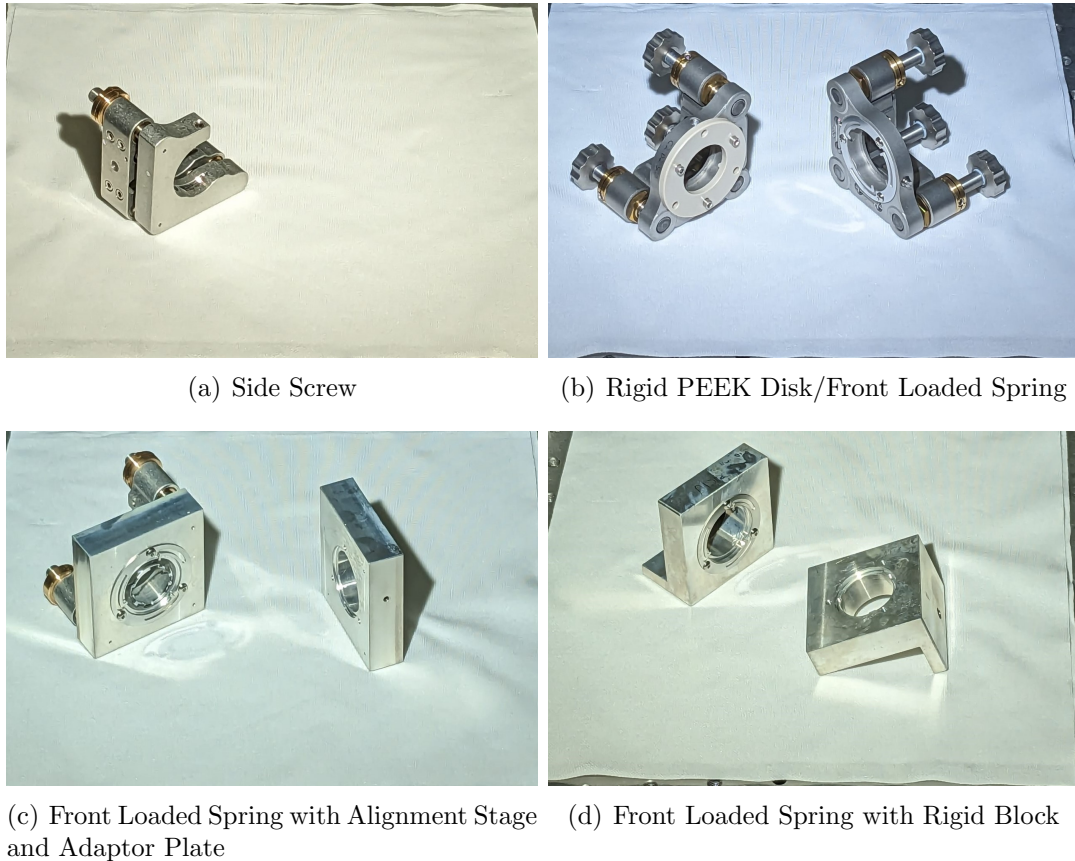
Figure 5.6: Photos of the different mounts for oscillators used to minimise clamping loss. The side screw (a) was the original style of mount used but had problems with clamping loss. The rigid PolyEther Ether Ketone (PEEK) disk was also difficult to use, and so was discarded after a sample was damaged. The front-loading springs were much less likely to damage the mounts as three softer contact points were made. The mount shown is a stock mirror mount from Newport (b). Adaptor plates were made to accommodate samples of different thicknesses (c). Eventually, it was suspected that the adaption to the alignment stage may be degrading Q factor; therefore, a rigid block was made instead (d).

Therefore, we decided to switch the design for the mount to one with a front restoring force, which would prevent the oscillators from moving relative to the frame. We tried a few mount designs, a front-loaded spring mount and a solid rigid mount, shown in Figure 5.6 (b). The rigid PEEK design never achieved a Q factor higher than 10800 and proved difficult to handle, even damaging an oscillator when mounting.

The initial prototype batch showed only a Q of up to 28000 with the front-loaded spring with the standard optical mount with three alignment screws. For the coating testers, an adaptor plate had to be made. Originally, this was designed to be attachable to the alignment stage of the three alignment screw mount, Figure 5.6 (c), but this was switched to a single solid mount as shown in Figure 5.6 (d). Tests with coating sample 01 showed a maximum Q factor of about 24000 in the adaptor to three alignment mount before an increase to 87700 in the solid one-piece mount. It is not fully clear that this increase was due to the rigid mount, as the sample had to be remounted with the front spring, and these samples were very susceptible to clamping errors. However, all coating test samples and 50 Hz oscillators were tested in rigid mounts after this. With rigid mounts, the sample's Q factors we measured in Section 5.4 were limited by fundamental Q factor limits, showing this clamping method works sufficiently. The ringdown experiment had plenty of degrees of freedom to steer the beam, so the rigid piece was suitable. Some setups require precise alignment with the optics, and attachment to an alignment stage is needed. For such cases, a more focused and in-depth experiment on this interface should be done.

5.2.3   Data Processing

The data had to be processed to enable a smooth and better fit of the ringdown. The first was an application of a bandpass filter, which was done digitally with the data, post acquisition. Low-frequency drifts in the raw data would corrupt the amplitude measurement by adding an additional slope to the oscillation amplitudes. The high frequencies are cut off to prevent transients and other resonances from coupling to the measurement. The bandpass filter uses a passband of 10-1000 Hz. The filter has some inherent oscillations, which typically corrupt the data's first and last few seconds. These are cut off. As we take between 10 and 20 ringdowns in one measurement set, and it is easy to take longer measurements, this is a manageable loss but does make the method unsuitable for measuring high f_0, low Q oscillations, with a ringdown time of less than a couple of seconds.

We only fit the envelope of the data, as when we consider the full sinusoid, it becomes computationally expensive to fit all the parameters. Therefore, a pre-fit routine runs that finds the highest peak of each oscillation and then fits these to equation 5.1. MATLAB's nlinfit routine is used to perform the fit using free parameters for the initial amplitude and Q factor. The resonance frequency was estimated from the highest peak in the ASD. Uncertainty on this value can be estimated as the half width, half maximum on the peak. Due to high Q factors this is typically smaller than the sampling frequency of the Moku:lab phasemeter. Hence, the measurements had an uncertainty of +/-0.01 Hz. As such precision is unnecessary and mounting stresses can change the f_0 by around 1 %, the f_0s here are quoted to three significant figures. The covariance of the fit data was used to estimate the uncertainty on the Q factor.

5.2.4 Ringdown Results

The samples from batch 1 were tested in the ringdown experiment and gave the results in the top six lines of table 5.1. The highest Q factor measured was 28000, whose ringdown is shown in Figure 5.8. After this measurement, the sample was damaged, probably due to insufficient cooling down periods between PZT excitements. Design 3 was the most promising, so a batch of five additional oscillators was ordered with this design for assembly into the first set of inertial sensors. A change was made to reduce the thickness to 6.35 mm as the standard wafer thickness so that the raw material was easy to acquire. Finally, the material was better specified from generic fused silica to high-grade Corning 7890-0F, due to its high purity and verified low bulk loss measured by [62].

The predicted drum mode oscillation frequencies from FEA and those measured in the ringdown setup are shown in Figure 5.7. The resonance frequency of DE3D3S01 was much lower than the other sample of the same design and also much lower than FEA suggested. Some time was spent trying to analyse the FEA models for accuracy. However, as other samples of the same design had f_0 that matched the expected values, miscommunication likely produced these first samples with the originally specified 50 µm thickness. Likewise, DE2D4S02 was etched with slightly thicker flexures than specified. Another interesting observation is that devices tend not to be spread over the whole possible range but rather have a narrow spread when produced in one set. The prototype batches were produced one at a time with re-calibration of the machine between each. The other designs were all produced in one batch. Once a machine is calibrated, it can produce large batches of consistent samples.

When the second batch arrived, they were immediately tested in the same setup; however, it was found that none of the samples was ringing after a kick. Figure 5.9 shows an example of this. The only known difference other than the small design changes between the first two batches was the transit method from the femtoprint to the AEI. After damage on one of the flexures in the first batch, the second was wrapped in a more protective layer. A glue-like medium was applied to the sheath where they sat to keep them in place. We believe glue dissolved into the flexures and heavily degraded the Q factor. Several attempts were made to clean them at the AEI with either a methanol bath or a mix 50/50 of isopropanol and acetone, where the device was left overnight, but this had no effect. They were shipped back to FEMTOprint, who treated the samples with an HF bath to etch off the top layer. Etching the top layer would have an unknown effect on the surface roughness of the flexures. The samples were then measured again and found to ring once again, as shown in Figure 5.10. The oscillators had a lower Q factor, which should be expected due to the additional treatments and residual contamination.

The oscillators from batch 2 and the remaining samples from batch 1 were then coated. The coating was anti-reflective on one side and highly reflective on the other. The reflectivity of the highly reflective side needed to be better defined as the surface of the oscillators was much rougher than a typical sample for coating. The coating was done using ion beam sputtering with a Niobia/ Fused Silica coating. A mask was applied to the surface of the oscillator, meaning that only a ring of 10 mm in the centre of the oscillator would be coated. A gap of 2.2 mm was left before the flexures to prevent contamination. A photo of the coated region is shown in Figure 5.11, where it
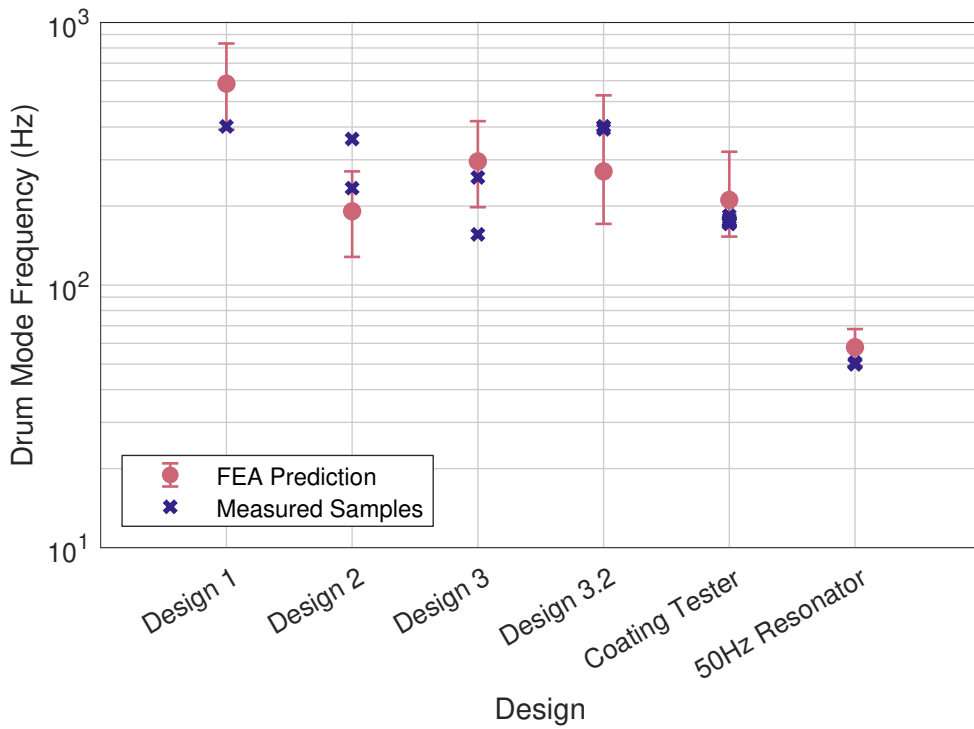
Figure 5.7: A comparison of the resonance frequency of the drum mode of oscillation predicted by FEA and measured in reality for each sample produced in this chapter. The uncertainties on the FEA values are derived by putting the upper and lower limits of average flexure thickness in the FEA model. Almost all samples lie in their expected regions except for DE3D3S01 and DE2D4S02.
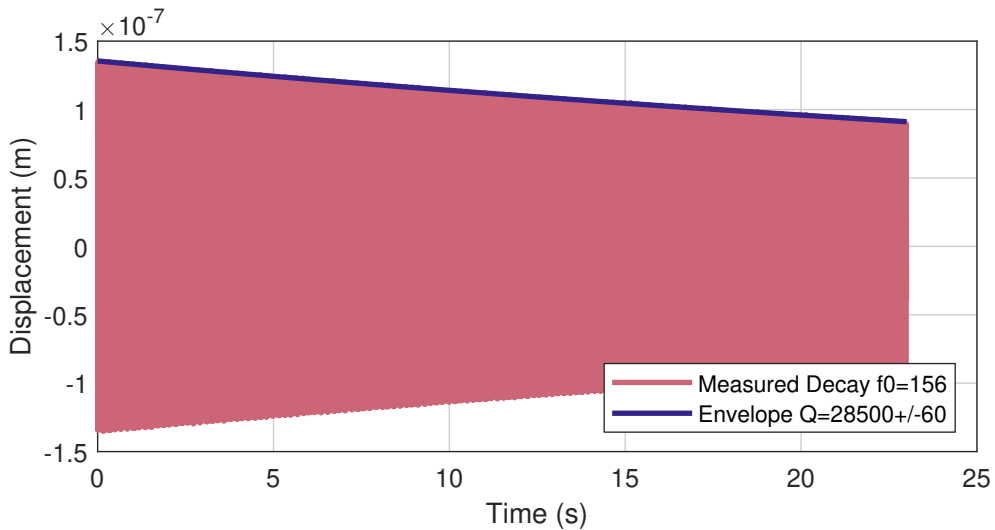

Figure 5.8: The measurement and fitted envelope of the sample DE3D3S01. The sample achieved a Q factor of 28500 with a natural frequency of 156 Hz.

Serial Number	Design	f_0 (Hz)	Best Q factor
Batch 1			
DE1D3S01	Design 1	402	5740±3
DE1D3S02	Design 1	No Ring	...
DE2D4S01	Design 2	234	10800±60
DE2D4S02	Design 2	360	8820±5
DE3D3S01	Design 3	156	28500±60
DE3D3S03	Design 3	257	6210±4
DE3D3Dam	Design 3	261	6980±10
Batch 2			
C2D2S01	Design 3.2	402	4230±6
C2D2S04	Design 3.2	401	5870±20
C2D2S05	Design 3.2	392	6990±6

Table 5.1: The results of the measurements of the test samples for the oscillators before coating show the natural frequency and highest Q factor measured. Design 3 proved most promising, as DE3D3S02 was only tested with the side clamp before it was damaged, but the Q factor that was measured was the limiting Q factor of this clamp. The second batch had very similar results, implying a common limiting factor to this batch. We suspected the HF bath needed to remove the residue, leaving a much greater surface roughness and residual contamination.
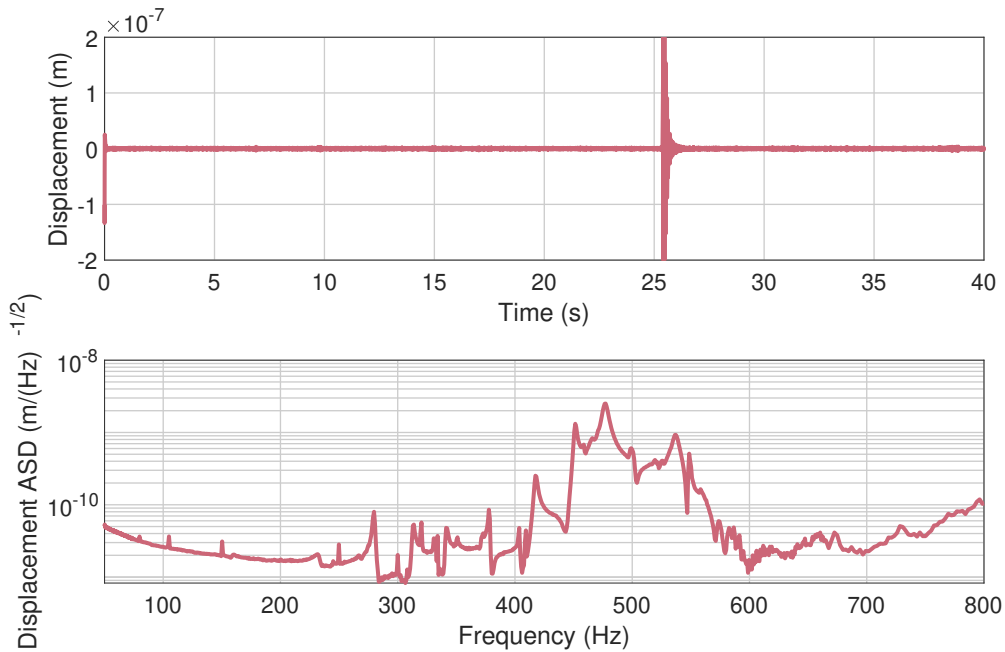
Figure 5.9: The measurements of sample C2D2S05 after production. It is representative of all samples in the batch. The time series of an excitement and a ASD are shown. No discernible ring can be seen in either series or spectral density.
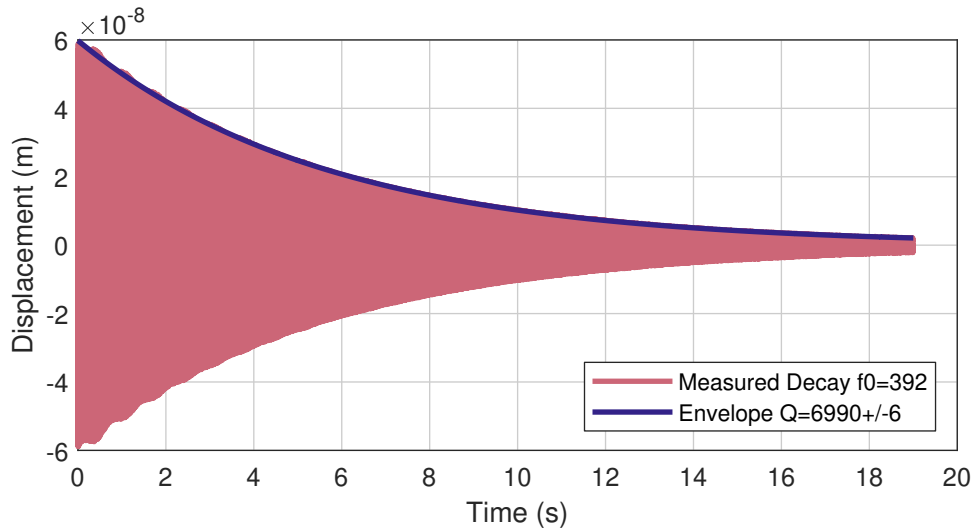
Figure 5.10: The Ringdown of Device C2D2S05 after treatment in a HF bath. A f_0 of 392 Hz was measured. The measurement took place at 1×10^{-2} mBar. The oscillation was restored when compared with Figure 5.9.

can be seen that the coating was contained to a region of 10 mm diameter. All coated samples were tested and showed the same result of a degraded Q factor of about 500. The results of one sample are shown in Figure 5.12, where a Q of 500 was measured. It was unknown exactly what led to this degradation in Q factor. Two hypotheses were proposed: the first was that the mask was insufficient, and the coating had escaped the masked region. The other was that the surface layer applied a strain to the surface and that this surface strain acted as an excess loss channel. The sample was sent for evaluation by a Energy-Dispersive X-ray spectroscopy (EDX) technique with kind help from the Laser Zentrum Hannover (LZH). EDX techniques can identify elemental compositions of samples in specific locations. A sample was tested for contamination of the coating material in 4 regions, shown in Figure 5.13. The tests showed the expected Niobium peak in the central region. However, no traces of Niobium were found in any of the other regions. Small amounts of organic contaminants were found in the Inner Anchor region. It is still being determined where these came from, but they are unlikely to originate from the clean environment in which they were coated. This left the troublesome problem that it was likely surface-induced stresses that caused the degradation in Q factor, requiring another solution for the coating.

Regardless, the oscillators were sent to the Fraunhofer-Institut für Angewandte Optik und Feinmechanik (Jena) (IOF) for final assembly with the spacer and in the coupling mirror as there was still some ring and a limited manufacturing time window. Meanwhile, an experiment was conducted to test different methods of coating oscillators.

5.3 Coating Test

An experiment was envisioned to test different methods of coating optics and see if any kept their high Q factor after coating. To conduct this test, an oscillator was designed
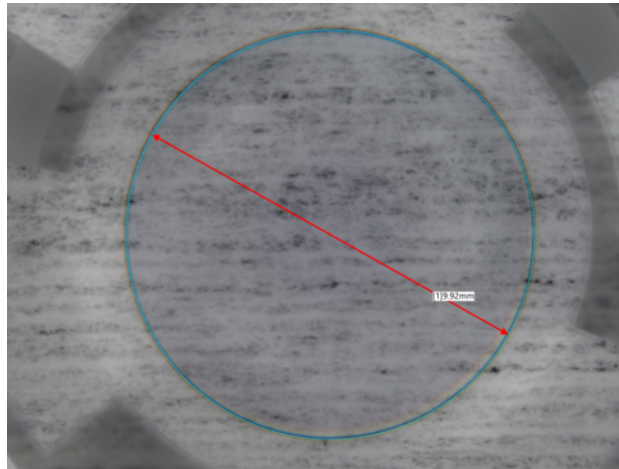
Figure 5.11: A photo taken with a microscope showing the coated region of the oscillator. The region is approximately in a 10 mm diameter ring, as was specified.
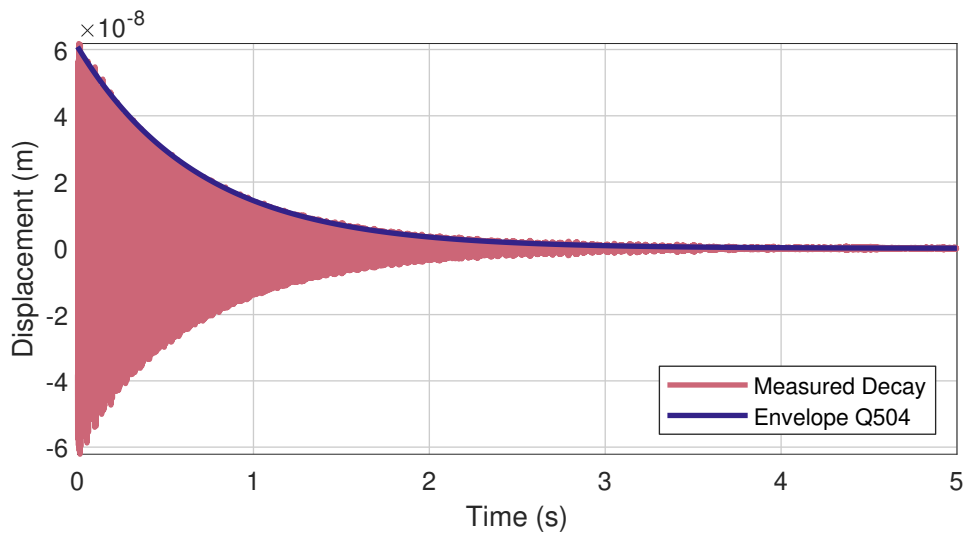

Figure 5.12: Ringdown of device C2D2S01 after Coating. The Q factor degraded by over an order of magnitude. As it was not exceptionally high before, this was incredibly surprising. It implies that coating losses were much higher than expected. The resonance frequencies of the devices were unaffected by coating.
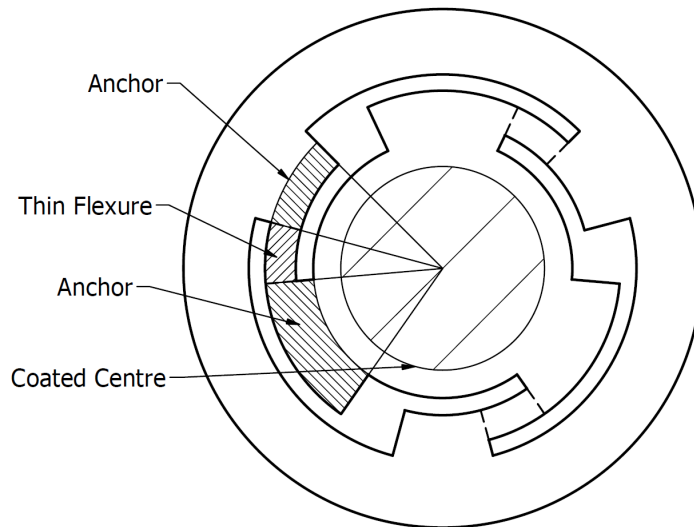
Figure 5.13: The different regions coated by the LZH on the coated sample with EDX. The centre was chosen as a test, while the other regions were chosen as the most sensitive to contamination. Niobium was found as expected in the coated centre region. No traces were found in either anchor or flexure region. Small amounts of organic contaminants were found in one of the anchor regions.

with a similar geometry to the oscillator used in the main devices while being cheap enough to produce in a large enough batch for testing. Uncoated samples would first be produced, and then different coating methods would be applied. Two main expensive features were identified that were not needed in the coating testers: the mass and the isolation from higher order modes. As the cost of the device is linearly correlated with the material removed, a reduction in sample thickness to 1.5mm would lead to a large reduction in costs. As the parallelogram flexure geometry proved difficult to produce, one from each pair was removed. This meant that the separation from the higher order tip mode was only a factor of two but increased the yield of the devices (the whole batch came out without visible defect, as opposed to the 20% failure rate in previous batches). The design is shown in Figure 5.14. It had a fundamental resonance of 208 Hz and a second tip-tilt mode at 512 Hz, which are shown in Figure 5.15. The expected Q factor for these devices was calculated as limited by a combination of TED losses and Material Losses. The TED losses were expected to be around 4×10^5 and material losses around 1×10^6 for a total Q factor of around 3.5×10^5.

5.3.1   Pre-Coating measurements

The full batch of samples was tested before coating so that the effect of the coating method could be distinguished. They were tested in the same ringdown setup as before with an adapted clamp for this size. The clamp used was a spring-loaded front clamp on a rigid mount shown in Figure 5.6(d). Table 5.2 shows the results.
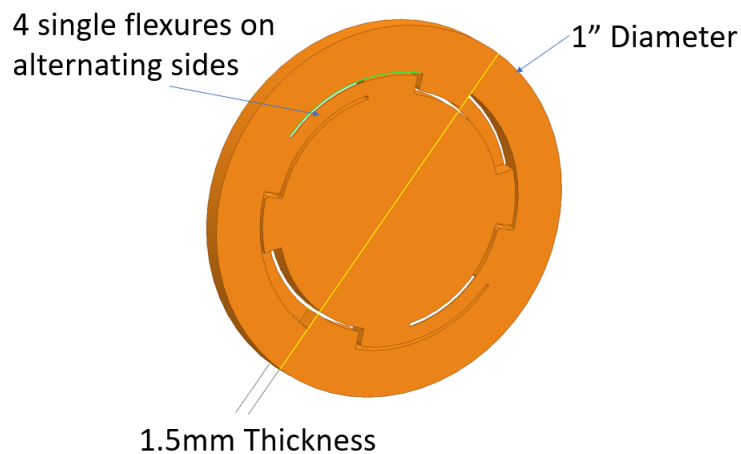
Figure 5.14: Geometry and key design elements of the coating tester sample. Three elements were changed from the main inertial sensor design to allow for a large batch to be produced. The flexures were changed to an alternating pattern so that all areas of the sample could be reached more easily, as this proved problematic in the past. The thickness was reduced to 1.5mm so that less material was removed. In order to keep the resonance frequency in the same range as the prototype oscillators, the geometry was changed to longer, less tall flexures.

(a) 208 Hz Fundamental Mode   (b) 512 Hz Tip-Tilt Mode
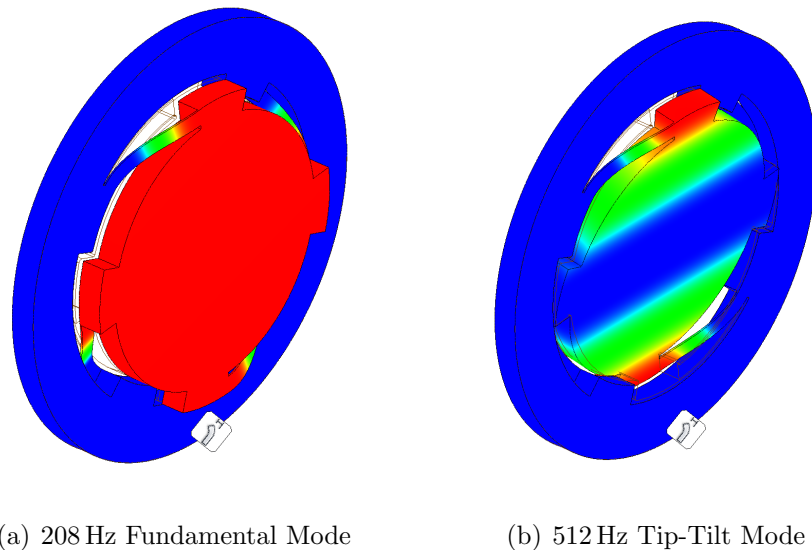
Figure 5.15: Fundamental and tip-tilt modes of the coating tester. The gap between the first and second modes is much smaller than in the standard oscillators. The tip-tilt mode is of concern when used as an inertial sensor. Exciting that mode would lead to beam pointing changes on reflection, coupling to position measurement errors. The scale of this problem would depend on the activity of the environment.

Across the whole batch resonance frequencies lower than the target were measured. All measurements were within a range predicted from tolerances quoted on average flexure thickness. Samples 01 and 02 were made as a tester batch and the last ten as one whole batch, which may account for the small difference in resonance frequencies.

The Q factors had a large spread, from 2000 to 148000. Mounting may have a large impact on the sample Q factor; screw tightness, the orientation of the clamp with respect to the flexures and the centring of the oscillator in the mount all likely have some influence on this. Sample 07 had a Q factor of only 10000 in its first measurement, and upon inspection, it was found that it was improperly centred and that the screws were slightly loose. Fixing both of these led to an improvement to the measured 37200. While this would imply clamping losses still dominated, an unusual effect was seen on the ASDs of the ringdowns in the sample. An ASD of the sample in the first mounting and optimal mounting is shown in Figure 5.16. When this sample was not mounted well, a second peak was present in the ASD. This is likely a peak created by rubbing against the clamp. Remounting the oscillators correctly removes this peak. A single peak can be used as a witness that the clamping is sufficient.

Sample 09 had a Q factor of only 2000. Multiple remountings yielded no significant changes in measured Q factor. Multiple peaks were seen in the spectrum that were not removed by reclamping. The sample appeared free of visual defects. It is unknown why this sample behaved as such, but was discounted as a sample for coating.

The range in Q factors in the samples is likely due to local contaminants on the flexure regions. The best-performing oscillators were beginning to hit the limit predicted by TED and material loss calculations. The entire batch performed better than the samples in the first two batches produced for the sensors. This was surprising as it was assumed that the lighter design would lead to greater frictional loss, but perhaps this is compensated for with a decreased moving mass. It is also possible that the dual flexures were acting against each other. However, the models have not predicted this effect. Furthermore, as a third party manufactures them, it could be that the optical handling and production methods have improved over time.

Regardless a batch of oscillators that can evaluate the effect of the coating method on Q factor was ready for testing. Three methods of coating were to be applied. A IBS coating with a smaller diameter, metallic coating, and attaching a mirror with adhesive glue. Only the latter two have so far been done.

5.3.2   Mirror Bonding

The effect of using Ultra Violet cured glue (UV glue) (Specifically Optocast 3553-LV-UTF) as an adhesive was tested. To do this, several mirrors with a highly reflective coating were ordered. The mirrors had a diameter of 7.75mm and a thickness of 2mm. They had a total mass of 0.2 g, which would increase the oscillating mass by about 28%. This should then take the resonance frequency down to about 155 Hz . These were bonded to samples 1,4,7, and 9. 9 was used as a test of the procedure to check the tool and show that the method did not crack the flexures.

Figure 5.17 shows a photo of the bonding tool. The bonding process used the following procedure.

- The white plastic mount was made parallel to the ground with the 4 axis mount

Sample number	f_0 (Hz)	Best Q factor
In Coming Test		
01	179	87700± 2000
02	184	24400±80
03	175	31000±100
04	173	118000±2000
05	175	11300±100
06	172	70700±700
07	173	37200±100
08	171	148000±600
09	176	2000±100
10	172	121000±80
11	174	39100±10
12	175	24500±300
Mirror Bonded (UV Glue)		
01	158	303000± 2000
04	151	322000±2000
07	136	84500±300
Metallic Coating		
06	172	198000± 1000
08	155	92100±700
12	174	156000±500

Table 5.2: The initial results of the measurements of the coating test samples, before any coating was applied to them. The corresponding graphs are shown in Appendix Figure C.3.
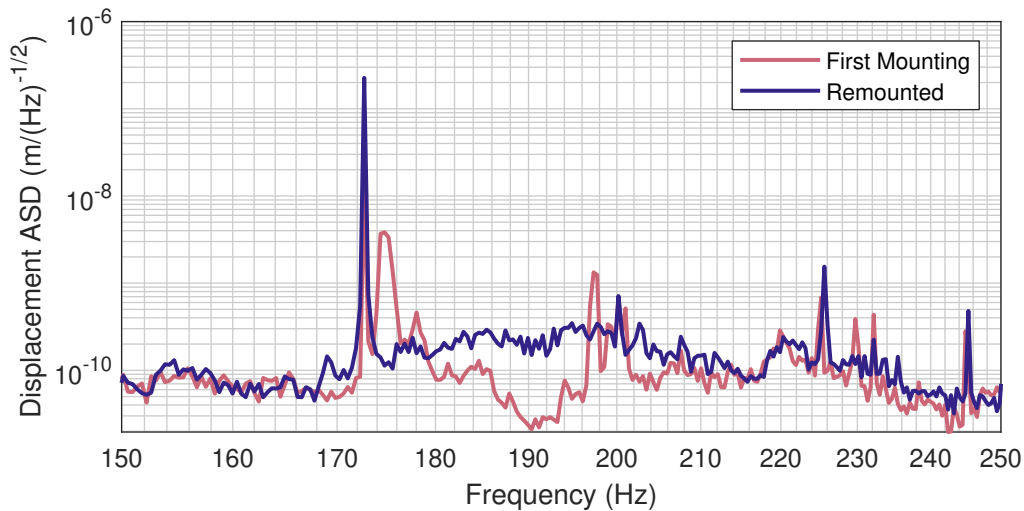
Figure 5.16: A comparison of the ASDs of one sample with the clamp slightly re-positioned and tightened. The initial measurement produced a Q factor of 10000, but by reclamping a Q factor of 37200 was achieved. A clear difference in the two ASDs can be seen, where the bad clamp creates a second off-resonance peak. This peak is likely the oscillator rubbing against the clamp. The removal of the second peak provides a witness for when the sample is clamped properly.

    using a ball bearing as a spirit level. Being level stopped the glue from sliding to one side of the mirror before curing.

- The mirror sat in the hole on the white plastic piece.

- The bonding surfaces were cleaned. A drop of isopropanol was placed on a lens tissue. The wet spot of the tissue was placed on the area to be bonded; then, the tissue was dragged to the dry area, much the same way an optical fibre is typically cleaned.

- The oscillator holder was then mounted in the rail shown in the bonding tool such that it could slide down with the vertical linear stage and make contact with the mirror.

- The linear and rotation stages were used to align the oscillator with the two mounts and centre the mirror on the oscillator. This is best done in a well-illuminated environment to see the edges clearly.

- The oscillator's linear stage is raised to its maximum height, and a very small amount of the UV glue was applied to the mirror. The flow rate from the syringe is very high, so it was wiped on a dry cloth several times after a very gentle squeeze.

- The oscillator was lowered back into position such that it just made contact and the UV glue evenly spread over the surface.

- A UV lamp was used to cure the glue. This usually took less than 1 minute under direct light.
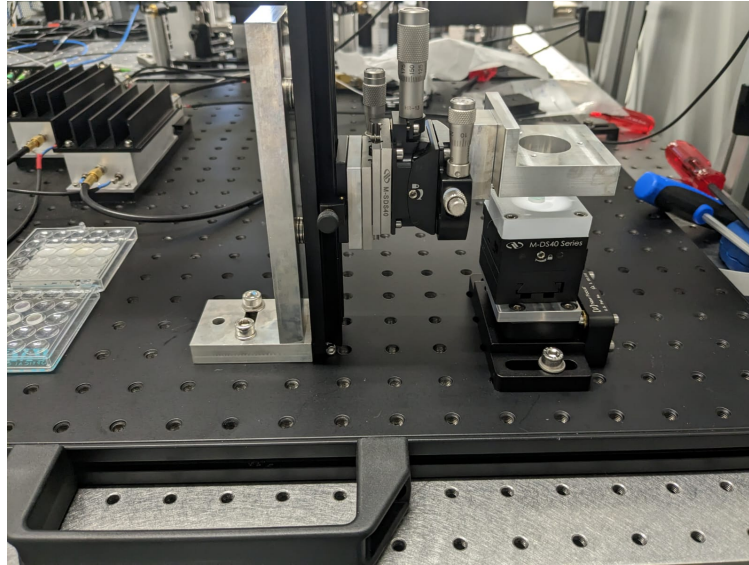
Figure 5.17: The bonding tool used to attach a mirror to the oscillator. The mirror sat in the white plastic piece, while the oscillator was held in the aluminium mirror mount. The various linear and rotation stages give full position control, enabling alignment and positioning between the surfaces.

No flexures were damaged in 4 bonded samples, and all adhered successfully. Large uncertainties on the centring of the mirror were caused as the alignment was done by eye. This would be problematic for inertial sensors as a miscentring would lead to a bias for one side. To bond an inertial sensor, a tool with defined tolerances would be required; however, these samples are a suitable test of the Q factor from the bond. Figure 5.18 shows the results of this process.

In total, three samples were bonded with mirrors. Samples 01, 04, and 07 were mounted and tested. The results after bonding are shown in Table 5.2 and the Graphs in Appendix C.4. Several surprising results were seen. First, each sample's Q factor increased from bonding. A possible cause of this effect is shown in Figure 5.19. After bonding, the samples had Q factors very close to their theoretical maximums as predicted during design. As the resonance frequency is decreased with the additional mass added, the TED losses will also decrease. This demonstrates a high likelihood that the devices are TED limited. However, neither point matches exactly the values predicted here, and the increase in Q factor is much sharper than would be expected. This highlights that the Zener approximation is insufficient for high-precision analysis and is better used as a loose tool for estimating oscillator behaviour.

The other strange effect was the change in resonance frequency for each sample. As the extra mass from the mirror was the same for each sample, the fundamental natural frequencies were expected to drop by approximately the same ratio. Sample 07's fundamental frequency dropped especially far. The sample was inspected for damage on the flexures from bonding, but none was found. It could be that a larger glue patch was used on this sample, as it was the first bonded, and refinement was still ongoing. This large patch could substantially add to the mass, lowering the fundamental frequency. The other samples experienced a drop in resonance frequency consistent with the extra
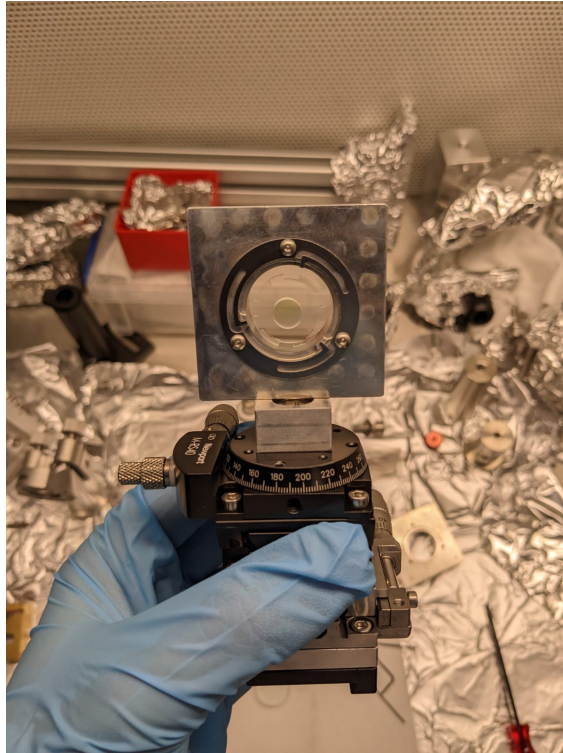
Figure 5.18: The results of a bonded coating tester and mirror. No residue glue has escaped the bonding area and all areas are bonded with no peeling.

mass.

The results from this test were very promising. A Q factor of 300000 was achieved with mirrors at least 4mm away from any flexure region. These Q factors allow for developing inertial sensors with very high performance on 1-inch optics. Furthermore, they push against the limits of what Q factor was expected from the fundamental terms limiting the oscillator. Reaching these Q factors, shows that the clamping solution works when properly applied.

5.3.3 Metallic Coatings

It was further decided to test the application of metallic coatings. While UV adhesive bonding was shown to maintain high Q factors, attaching mirrors runs the risk of misalignment, requires careful handling, and will cause a mismatch in mass distribution. Metallic coatings would allow for the direct application of coatings to the surfaces, preventing these worries. A key trade-off is that the maximum reflectivity of a metallic coating is limited when compared to layered coatings. Typically for $1064\,\mathrm{nm}$ laser light, the best reflectivity is from silver coatings, which can achieve reflectivity of $98\,\%$. Although not as high as other methods, achieving an impedance-matched optical resonator with a finesse of 155 is still possible. As silver tarnishes in air, coatings often have an additional protective layer on top, typically a material that is transparent to the light. In this test, a SiO_2 coating was applied. It was therefore decided to test adding a silver coating with a protective layer to three coating tester samples.

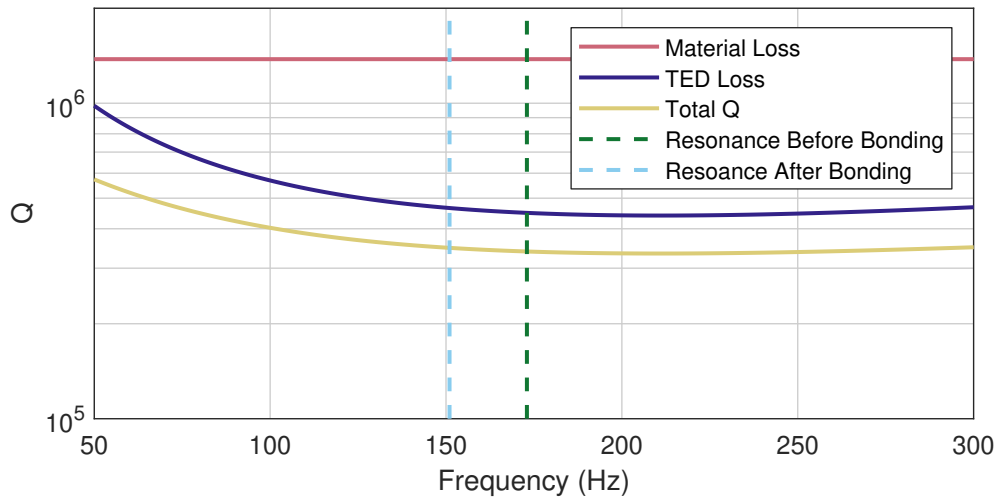Samples 06, 08, and 12 were selected, as they both offered a spread of results from

Figure 5.19: The contributions to Q factor from TED and Material loss as a function of frequency. Two points relating to the resonance frequency of sample 04 before and after bonding with a mirror are highlighted. They show how the TED losses decrease with the decreasing frequency and suggest why an increase in Q factor was measured when the mirror was attached.

the incoming test, and appeared visually clean and free of scratches. The samples were coated using an electron beam evaporation technique [117]. The coating was specified such that a disk a 4 mm diameter in the centre was fully covered to a 98% reflectivity, a disk of 9 mm diameter had an arbitrary amount of coating as a spillover, and outside this 9 mm disk no coating was present. The specification would leave at least 3.8 mm between flexure regions and any residual coating. Visual inspection of each sample showed this to be true.

The samples were then tested in the ringdown setup. Results of the coated samples are shown in Table 5.2, and the rings are shown in Figure C.5. All samples showed that a high Q factor was maintained, but each sample also showed quirks. Samples 06 and 12 behaved in a similar way to each other. Both saw very small decreases in resonance frequency, which can easily be attributed to the extra mass of the coating. Both also saw an increase in Q factor after coating. This strange quirk is attributed to insufficient mounting in initial uncoated measurements.

Further inspection of the ringdown data showed small secondary peaks in both, similar to Figure 5.16. The extra peaks were not noticed until after the samples returned from coating. Sample 08 did see degradation in Q factor, but also a substantial change in f_0. A closer inspection of the sample revealed a fracture along the flexure. It is unknown when this arose, but most likely during shipping. This sample can therefore be discarded.

The results suggest that silver coatings with at least 3.8 mm between coating and flexure do not limit Q factor. Other metallic coatings likely follow similar behaviour. Metallic coatings avoid mass asymmetry and alignment problems that can occur when attempting to bond mirrors to oscillators. It is therefore suggested that if a lower reflectivity can be tolerated, this method can be used to coat optics. However, bonding optics is the advised course of action if a reflectivity greater than 98% is needed.

5.4   50Hz Oscillators

With the prototype oscillators showing both high Q factor oscillations and a method of getting high reflectivity coatings attached, the next step was to produce samples which could match the performance of an L4C geophone in a 1-inch optic. Both thermal noise and displacement noise had to be considered for this. When this sensor was being designed, the results from the frequency reference discussed in Section 6.2 were produced. This gave us an estimate of the currently achievable noise floor for the readout scheme, which could be used to estimate the required natural frequency of oscillation needed to meet the targeted acceleration performance. A frequency of 50 Hz was found to be necessary. Assuming the worst Q factor from the coating tester results of 10000 (ignoring sample 09), a test mass of approximately 3 g was needed to match the performance of the L4C. These oscillators could be combined with the readout technique shown in Chapter 6 to produce a compact inertial sensor capable of meeting an L4C's performance. Work is ongoing to validate the performance of the readout with the 50 Hz oscillators together.

With these mechanical specifications in place, the question was how to design this to be both manufacturable and survivable. The configuration shown in Figure 5.20 was determined to be a viable solution. Here a few things were key to the design. It was decided to use a mirror bonding method to attach a mirror to the optic, as results from the metallic coating were unavailable when the design was made. A cut-out was made on the test mass to help with mirror positioning. The flexures were designed with the parameters shown as a combination that could withstand 5 g of acceleration. The length was chosen using Figure 4.7 as a reference. The substrate thickness (6.35 mm) was chosen as that was available, and flexure height (1 mm) was then tuned to give the correct resonance frequency. The flexure width was chosen considering the studies in Section 4.4.2 as 180 µm. This was chosen as was the optimal thickness for Q factor on resonance at 50 Hz and also gave a good sensitivity across the whole control bandwidth (0.1-100 Hz). To achieve an oscillator that could survive under 5 g of load, the inner edge of the flexure had to be at least 22 mm long. As the flexures were so long, they had to be folded into each other to fit them in a 1" optic. This left island anchors in-between folded flexure regions. Leaving these regions meant less material had to be removed, reducing manufacturing costs substantially. The gap left between the island and the flexures was chosen such that they would not inhibit oscillator motion even when used as a vertical sensor. Manufacturing some internal motion stops with these islands would also be possible. This would likely be necessary if a lower frequency oscillator were desired in such a compact geometry. The regions where the flexures contacted the anchors had 0.1 mm inner radius to distribute the strain in the corners better. The flexures on both sides are offset 90 degrees to allow for room for these anchor structures.

The oscillator's lowest two oscillation modes are shown in Figure 5.21. The fundamental in-out mode occurs at 50 Hz. A similar tip-tilt mode to that of the coating tester was also present, although by making the sample thicker, the relative gap between the tip-tilt and fundamental mode was increased to a factor of four. This is still not as rigid to this mode as the prototype design, which could typically achieve relative ratios of about 20, but it should still be enough to filter out this mode and prevent cross-coupling

1 Inch Diameter

Inlets to Minimise
Etched Material

Cut Out to Seat Mirror

3g Test Mass

Flexures:

180μm Thick
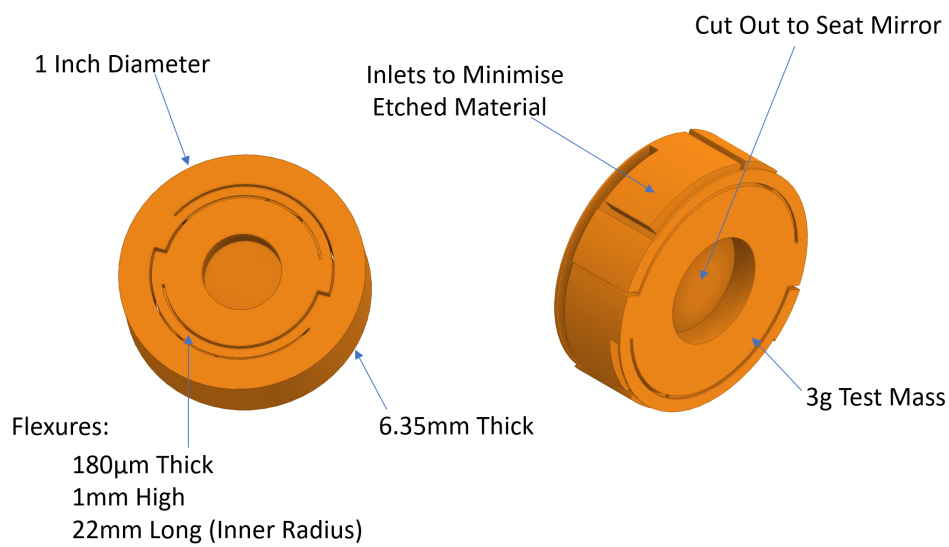1mm High
22mm Long (Inner Radius)

6.35mm Thick

Figure 5.20: An annotated drawing of the key design features of the 50 Hz oscillator. The flexure parameters were designed to make the oscillator withstand 5 g of acceleration while providing a 50 Hz resonance frequency. The second picture shows the inner structure. The material was carefully chosen to minimise the amount etched and prevent the flexures from contacting under maximum load. The inner radius of curvature was set at 0.1 mm on all corners to distribute better the stress localised in these areas.

(a) 50 Hz Fundamental Mode   (b) 209 Hz Tip-Tilt Mode
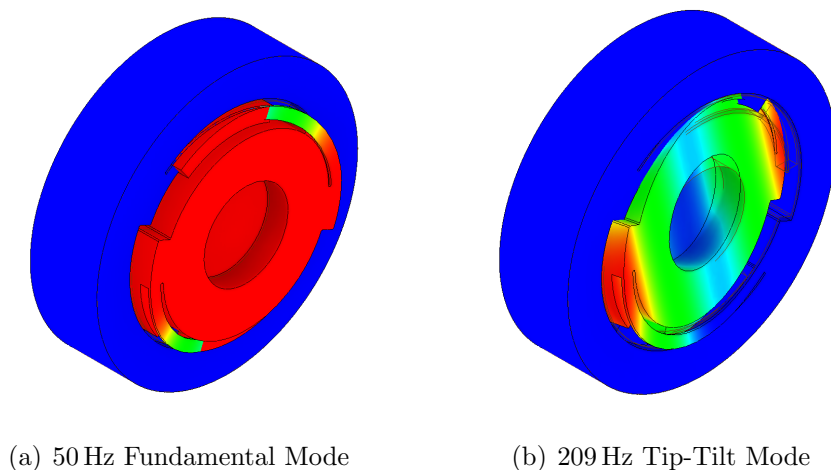
Figure 5.21: Fundamental and tip tilt modes of the modes of the 50 Hz tester.

Serial Number	f_0 (Hz)	Best Q factor
B01	50.2	657000 ± 3000
B02	50.7	621000 ± 3000
B03	49.9	634000 ± 1000

Table 5.3: The ringdown results from the 50 Hz oscillators. The ringdowns these are extracted from are shown in Figure C.6.

between modes.

The expected Q factor from these oscillators is again estimated as having contributions of TED and Material Loss. The TED $Q_\mathrm{TED}{}^{-1}$ on resonance was estimated using Equation 3.19 as 1.5×10^{-6}. The material loss was slightly lower than in previous designs as the flexures in this sample were thicker, leading to an estimated $Q_\mathrm{Mat}{}^{-1}$ of 5×10^{-7} We estimated an expected total Q factor of 520000.

5.4.1 Ringdown Results

Three samples of this design were produced, and the results of their ringdown tests are shown in table 5.3. Each sample had a f_0 close to the target of 50 Hz. However, that target was designed to account for the additional mass from the attached mirror, while these measurements were taken before the mirror was attached. This suggests that the flexure thickness was slightly less than the quoted value but still within tolerance. Each sample had a Q factor of above 600000, above what was predicted. The slightly thinner flexures than the design do not fully explain the discrepancy when only considering TED, it could, however, be consistent with the surfaces of these samples being better than our estimates from section 4.4, suggesting a d_s closer to 5 µm.

Figure 5.22 shows a longer time series ringdown of sample B03. The initial contribution of the tip-tilt mode can be seen decaying in the first 500 seconds of the ring, which is why it does not fit the envelope well there. The presence of these modes can
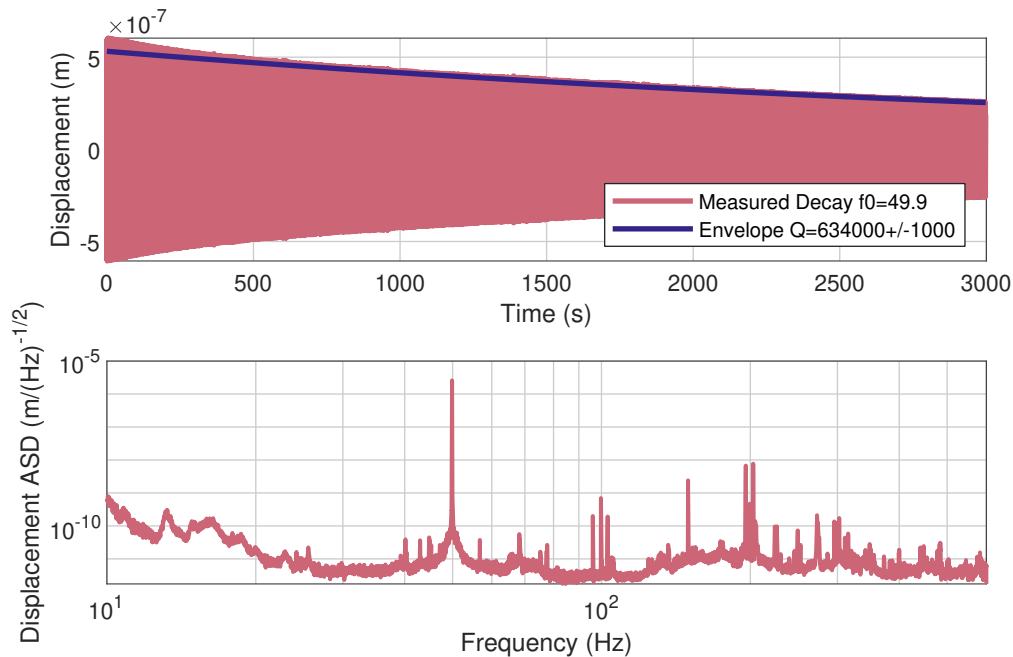
Figure 5.22: (Top) A long term ringdown of sample B03. (Bottom) The ASD of this measurement. The ringdown is plotted with a fitted envelope. The envelope is estimated from just a subsection of the ring, from 1500-2100s, as more than this many data points caused the fitting routine to crash. This section was chosen as the first 600 seconds did not fit this envelope well because there were still contributions to the measurement from the 200 Hz tip-tilt modes, which can be seen in the ASD.

be seen in the ASD as the collection of smaller peaks at 200 Hz. The higher orders of the oscillators had a Q factor of about 184000.

With these measurements, the thermal noise of the device can be estimated and compared to what other sensors available achieve. Figure 5.23 compares the thermal noise performance of these oscillators with the overall performance of other state-of-the-art inertial sensors. It can be seen that these oscillators achieve their primary goal of being competitive with these sensors' performance in the typical control bandwidth of 0.1-100 Hz, even showing a lower possible noise floor than these sensors, particularly at higher frequencies. With the thermal noise hitting the required level, a readout scheme capable of reaching it must be developed, and the progress towards this is discussed in Chapter 6.

5.5   Designing 15Hz Oscillators

The 50 Hz resonator achieves sufficient thermal acceleration noise for a good inertial sensor, but the readout requires an optical cavity, low-noise lasers, excellent controls and other specialised equipment. While this is acceptable for use in optical laboratories, it would be an issue for use outside of specialised physics laboratories. A question is raised as to whether it is possible to design an oscillator that can be used with a much simpler optical readout. The aim here would be to produce something that a plug-and-
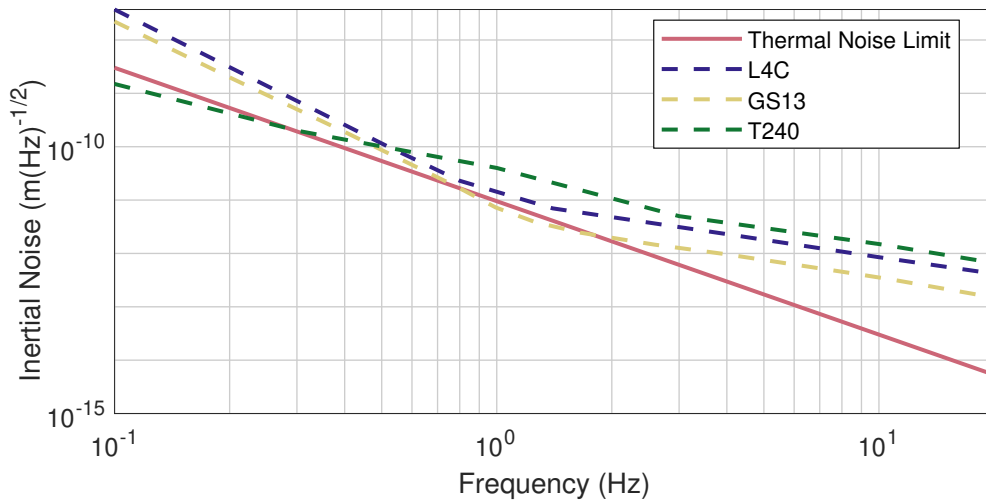
Figure 5.23: A comparison of the projected thermal noise of the 50 Hz oscillator with several state-of-the-art sensors used at LIGO. The thermal noise of the oscillator surpasses or matches all these sensors across most of the relevant control bandwidth for gravitational wave detectors. (Acceleration Equivalent D.2)
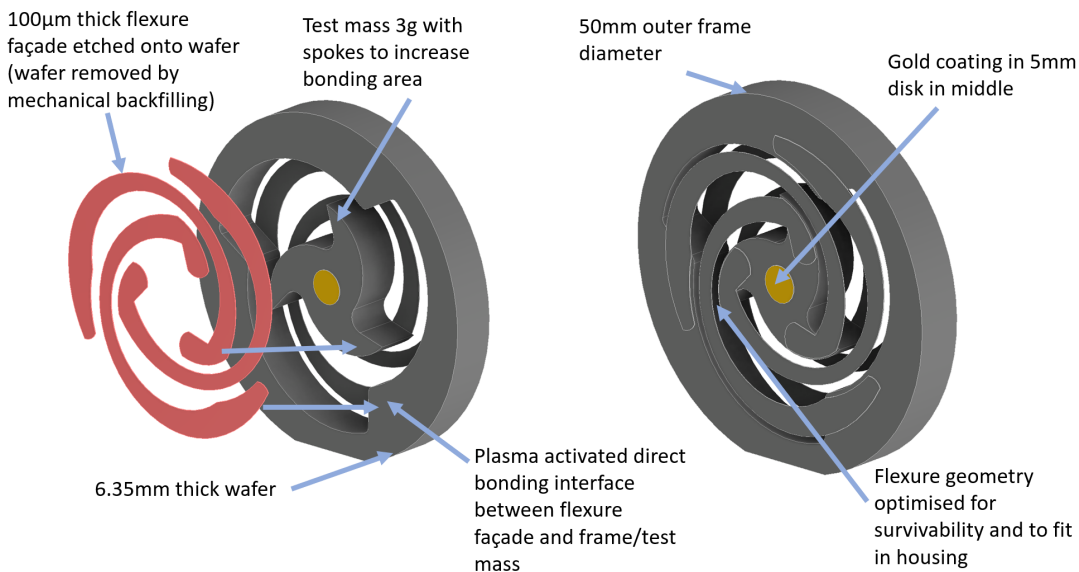

Figure 5.24: The design schematic of the 15 Hz oscillator. The design uses a different manufacturing techniques to others presented so far. The outer frame and test mass are made form one wafer, but make into two pieces by ultra sonic milling. Then the flexure façade is etched into a separate wafer which is then bonded to the outer frame and test mass by plasma activated direct bonding. The wafer the flexures are etched into is removed by a back thinning process.
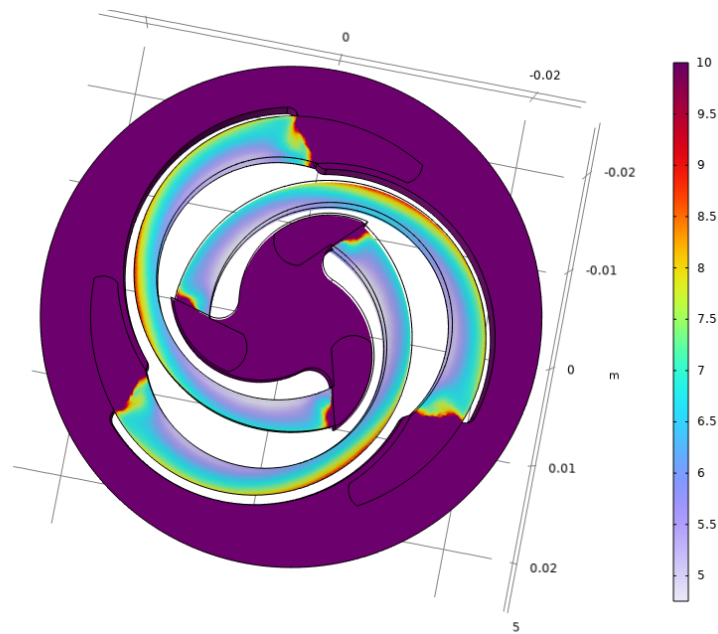
Figure 5.25: Safety factor distribution across the oscillator under 1g of load. Spiral arms allow a longer flexure to fit into a more compact disk, keeping the overall diameter to a 50 mm wafer. Using varying flexure height across its length, the stresses can be best optimised to be balanced across the whole flexure. Doing so produced a 15 Hz oscillator with a safety factor of 5 at 1g.

play interferometer could use to match high-performance seismometers. We show in Figure 4.1 that resonance frequencies of 15 Hz and below are of interest when combined with many compact readout schemes.

Low frequency means higher stress when under load from gravity, which must be accounted for in the design. Figure 4.7 already suggests a flexure length of 40 mm would be necessary. Flexures this long would not fit well in a 1" optic, so this design requirement was lifted. Instead, a 50 mm diameter, 6.35 mm thick disk was chosen as it was also a standard wafer size readily available. Fitting a 15 Hz oscillator on such a wafer still proved difficult. To do this, we had to optimise flexure geometries to better distribute the stress across the whole flexure using results from Section 4.3.1. The resulting design is shown in Figure 5.24. The flexures were folded inwards in a spiral geometry, effectively giving longer flexures in the space available. The height of the flexures varied across their lengths. The flexures were designed as the gap between two non-concentric circles of differing radii. The gap between the centres and radii was then optimised through manual iterative design. The resulting façade and stress distribution of the 15 Hz oscillator is shown in Figure 5.25, while the fundamental and first higher order mode of oscillation are shown in Figure 5.26. Instead of the usual tip-tilt mode, the first order is a bounce mode. As there is an order of magnitude between the two, this should not cause issues. The flexure widths were tuned to 100 µm. In part, this was to help balance the resonance frequency and load on the flexures. Furthermore, as the manufacturing method was to be changed, the surface losses were less of an issue, so the results derived in Section 4.4.2 were not directly applicable.

Conversations with companies about getting this design made tended to be quite

(a) 15 Hz Fundamental Mode   (b) 223 Hz Bounce Mode
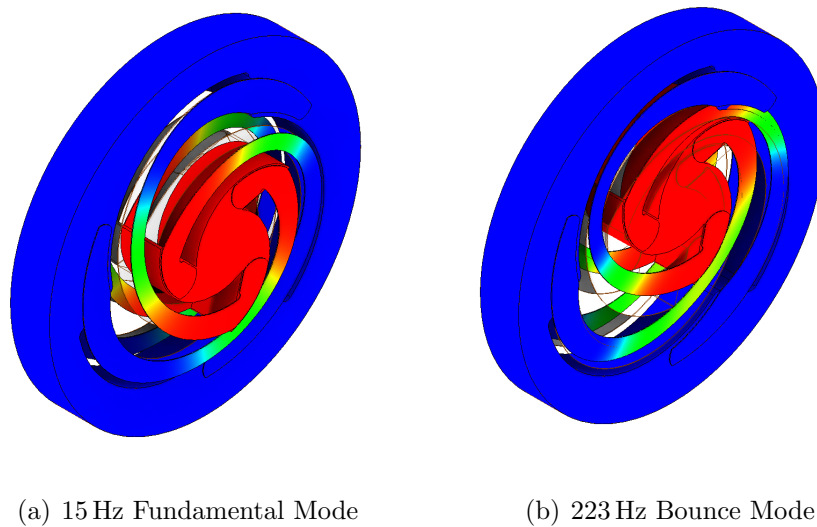
Figure 5.26: Fundamental and bounce modes of the 15 Hz design. Instead of the usual tip-tilt mode, a bounce up and down mode is seen in the flexure at 200 Hz.

brief. With the long flexures, the chance for defects in manufacturing was high. Furthermore, the sheer amount of material that would have to be removed was so high that the etch time would have been problematic as the acid would begin to etch undesired areas. Therefore, an alternative manufacturing process was desired.

The proposed process was jointly developed by the IOF and AEI. Instead of subtractive manufacturing, the test mass and outer frame will be produced from one wafer. Then, the flexure geometry will be imprinted on two wafers, and the result will be bonded by plasma activated direct bonding [118] to the two pieces. The remaining wafer will then be removed with an abrasive polish as a back thinning process, leaving the final structure. Doing the manufacture this way allows for a much better scaling with design size, so in principle, the process could be adapted for even lower frequency oscillators. However, the method also has a direct bond where the flexures intersect with the outer frame and test mass. Furthermore, a filament had to be inserted between the flexures due to the forces in the back thinning. A solvent would then remove the filament. The manufacturing will be conducted at the IOF.

With these suspensions designed, it was possible to fit a 3 g target as a test mass. A Q factor of 3,000 was required to achieve sufficient thermal noise to match the readout performance between 0.1 and 100 Hz. This gives a significant safety margin over results from both the coating test and 50 Hz oscillators, and the theoretical limits of TED with an expected $Q_{\text{TED}} = 2,000,000$. As a high reflectivity is not needed for the readout method, a metallic gold coating has been applied before bonding. The material loss may be considerably lower than in other samples due to the change in manufacturing steps. With the current proposed design process, a highly polished finish can be achieved on both flexure sides. However, this method will also lead to a worked region on the flexure interface with both the test mass and the frame. The direct bond area could act as a lossy area, so some test samples must be produced and tested. A prelimminary test sample has so far been produced at the IOF, a photo of which is shown in Figure 5.27.
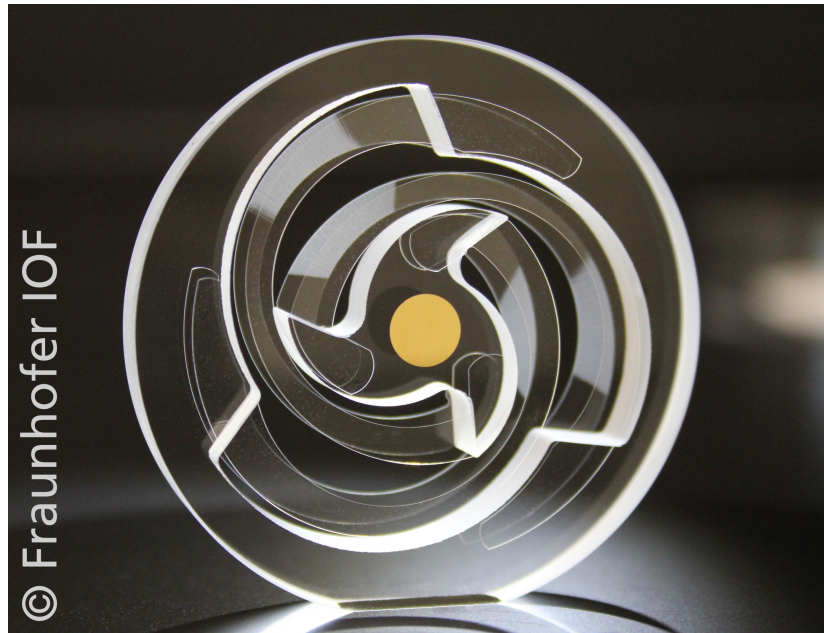
Figure 5.27: A photo of the initial 15 Hz oscillator produced at the IOF. Small residues of the filament can still be seen along the flexures indicating a refinement of the cleaning method is needed. In the preliminary sample, a small misalignment in the flexures causes a small overhang on the outer frame. (Picture Credit: Pascal Birckigt @ IOF)

The initial sample still has some residual contamination

The results of these ongoing assembly tests will be critical for the future of the field. If the process can produce high Q factor oscillators, these designs become very scalable, and a vast design space will be open.

5.6   Chapter Summary

The main aim of this chapter's work was to achieve compact oscillators with excellent thermal noise performance. We have achieved this goal. The 50 Hz oscillators offer a promising look at what is possible in this field, with high-performance devices being possible. The thermal noise of the devices is lower than the noise of most state-of-the-art sensors available today, while fitting inside an optic of 1" diameter and 6.35 mm thick.

Achieving this goal was not easy. Many samples were damaged or never performed well. Some general advice for other groups working on this would be to start with high-frequency oscillators designed to minimise material removed. Learning how to handle oscillators with cheap and more robust samples will reduce the cost of the lessons that must be learned while handling them.

The Q factor destroying effects of IBS coatings were unexpected. Thankfully, alternative methods of achieving high reflectivity coatings without degrading the Q factor were found. Mirror bonding allows for an arbitrary coating reflectivity but does require careful alignment, handling, and positioning. Metallic coatings applied directly to the optic with electron beam evaporation have also been shown not to degrade Q factor

when a gap of at least $3.8\,\mathrm{mm}$ is left between coating and flexure. This allows for an easier-to-apply coating, but can only reach reflectivity of 98%. Between these two methods, a solution for most applications can be found.

The alternative design method discussed in section 5.5 is very important to the future of the field. While the results from the $50\,\mathrm{Hz}$ oscillator are very promising, the need for a sophisticated readout technique limits their application to environments where high-precision optics are already available, such as gravitational wave observatories. The ability to scale the design for use with simpler readout techniques would open a wide range of applications for such oscillators.

Chapter 6

Development of an Optical Readout Scheme

A oscillator with excellent thermal noise performance has been achieved. Now we must find a way of reading out this motion. Due to the high resonance frequency, a very precise readout method was needed. It was therefore decided to use a Pound Drever Hall (PDH) readout scheme on an optical resonator. These schemes require operating in a very narrow operating point, which is difficult to maintain with a high Q factor oscillator. This chapter discusses ongoing developments to the readout scheme and testing overall inertial sensors. While there is still much work to do on the readout scheme it is shown that PDH readouts are a suitable scheme, for meeting the stringent noise requirements. Throughout this chapter I will stick to the earlier established convention of using resonance to describe optical effects, and oscillation to describe mechanical motion.
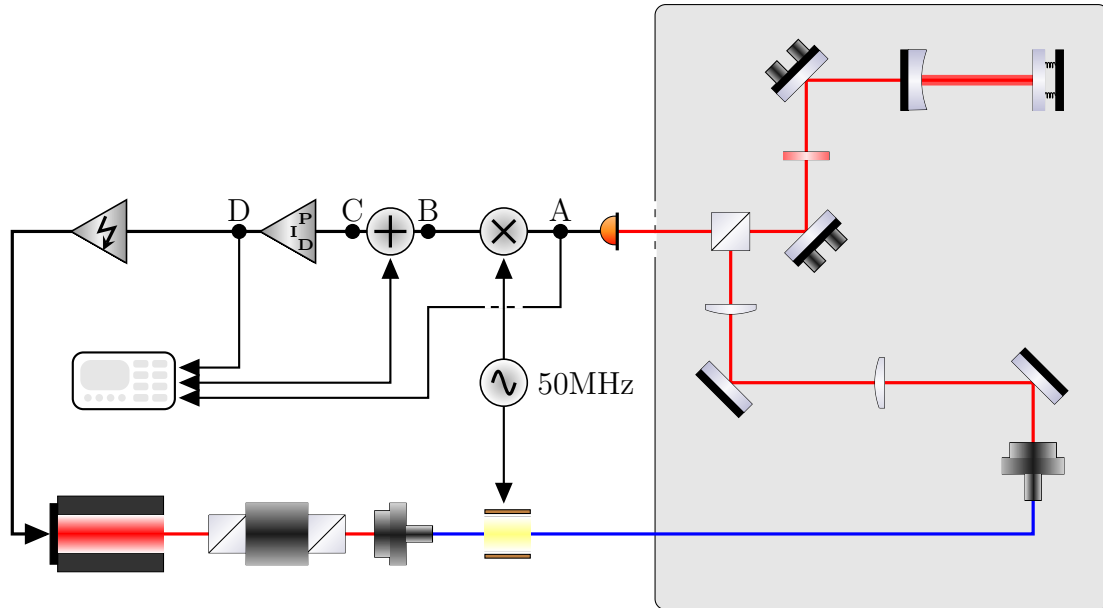
Figure 6.1: The experimental setup of the initial readout scheme. A feedback control loop uses a PDH signal to keep the optical resonator at its operating point by driving the lasers frequency. The side bands are generated with a 50 MHz signal driving an Electro-Optic Modulator (EOM) and a mixer in the photodetector. Fast response is controlled by attenuating the laser piezo, while the long term drifts are controlled by attenuating the laser temperature. The feedback signal is generated by an analogue electronic circuit. The optical setup occurs in vacuum. Two lens focus the beam onto the test mass with the correct beam parameters to match the optical resonator geometry, while piezo driven mirrors can be used to align the beam to the cavity. The photodetector is placed outside the vacuum behind a viewing port. The signal readout is done using a CDS. Four channels are read out: the dc signal from photodetector (A), the demodulated signal before (B) and after an adder (C), and the fast signal feedback (D). The CDS can inject signals into the adder to measure loop response of the system.

6.1   Prototype Optical Readout Scheme

The displacement of the test mass was read out using a Pound Drever Hall (PDH) scheme as shown in Figure 6.1. A Control and Data System (CDS) was used as a data logger but could also inject signals for measuring the transfer function of the control loop. The PID controller was a custom-made analogue electronic used throughout the AEI. The laser temperature and PZT controllers were used to modulate the frequency. The temperature controller was used to compensate for long-term drifts in the lab conditions, while the Piezoelectric (Stack) (PZT) controller was used for fast response controls. Sample C2D2S04 was used as the test sample here, even though the optical coating degraded its Q factor. The prototype aimed to highlight problems with the readout scheme and test electronics while a high Q factor sample was being studied. Without the high Q factor, insufficient thermal noise reduction could be achieved to make a high-performance sensor.
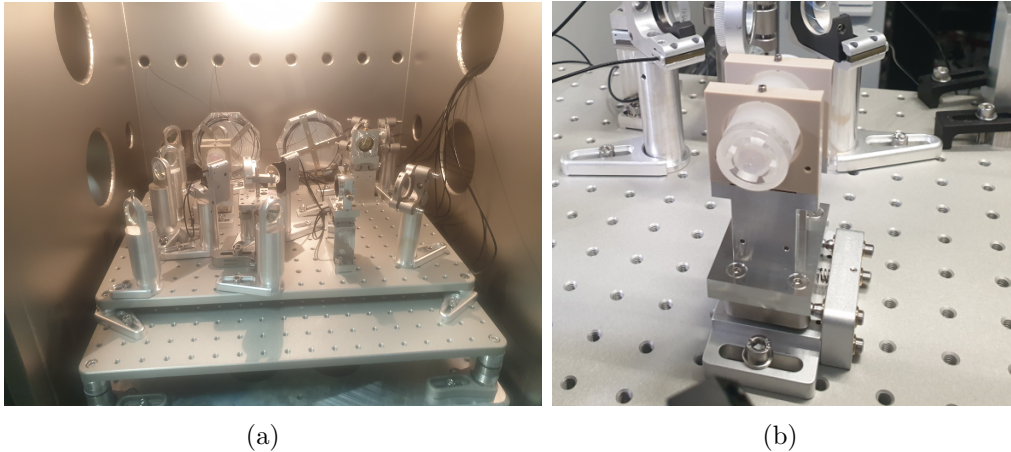
Figure 6.2: a) A photo of the optical readout breadboard with one inertial sensor in the vacuum tank. Readout is done with a photodiode mounted outside the viewport on the left. b) A photo of an inertial sensor in its mount. The mount is affixed to a 4-axis alignment stage to allow for adjustments in positioning. The inertial sensor is held in place by two screws from above with Viton tips to prevent damage to the glass.

6.1.1   Optical Breadboard

Figure 6.2 shows a photo of the breadboard produced. The whole setup was constructed on an aluminium breadboard so that the experiment in a vacuum could be easily switched. The breadboard was placed on Newport VIB100 isolation feet. These feet have a resonance frequency of around 8 Hz, depending on the load, offering excellent isolation at the resonance of the oscillators. For measurements at around 8 Hz, four posts were produced that were slightly longer than the suspensions. These could be placed under the breadboard to act as a mechanical short, rigidly attaching the breadboard to the ground. The cables in the photo belong to piezo-driven mounts for alignment mirrors and linear stages. These were used as it was uncertain if the setup would need readjustment after pumping down to vacuum pressure; this proved not to be a problem, so future design iterations did not include them. We left the cables a little slack to prevent a mechanical short.

The inertial sensor itself was mounted in 4-axis alignment stage to allow for positioning the sensor's central axis along a beam. The sensor sits in a peak mount to reduce the chance of damaging the spacer during tightening. Mounting points were chosen to be at the Airy points [119]. When mounted at Airy points, the end faces of a suspended beam stay parallel. These points were chosen due to the strong coupling of tilt to length changes in the cavity. These points were forund for this geometry using COMSOL Multiphysics in the method discussed by Dovale in his thesis [120].

6.1.2   Optical Resonator Design

The resonator itself consisted of a ULE spacer 5 cm long, giving a free spectral range of 3 GHz. This length was chosen as a compromise between longer cavities introducing more frequency noise, shorter cavities being increasingly difficult to control with
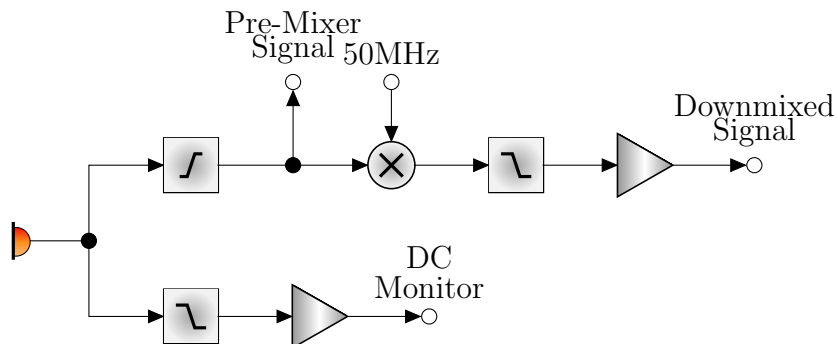
Figure 6.3: A block diagram of how the photodetector with a down mixer works. The photocurrent from the photo diode is converted to a voltage by a Trans Impedance Amplifier (TIA). This signal goes through two paths; one goes through a low pass filter and a gain stage to produce a DC signal on the PD. This is useful for alignment of the PD and also for measuring finesses of the cavities for line width. The other path goes through a high pass filter, and is then mixed with a 50 MHz signal. The result is then low passed to remove higher harmonics, before a gain stage converts the signal to the desired amplitude.

frequency shifts, and with desire for a compact sensor design. Sideband modulation frequencies were chosen at 50 MHz to sufficiently separate them from the central band and produce a clear PDH signal for phase matching. The modulation depth could easily be controlled by changing the voltage applied to the EOM. The modulation depth was therefore chosen such that the peak-to-peak signal of the PDH signal matched the range of the data acquisition system (+/-10 V for the CDS).

The in-coupling mirror was a 1-inch optic. They had an outer flat annulus for bonding, and the inner half-inch was coated for bonding. The coated area was not flat, but rather had a 0.25 m radius of curvature to ensure the resonator's stability as the oscillator had a flat surface. They were coated with a coating specification of 98.5% reflectivity. The oscillators were coated with the same specified coating of 98.5% reflectivity to make the sensor impedance matched. These samples had been treated in a HF bath to etch the top layer away after their contamination, so the surface roughness was unknown. It was unclear if the specified coating could be achieved and the quote was for best effort. The finesse of the resonator was measured as 200 using the peak width of the resonator, which matches well with the theoretical finesse of this resonator of 207, suggesting both mirrors were close to specification.

The cavities were operated with 30 mW of input power. This power was chosen such that neither shot noise nor radiation pressure noise was a limiting factor, and so that the electronic gains were not so high to limit bandwidths. The NPRO laser was operated at full 500 mW operating power so that its noise-eating control loops worked and optimal noise performance was achieved. The desired power was tapped.

6.1.3   Photodetectors

The photodetectors used in this experiment were again custom AEI designs, made by Fabian Meylahn and Henning Vahlbruch. The designs are as yet unpublished and are not mine to publish. However, they incorporate two channels, one a DC channel for
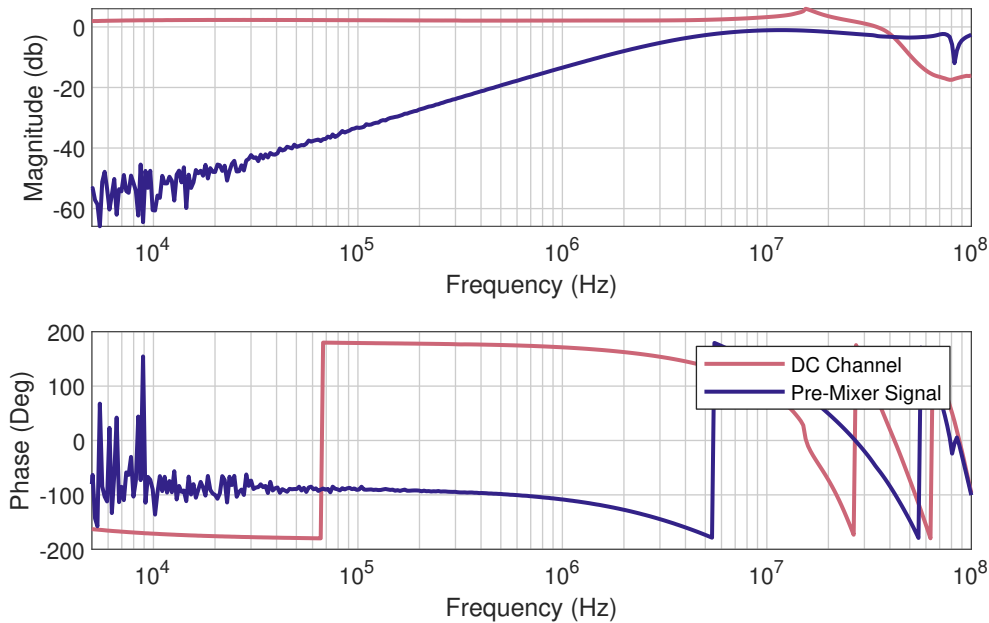
Figure 6.4: Transfer function of the photodetector used for resonator readouts DC and Pre-mixer Channels. The Pre-mixer channel achieves a flat response at the critical frequency of 50 MHz, while also not introducing any significant resonances, which could down couple into the demodulated signal or lead to instability.

monitoring power levels and the other a demodulated channel with an integrated mixer and corresponding low and high pass filters. Figure 6.3 shows a block diagram of how the photodetector works. An InGaAs 1000 photodiode is used for detection, with a high sensitivity Op-Amp used as a trans-impedance amplifier to achieve a bandwidth sufficient for detecting the 50MHz beat note of the reflected signal. The transfer function of the DC and Pre-Mixer channels is shown in Figure 6.4. This was measured by sweeping over a frequency applied to a laser. A flat response and sufficient gain of the Pre-Mixer signal are seen between 10 MHz and 80 MHz. The down-mixed channel response was measured by injecting a signal directly after the mixer and measuring the outputs; the result is shown in Figure 6.5. The result shows a pole at about 8 kHz, which allows for a control bandwidth of about 20 kHz before phase delay becomes problematic. The dark noise of the photodetectors channels is shown in Figure 6.6, which shows that the dark noise of the photodetectors is below that of the Analogue to Digital Converter (ADC) Noise of the CDS.

6.1.4 Control Electronics

The sensors were controlled by custom in-house electronics designed by Fabian Mey-lahn. The electronics were highly customisable with the base design shown in a block diagram in Figure 6.7. The servo control has two outputs: a slow control loop and a fast control loop. The slow control loop can act on slow controllers, such as the laser temperature controller. This channel was only used on the prototype sensor and frequency reference resonator discussed in Figure 6.2. The fast controller has three successive PID controllers. In practice, only proportional/integrative controllers were used. In
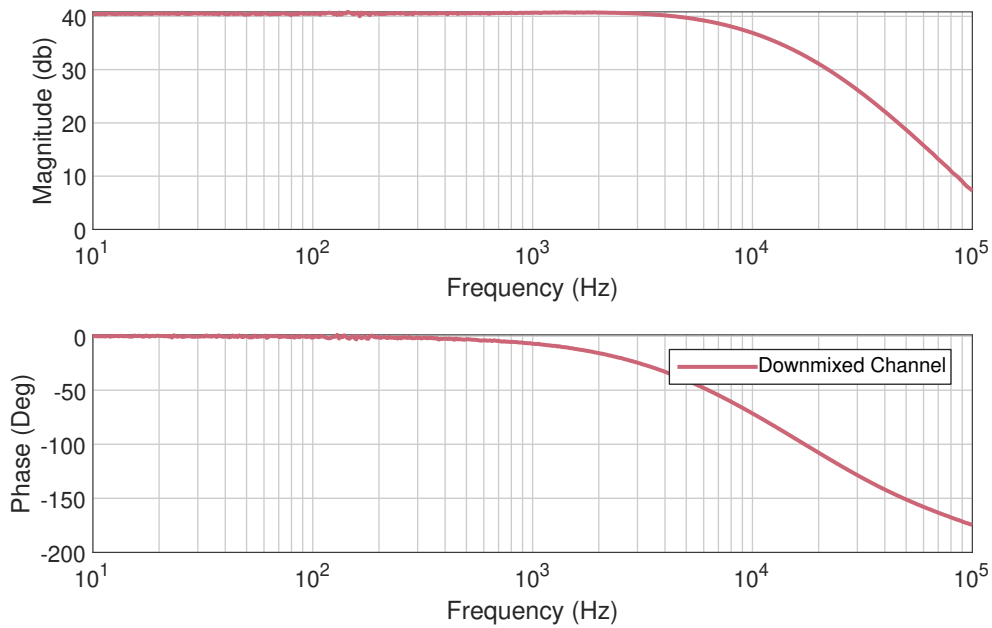
Figure 6.5: The transfer function of the down mixed channel that provides the error signal of the resonator. The response of the PD has a single pole at around 8 kHz, which easily covers the bandwidth of interest for these sensors. The phase lag of the scheme allowed for a single pole integrator to have sufficient phase margin for a unity gain frequency of about 20 kHz.
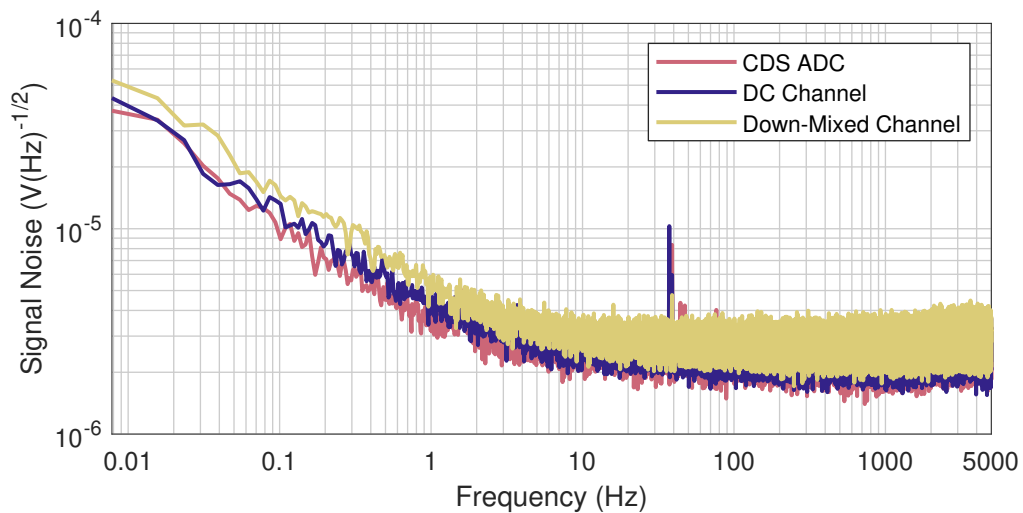

Figure 6.6: The dark noise of the two channels in the photodetectors sued for resonator readout. They were measured with a CDS, whose ADC noise is also shown. The dark noise of the photodetectors measured follows the CDS ADC noise well, suggesting this is the limit for the measurement and that the photodetector dark noise is not an issue of concern for this experiment.
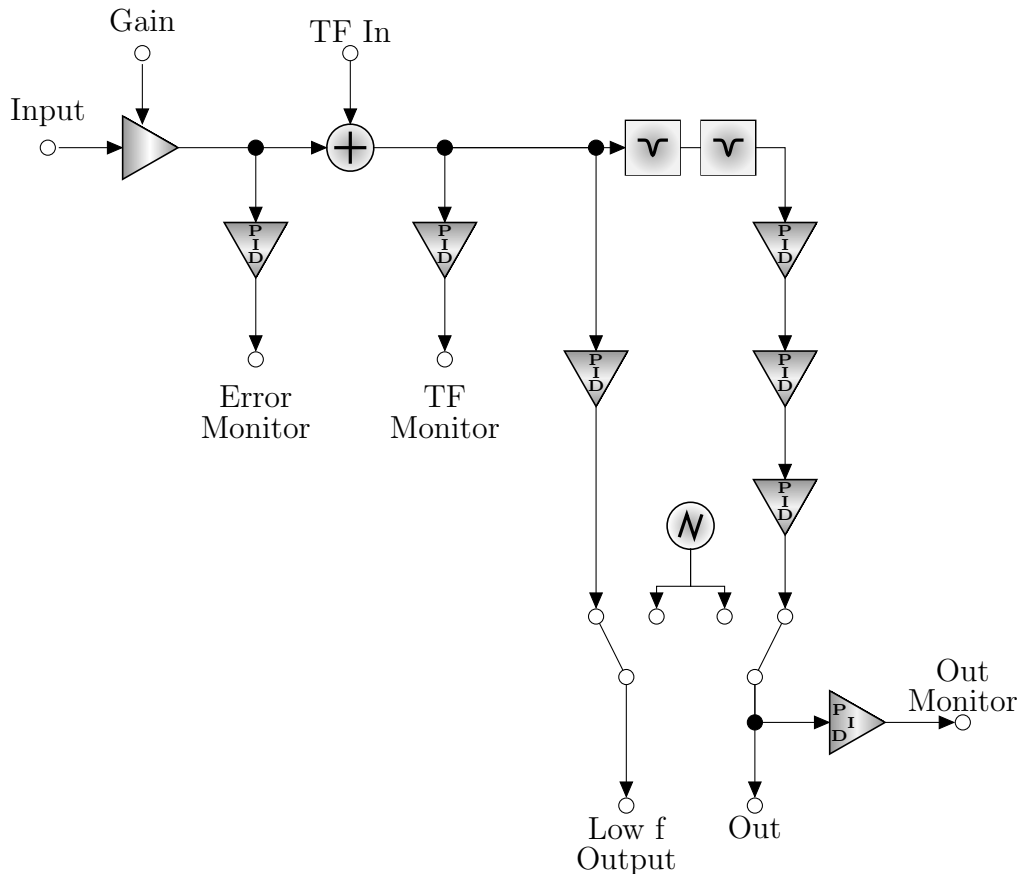
Figure 6.7: A block diagram of the control electronics. The control has two outputs, one for a slow control loop and one for a fast control loop. The fast controller has three successive PID controllers. In practice only proportional/integrative controllers were used. The TF input can be used to add small small signals for measuring the transfer function of the loops. Every output monitor had a PID controller that could be used to optimise readout signals.

the Prototype sensor and the Frequency reference resonator, only one of the stages was used, and the others were bypassed. In the 2 OIS sensors discussed in Section 6.3, two of the stages were used for two poles in the controller as good control was needed at low frequency to suppress ground motion, but trying to drive the controller to compensate for the free oscillator motion close to f_0 lead to problems with the PZT. Notch filters were used to suppress motion from higher-order modes in the oscillators.

The electronics also had several small quality-of-life additions that made working with them much easier. Each output had a PID stage typically used as a purely proportional gain to optimise the signal readouts for dark noise without amplifiers or attenuators. The piece also had the option to add a second signal. This is intended to be used for measuring loop transfer functions. Finally, either output could be switched to a ramp. This allowed for the scanning of the resonator and calibration of the sensors.

6.1.5 Calibrating a Self-Calibrating Sensor

Mechanical inertial sensors with optical readouts are often described as being self-calibrating [61, 121]. This does not mean that the sensor does not require calibration, but rather that it is possible to calibrate the sensor without reference to another sensor. By calibrating using only sensor signals and fundamental constants, the sensor should be immune to systematic calibration errors. There are several ways that the voltage output of the error signal can be converted to a calibrated displacement; the one detailed here was found to be easiest in practice.

In order to calibrate the sensor, a few things are needed; the mechanical transfer function of the oscillator, the wavelength of the laser light, λ, the length of the resonator, L, and a scan of the PDH signal used to generate the error signal of the resonator. The latter of these takes a form similar to that in Figure 6.9. The two factors of relevance are the distance between 0V crossings, $\Delta V_{\mathrm{PDH}=0}$, and the gradient at the central 0 crossings, $d\mathrm{PDH}/dV$, where V is the scan voltage.

From here, the first calibration factor can be derived. This is the difference in test mass position measured per volt applied to the laser. The distance $\Delta V_{\mathrm{PDH}=0}$ gives the distance between the resonance of the sidebands and the main band. This directly correlates to frequency as the crossing distance equals the modulation frequency of the sidebands, f_{sb}. Now it must be converted to distance using the free spectral range giving the calibration factor

$$\mathcal{C}_1 = \frac{f_{\mathrm{sb}}}{f_{\mathrm{FSR}}} \frac{\lambda}{\Delta V_{\mathrm{PDH}=0}}. \tag{6.1}$$

The sideband frequency was the reference frequency used to calibrate the fast feedback channel. This factor can also be calculated using a full Free Spectral Range (FSR), but this may not be possible when driving the laser piezo with a cm scale resonator and a NPRO laser. However, the full FSR was used to calibrate the laser temperature controller. Typically it is a good idea to use the largest reference frequency for the scan possible and average over multiple scans to provide a more accurate measurement.

The gradient of the PDH signal gives the second calibration factor at the 0 crossing

$$\mathcal{C}_2 = \frac{d\mathrm{PDH}}{dV}|_{\mathrm{PDH}=0}. \tag{6.2}$$

These two factors can then be combined to give the appropriate displacement of the test mass using the error signal ϵ, and feedback signal, β. These are given by

$$\Delta x(t) = \frac{\mathcal{C}_1}{\mathcal{C}_2}\epsilon(t) + \mathcal{C}_1\beta(t) * \eta(t), \tag{6.3}$$

where $*$ denotes the linear convolution, $\eta(t)$ is the controller response normalised to its response at the scan frequency. To measure this, a simple setup was devised. A resonator without an oscillator was set up and aligned, but a mechanical PZT stack was behind the in-coupling mirror. Applying a voltage here would drive the mirror forwards and backward and so could be used to control the resonator. The resonator was then controlled with the mechanical PZT, and the laser had a swept sine applied. The response of the control loop could then be used to determine the frequency response of the laser. This is shown in Figure 6.8, which shows a flat laser response in terms of
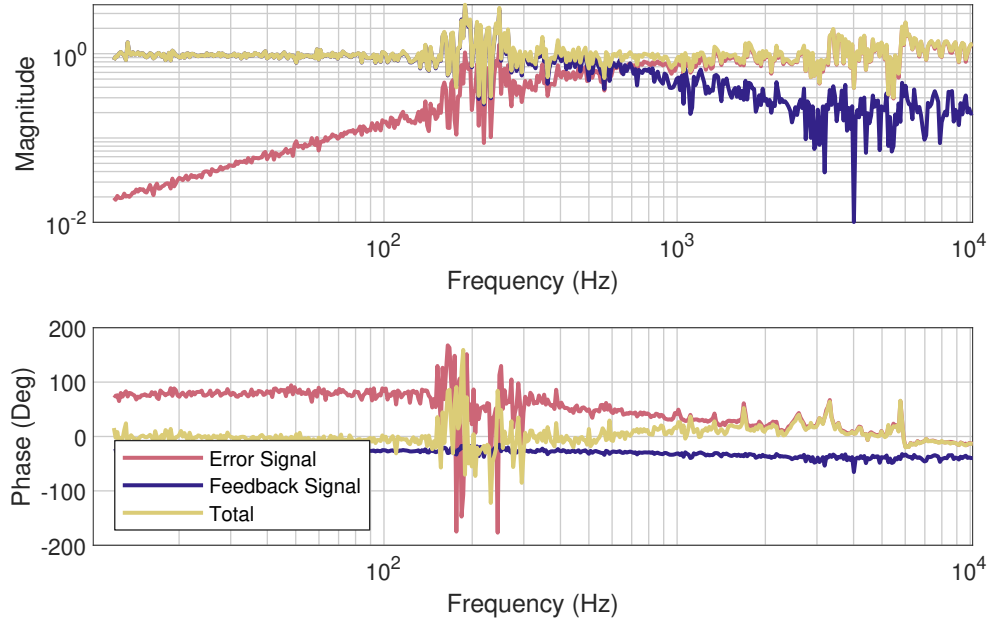
Figure 6.8: Calibration scan of the laser PZT fast response. The calibration has been normalised to its response at 20 Hz to reflect the shape η would look in the calibration. This can be converted into MHz/V by scaling the total by 1.8 MHz/V, or resonator length change with the conversion factor 0.6 nm/V.

magnitude and phase response across the whole region of interest. The actual magnitude of the frequency response was about 1.8 MHz/V, which is in the region as stated in the laser manual (which gives a value ≥ 1 MHz/V).

Equivalently the calibration can be achieved using the open loop response γ of the control loop

$$\Delta x(t) = \frac{\mathcal{C}_1}{\mathcal{C}_2}(\epsilon(t) + \epsilon(t) * \gamma(t)). \tag{6.4}$$

The motion of the test mass must then be converted to an inertial equivalent using the transfer function of the oscillator given by Equation 2.9. This gives

$$x_{\mathrm{g}}(t) = \Delta x(t) * \frac{x_{\mathrm{g}}(t)}{\Delta x(t)}. \tag{6.5}$$

If only noise performance is of interest, then the equations from Section 3.1 can also convert to the desired unit.

6.1.6   Results From the Prototype Optical Inertial Sensor

The first prototype Optical Inertial Sensor (OIS) was measured on a seismically isolated table. This prototype was made from Sample C2D2S04, even though its Q factor was degraded during coating. The reason for this was parallel development between bonding and building readout and the development of the coating testers and final sample. This table used VIB100 isolation feet to isolate the table from Seismic disturbances above 10Hz. The feet could be shorted by the insertion of slightly longer posts than the isolation feet. This was done in measurements where a direct comparison to a L4C
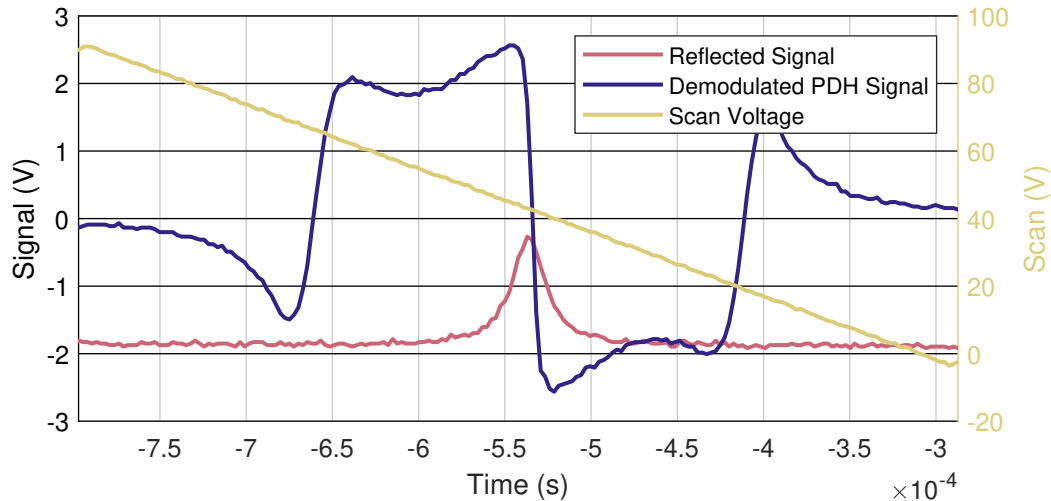
Figure 6.9: PDH Error signal used to calibrate an inertial sensor. The scan voltage is also applied to the laser's piezo. The reflected signal is the reflected DC signal on the photodetector, and the PDH signal is the demodulated signal used in the control loop. These correspond to simultaneous readouts of points A, B, and D, respectively, from Figure 6.1.

geophone was desired. These geophones sat on top of the tank as close as possible to the sensor.

Free Floating Measurements

The first measurements of the floating inertial sensor are shown in Figure 6.11. From the graphs, the limiting factors can already be seen. Between 30 and 100 Hz and below 1 Hz the sensor is limited by frequency noise. The frequency noise shown in the displacement noise plot in Figure 6.11 (a) was measured using an unequal arm length interferometer in the style described by [47]. Between 1 and 30Hz, the isolation feet offer insufficient suppression to cancel out the noise that dominated the sensor readout. A second sensor would be needed to subtract this out. This would, however, also subtract out frequency noise from both if the same laser is used, so a frequency reference would also be needed to evaluate this. This setup was used in Section 6.2. The sensor is then limited by thermal noise above 100 Hz. The prototype sensor had an excellent performance at high frequencies, and it is a clear result that levels of displacement noise of sub-10 fm were achieved above 50 Hz. A frequency reference will be required to perform better at a lower frequency. Furthermore, the ground motion will prevent us from reaching the expected noise floor, so two sensors will be needed simultaneously to do coherent subtraction of ground motion.

Clamped Measurements

The first measurements, shown in Figure 6.12, show the results of the inertial sensor when the table was clamped. The isolation feet were shorted by putting long posts next to them and screwing them to the table. Two L4C geophones were placed on the vacuum chamber so that measurements from them could be compared with each other
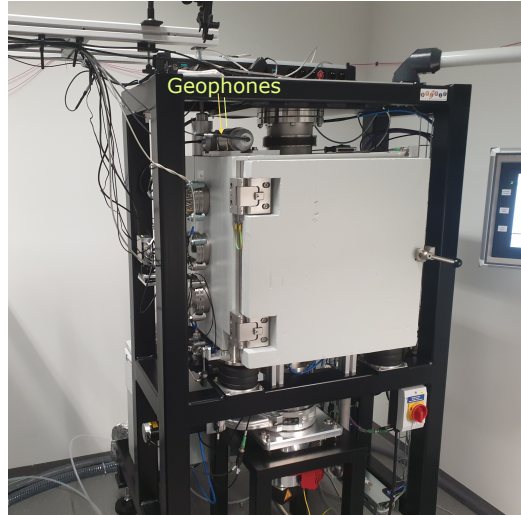
Figure 6.10: A photo showing the positioning of the geophones on the vacuum tank where the experiments occurred. They were positioned as close as possible to the oscillator. This meant they benefited from the tank's isolation from floating on compressed air but did not sit on the VIB100 isolation feet that isolated the optical breadboard above 8 Hz. Furthermore, the geophones sat on the Styrofoam that was used to insulate the vacuum tank, and this likely had some effects on measured motion compared to the tank.

and with the OIS. The geophone-to-geophone coherence shows the range over which they provide reliable data. Above 100 Hz, the geophones are self-noise limited as they no longer agree with each other.

6.2   Frequency Reference

The measurements in Figure 6.11 showed that this frequency reference was required. We must therefore discuss how laser frequency noise couples to our measurements.

6.2.1   Frequency Noise

The intrinsic fluctuations of a lasers frequency directly couple to variations in the length measurement of a cavity. This relation can be derived starting with the relation between phase and length fluctuations of a cavity

$$\tilde{S}_{\mathrm{d}}(f) = \tilde{S}_{\phi}(f)\frac{\lambda}{(2\pi)} \tag{6.6}$$

where d is displacement the apparent length change in the cavity, effectively displacement noise, and ϕ is the phase. Here $\tilde{S}$ is the one sided amplitude spectral density. The phase variations are related to frequency variation by

$$\tilde{S}_{\phi}(f) = (2\pi\tau_{\mathrm{d}})\,\tilde{S}_{\mathrm{c}}(f)\frac{\sin{(\pi f\tau_{\mathrm{d}})}}{(\pi f\tau_{\mathrm{d}})} \tag{6.7}$$
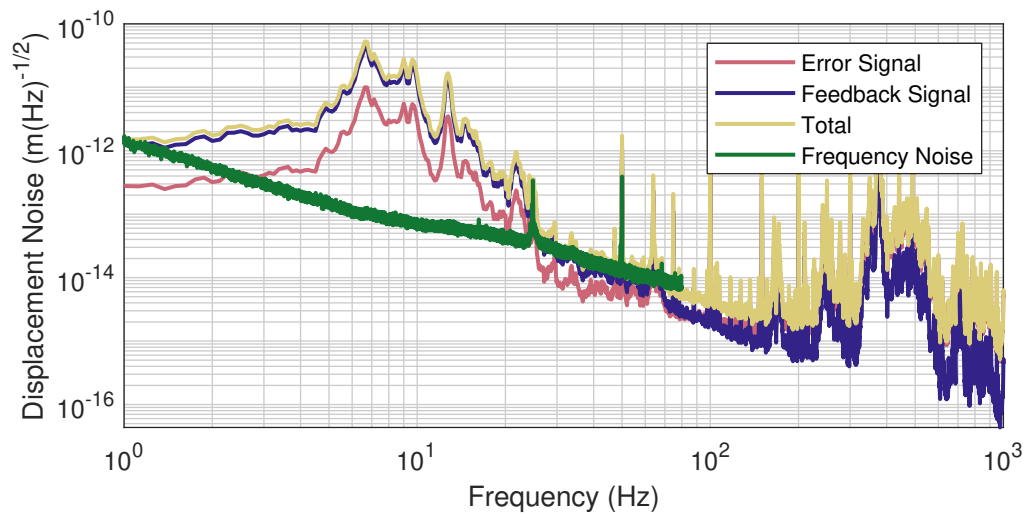
(a) Displacement Noise
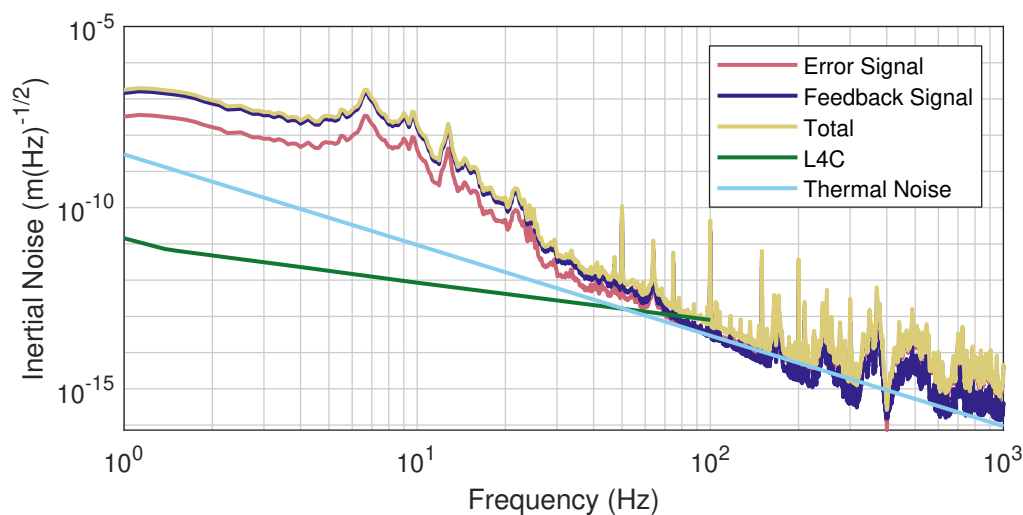
(b) Inertial Noise

Figure 6.11: (Top) The results of the prototype inertial sensor in displacement noise and (bottom) inertial noise. Additional noises are also plotted on the graph. The frequency noise in the displacement noise plot was measured using an unequal arm length interferometer in the style described by [47]. The Thermal Noise contribution was estimated using measurements from Chapter 5. The L4C is a sensor commonly used in LIGO and is representative of the levels the sensors need to achieve to be useful to the Gravitational-wave community. The performance of this sensor was not going to reach those levels due to its high resonance frequency and degraded Q factor.
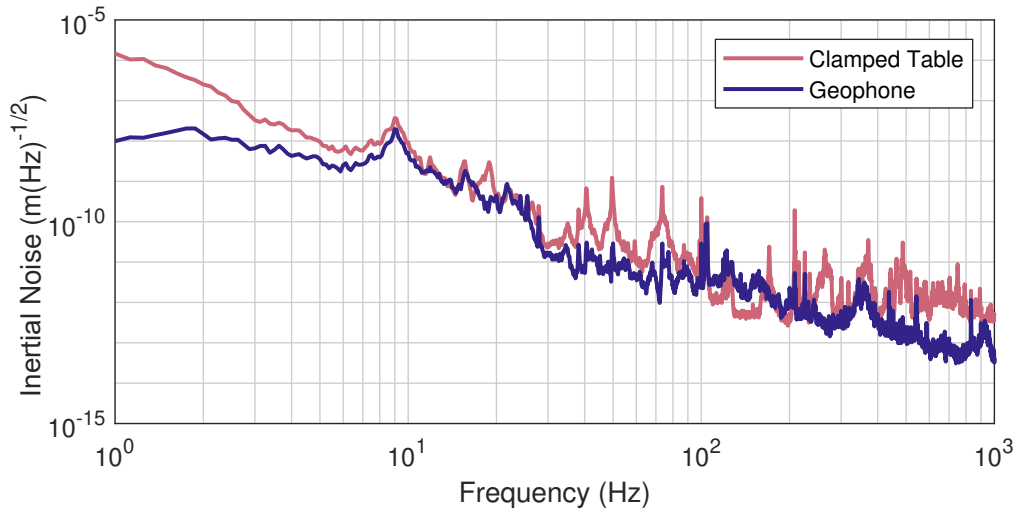
(a) Clamped table Noise
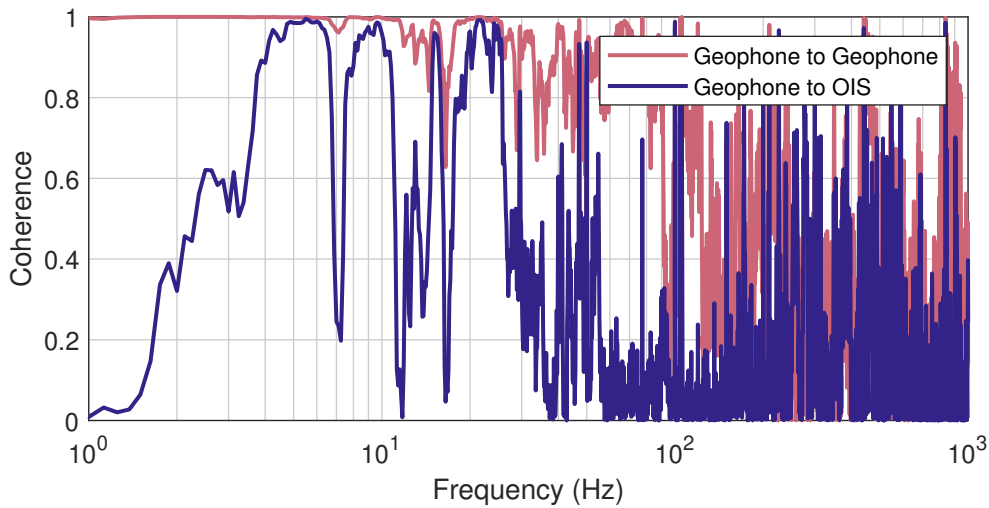
(b) Coherence

Figure 6.12: (Top) The noise of the inertial sensor on a table clamped to the ground and a L4C geophone measuring simultaneously. (Bottom) The coherence between those signals. A strong coherence is measured in the frequency range not dominated by frequency or thermal noise. The loss of coherence above 10 Hz is unexplained. However, it is possible the isolation feet were not perfectly shorted and were somehow driving the table, as that is their resonance frequency. The noise level of the clamped sensor and the geophone agrees well, except some peaks between 10 and 100 Hz, which are much higher in the OIS. This is possibly due to optical mounts vibrating and suggests a need for stiffer mounts.

with c referring to the noise reflected from the cavity, and τ_d being the time delay in the cavity, which for a high finesse Fabry-Perot cavity is

$$\tau_{\mathrm{d}} = \frac{2L}{c} = \frac{1}{\nu_{\mathrm{FSR}}}, \tag{6.8}$$

where L is the cavity length, c the speed of light and ν_{FSR} being the free spectral range of the cavity. As the cavity acts as a low pass filter for frequency noise, normally the value of S_{c} must be scaled by a factor $(1 + f/f_{\mathrm{HWHM}})$ to give the required laser frequency noise, but due to the short length and high finesse of this cavity, this term is negligible for all frequencies that are studied, giving $\mathcal{S}_{\mathrm{c}} = \mathcal{S}_{\mathrm{L}}$, where $\mathcal{S}_{\mathrm{L}}$ is the frequency fluctuations of the laser. From here the final expression of the laser frequency noise can be calculated as

$$\tilde{S}_{\mathrm{d}} = \frac{\lambda}{2\pi} \frac{(2\pi\tau_d)}{(\pi f \tau_d)} sin(\pi f \tau_d)\tilde{S}_{\mathrm{L}}. \tag{6.9}$$

If frequency noise is limiting the performance, a frequency reference can be used to suppress it. A frequency reference is anything which can provide an absolute measure of frequency deviations. These can be created by interfering the beam generated at two different points in time; for example using an optical cavity [89] or an unequal arm length interferometer [122], or by other reference to some frequency standard, such as a atomic transition in a gas [123]. In first two cases, the level of stability that can be gained is given by considering the inverse of equation 6.9. It can be seen that the length variations of the interferometer must be minimised. Hence, state of the art frequency references use low thermal expansion glasses and rigid materials to provide the best possible reference. Typically, the laser is then locked to the error signal from the stable interferometer, meaning that if additional feedback controls are required other sources of modulation must be considered.

6.2.2   Development of a Frequency Reference

The first attempt used an unequal arm length interferometer of the style described in [91], but it never reached the performance. It seemed to be the case that the interferometer was not stable enough as a reference, although it is unclear as to why. The level of frequency noise suppression needed to match a L4C sensor is shown in Figure 6.13. Even to reach performance for a 50 Hz resonator, a suppression by a factor of 200 is needed at 1 Hz over an NPRO lasers free running noise.

The successful attempt presented here achieved the necessary performance to suppress the noise to an acceptable degree. This method used a rigid ULE glass resonator, with a PDH modulation, to suppress frequency noise on the laser by driving the laser's PZT. The scheme used the same readout technique as shown in Figure 6.1.

The resonator used a ULE spacer with the same design as the prototypes, but with four additional venting holes drilled through the centre of the spacer. The in-coupling mirror was identical to the one in the prototype inertial sensor. The output mirror was coated in the same run as the first-generation oscillators. These oscillators were quoted to have a matching reflectivity to make an impedance-matched resonator. As the oscillators had a higher roughness than the blank samples, the reflectivity of these mirrors would be higher. These mirrors, therefore, made the resonator slightly overcoupled. A finesse of 530 was measured in this resonator using the full width at half
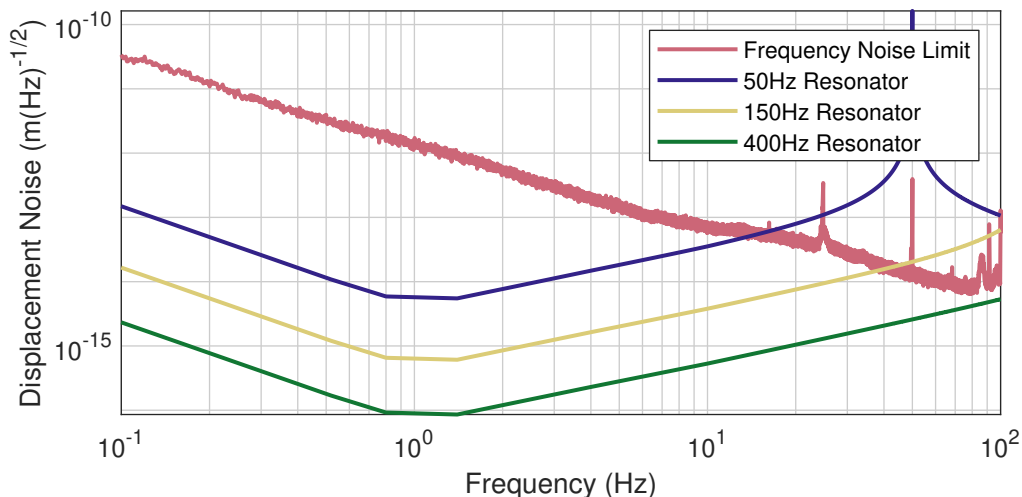
Figure 6.13: A plot showing the frequency noise limit of the inertial sensors and the level of displacement noise necessary to achieve the same performance as an L4C geophone for oscillators of differing resonance frequency. A suppression of frequency noise by a factor of 200 at 1 Hz is needed to match performance in a 50 Hz resonator to an L4C.

maximum as an estimate. Achieving this high finesse suggests that even the in-coupling mirrors had a higher reflectivity than quoted. Taking the upper tolerance limit on a coating of 98.9% reflectivity on the in-coupling mirror suggests the output coupler had a reflectivity close to 99.9%.

The parts were bonded together using the same tool used to bond the coating testers to their mirrors, with an adaption to hold 1-inch optics. UV glue was also used for this process. As the UV glue acts as a thick medium, the mirrors are not perfectly flush. It was found that this led to a approximately 10% Hermite Gauss (0,1) mode contribution, which could not be suppressed via alignment, but this was still smaller than the fundamental mode. This mode is consistent with relative mirror tilts as described by Black et al. [124].

The resonator was setup up as the prototype resonator but with an additional 50/50 beamsplitter in the output. Each port of this output then had a photodetector and demodulation step. One photodetector was used to control and suppress the motion. The other was used as a "semi-out of loop" sensor. As we wanted to suppress frequency noise with this device, any measurement would be limited by the stability of the reference. As the ULE resonator was the most stable reference available, a true out-of-loop measurement of the device's performance was not possible. Using this second photodetector as a measure of performance could show the electronic noise and what was not suppressed by the loop but could not assess the stability of the resonator itself as length fluctuations in the resonator would also be suppressed by the loop.

The out-of-loop noise has been suppressed to an acceptable level. The loop transfer functions were taken and are shown in Figure 6.15. The necessary suppression of 200 was achieved in the required frequency band below 100 Hz. The loop was stable and had some phase margin for extra suppression if necessary. Figure 6.16 shows the resulting displacement noise.
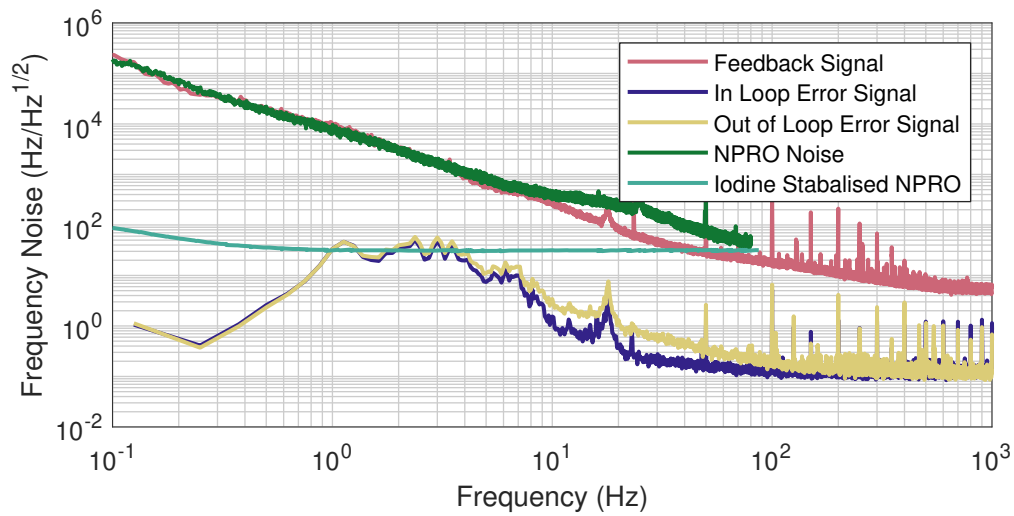
Figure 6.14: The frequency Noise of the free running NPRO, measured both with an unequal arm length interferometer (Green) and this frequency reference (Pink). The in-loop error signal and out-of-loop error Signal are also shown. A residue of some motion is transmitted to the device. Clamping the breadboard did not remove the bump.
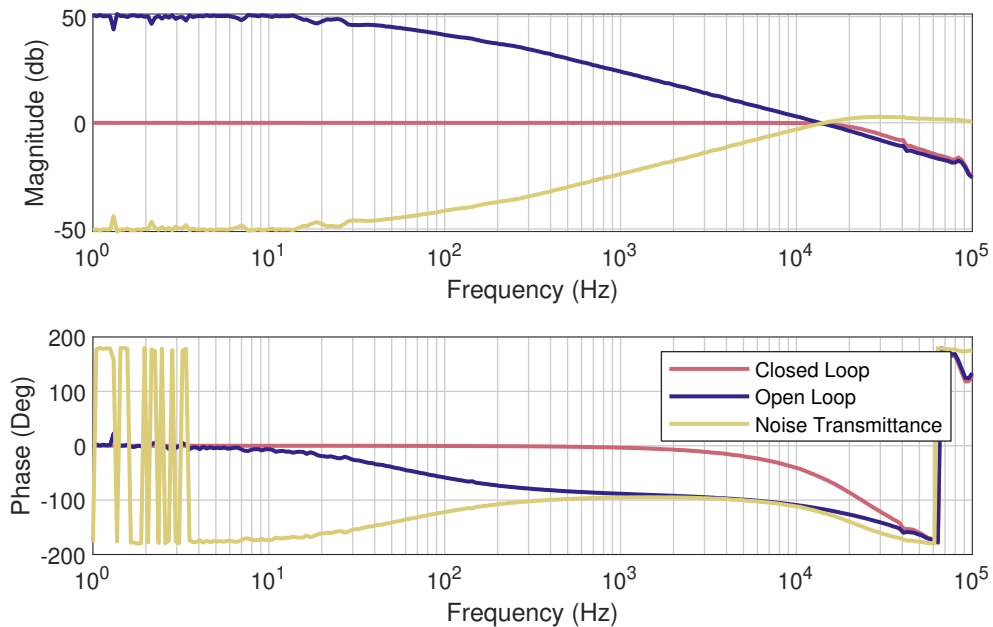

Figure 6.15: The loop transfer function of the controls for the frequency reference. An open loop gain of 50 db at 1 Hz was achieved. A phase margin of 60 degrees is maintained at the unity gain frequency.
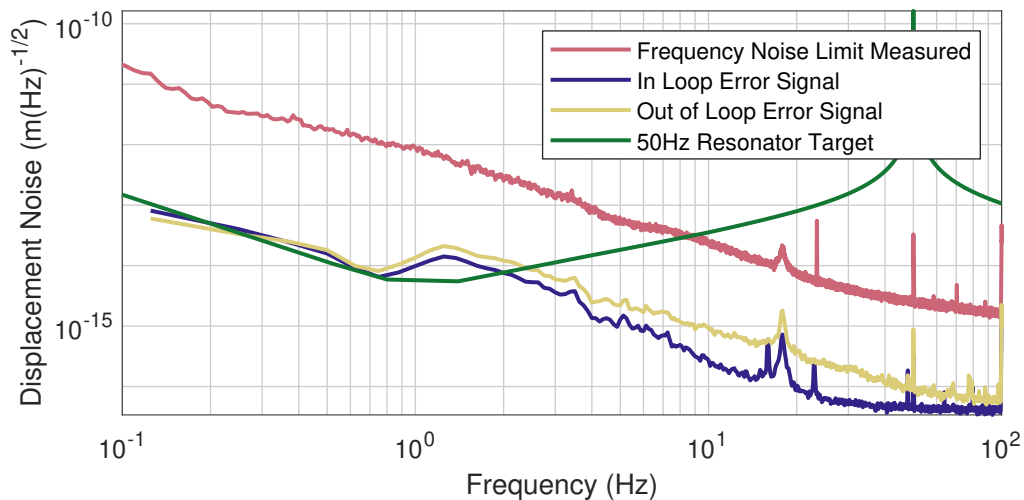
Figure 6.16: The displacement noise of the frequency reference sensor. With the feedback engaged, the out-of-loop sensor measures a noise floor a factor of 2 higher than an L4C in its optimal frequency region.

6.3 Two Inertial Sensors, One Frequency Reference

With the frequency reference reaching the necessary performance, it was time to integrate this with two other inertial sensors to perform a huddle test. Here the performance of two or more sensors is compared, and the coherent motion between them is subtracted. The method is described in some detail by Kirchhoff et al. [125]. To do this kind of measurement, the two sensors must be as close as possible and oriented parallel to each other.

Figure 6.17 shows the scheme for this readout technique. Three beams are injected into the vacuum. One goes to the frequency reference, the others to the inertial sensors. The readout is still done out of vacuum by the same photodetectors as discussed in section 6.1.3. Each sensor has its own photodetector. A mechanical Piezo stack was mounted behind the input mirrors of each of the two OIS. This allowed for tuning and control of the resonator length. Figure 6.18 shows a photo of this setup in vacuum.

The two oscillators used in this experiment were coating testers bonded with mirrors. When this experiment started, the 50 Hz resonators were not ready, and it was unclear if they would be manufacturable. This means the current performance shown here cannot reach the true high-performance sensitivity possible with the 50 Hz samples. Samples S01 and S07 from the coating tests were used, as S04 was damaged when mounted in the setup. The inertial sensor using oscillator 01 will now be referred to as OIS 1 and the one using 07 as OIS 2. With the oscillators installed, the finesse of the cavities was measured. Using the quoted reflectivity of the mirrors (98.5% for incoupling and 99.4% for output), a finesse of 300 was expected. For OIS 1, a finesse of 370 was measured, and 340 for OIS 2, suggesting a higher reflectivity on the coatings than expected (although a configuration of mirrors within manufacturing tolerances could give this).
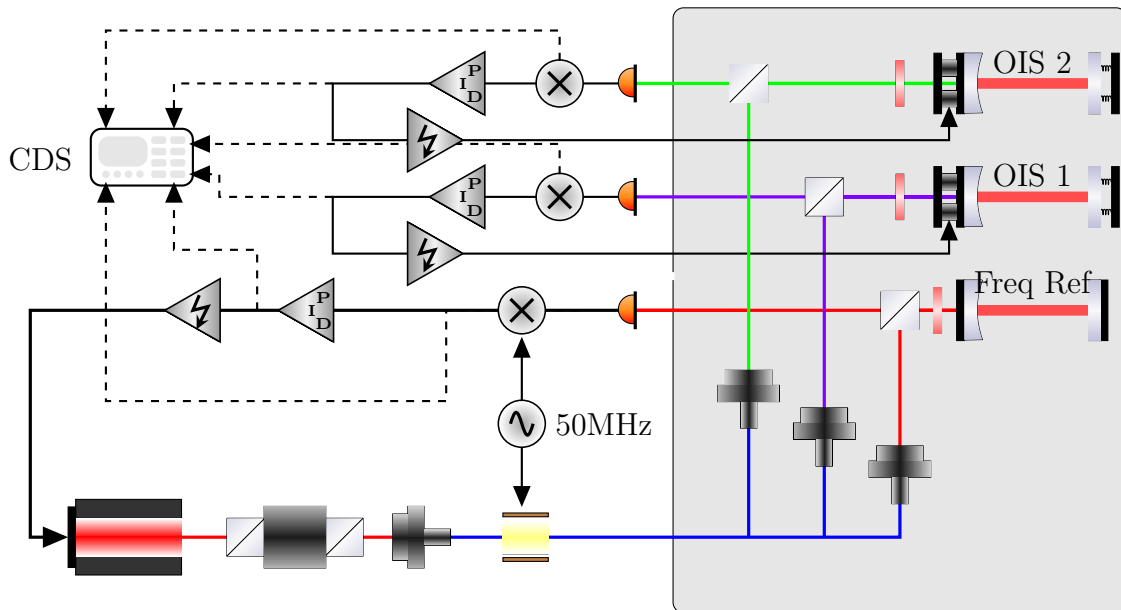
Figure 6.17: The layout of the huddle test measurement setup. 3 sensors sat next to each other in vacuum. One frequency reference and two OIS. Each Sensor had an identical input and readout setup, which is the same as that shown in Figure 6.1. Each sensors had approximately 30 mW of injected power, and where modulated with 50 MHz sidebands, which were demodulated in the photodetector. The frequency reference cavity was used to lock the laser using is fast PZT response. Meanwhile the two OIS were held in lock by using mechanical PZT stacks mounted behind the in-coupling mirror to control the cavity length.
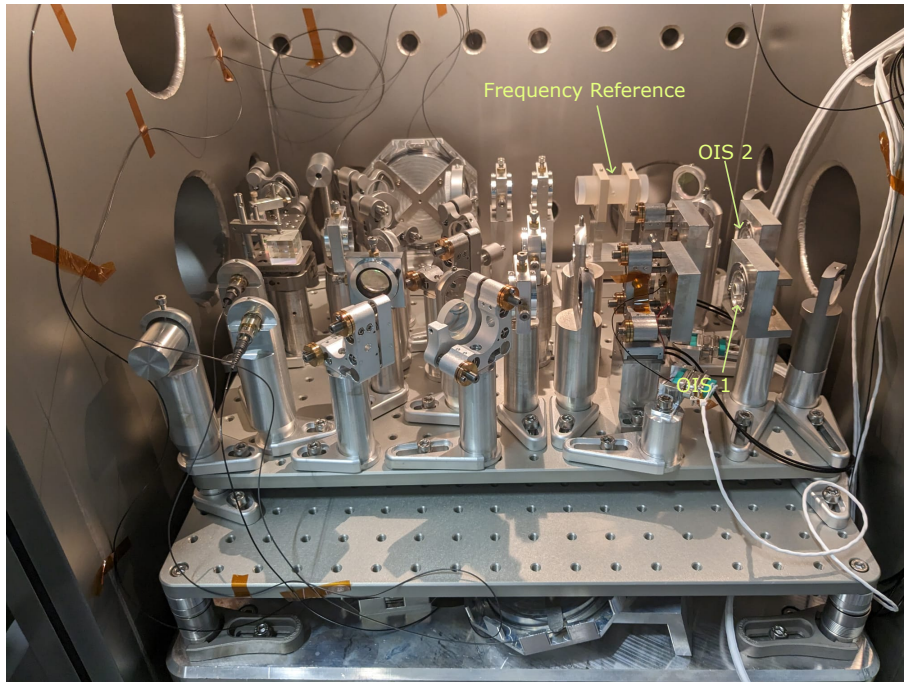
Figure 6.18: A photo of the in-vacuum part of the huddle test. The experiment again sits on the isolation feet for seismic isolation above 8 Hz. The positions of OIS 1,2 and the frequency reference are highlighted.

6.3.1 Laser Modulation of Optical Inertial Sensors

The first step with this setup was to confirm the results from the prototype that frequency noise limits in the low and high-frequency regimes and seismic motion is measured from 1 Hz to about 30 Hz. As each was deliberately scaled to have a similar response to the frequency reference, they were effectively interchangeable with the frequency reference. Using its PZT control, the laser could be brought to resonance with each sensor in turn. This is shown in Figure 6.19. Using the laser to control the OIS prevented the suppression of frequency noise, so frequency noise was limited at high and low frequencies. Between 1 Hz and 40 Hz the sensors are limited by seismic disturbances. As this is a measure of displacement noise and the resonance frequencies of the two sensors are slightly different, the response should be different; however, as OIS 2 had the lower natural frequency, its displacement coupling to the ground should be expected to be higher. This is not seen; infact, OIS 1 measured a higher signal and suggested something was not quite right with OIS1, a recurring problem throughout this section.

6.3.2 AOM controlling the Optical Inertial Sensors

The "cleanest" solution to keep the light resonant with the cavities, and the originally planned method, was to use an AOM to shift the laser frequency to each OIS, then modulate this frequency to keep the light resonant with the resonator. AOMs typically have very good frequency stability, and by doing it this way, the calibration on the feedback control is very trivial. However, they were not used here as a servo control. Initial
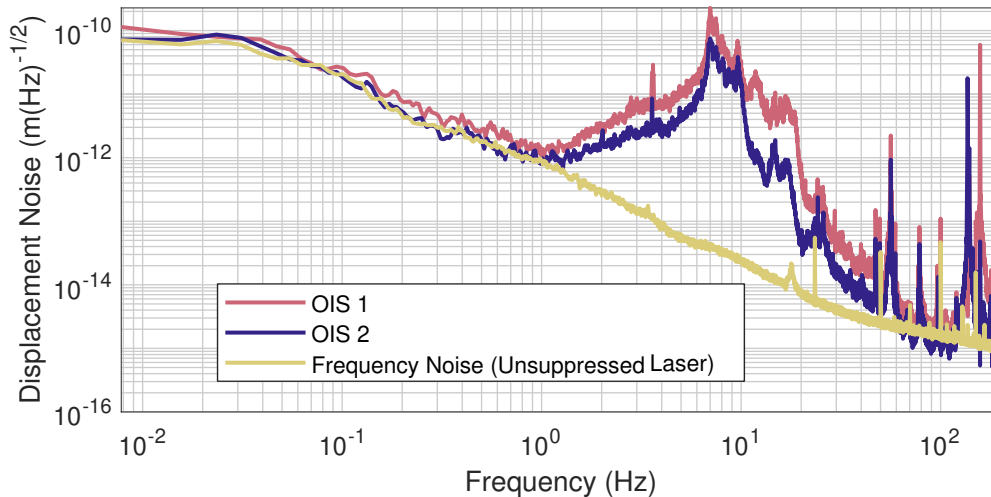
Figure 6.19: The displacement noise of each sensor sequentially along with the noise floor of the frequency noise of an unstabilised NPRO laser. The sensors are frequency noise limited at sub-1 Hz and above 40 Hz. The peak above this is likely due to seismic disturbances.

tests in the lab found they very quickly ended up in unoptimal operating points. The displacement noise of the OIS can be converted to a frequency needed to compensate for that motion, shown in Figure 6.20. This is quite a large level of frequency shift for the AOM to compensate. The AOMs available here were fibre coupled, and so had a very limited bandwidth. The responsivity of the AOMs is shown in Figure 6.21. The large motion, therefore, caused constant shifting of the DC operating point level of the OIS. These problems would be made an order of magnitude worse if a 50 Hz Oscillator was installed instead.

A non-fibre coupled AOM would likely have a greater range of operation, but its output angle would change substantially. With the stringent requirements on alignment in an optical resonator, they would have to be coupled with a retroreflector to be used in such a setup. However, due to availability, a piezo stack was used to change the length of the resonator instead. These would require calibration and are prone to various non-linearities and hysteresis effects. Ultimately, the intended use case of these sensors is in some of the most seismically quiet locations on earth, the isolation systems of gravitational wave detectors. When used in this environment, the AOMs are likely the best solution, but space in such an isolated environment, for example, the AEI 10m prototype [126], must be made available.

6.3.3   Piezo Controlling the OIS

Unable to achieve a good operating point with the AOMs, we sought an alternative means of keeping the laser resonant with the cavity. The mechanical PZT stacks were originally only intended to fine-tune the DC free spectral range so the AOMs did not have to push too far out of their operational range; however, it was decided to try to use them to control the optical resonator length to match the laser frequency. Using the PZT stacks in this way leads to some problems with calibration, namely, hysteresis
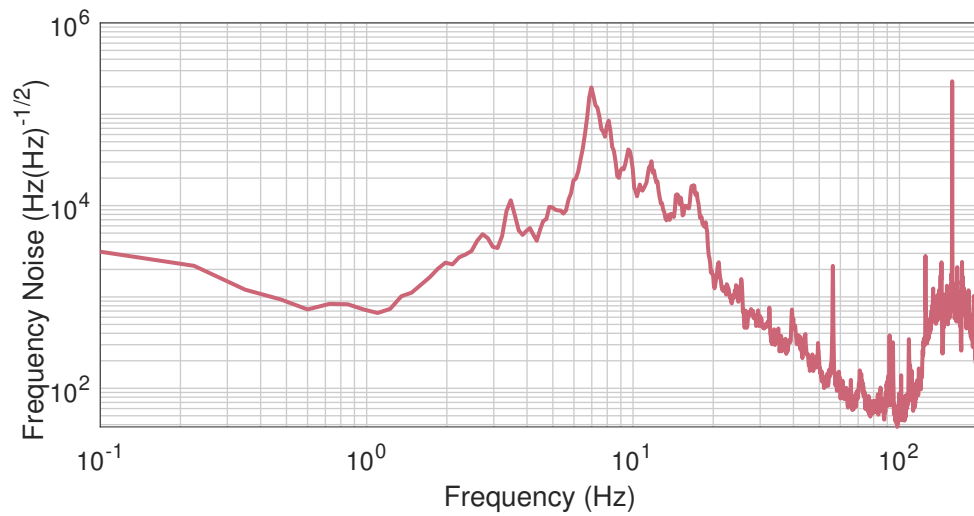
Figure 6.20: The frequency shifts needed to compensate for OIS test mass motion while sitting on the ground. The large motion at sub-10 Hz leads to problems with AOM control of the laser frequency due to AOM efficiency.
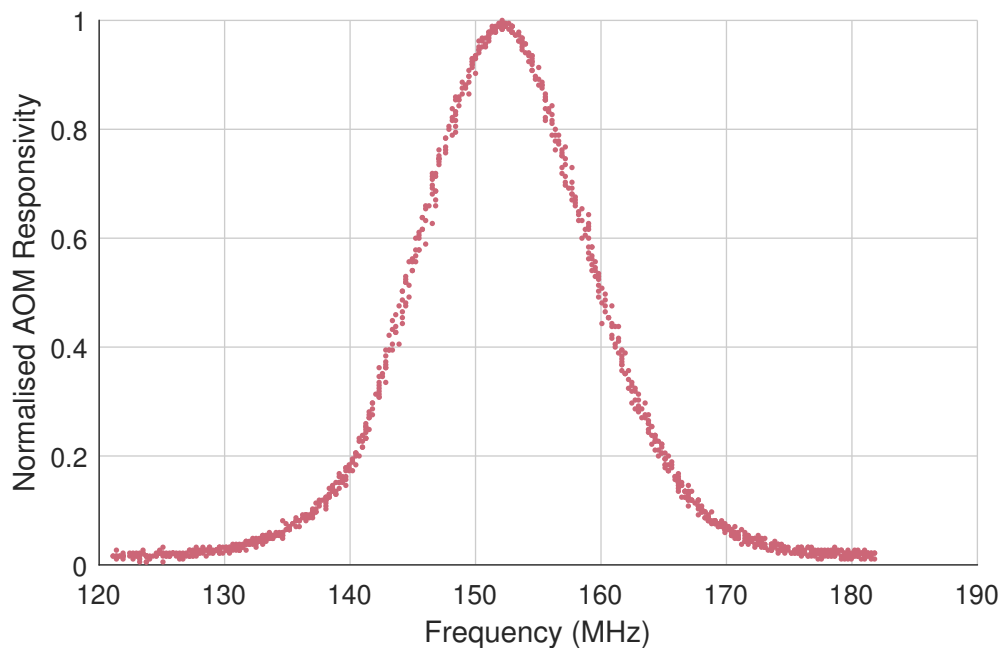

Figure 6.21: The efficiency of a fibre coupled AOM to frequency shifts. A photodetector was connected directly to the fibre, and the frequency across the AOM was scanned. The response is normalised to the maximum responsivity measured.
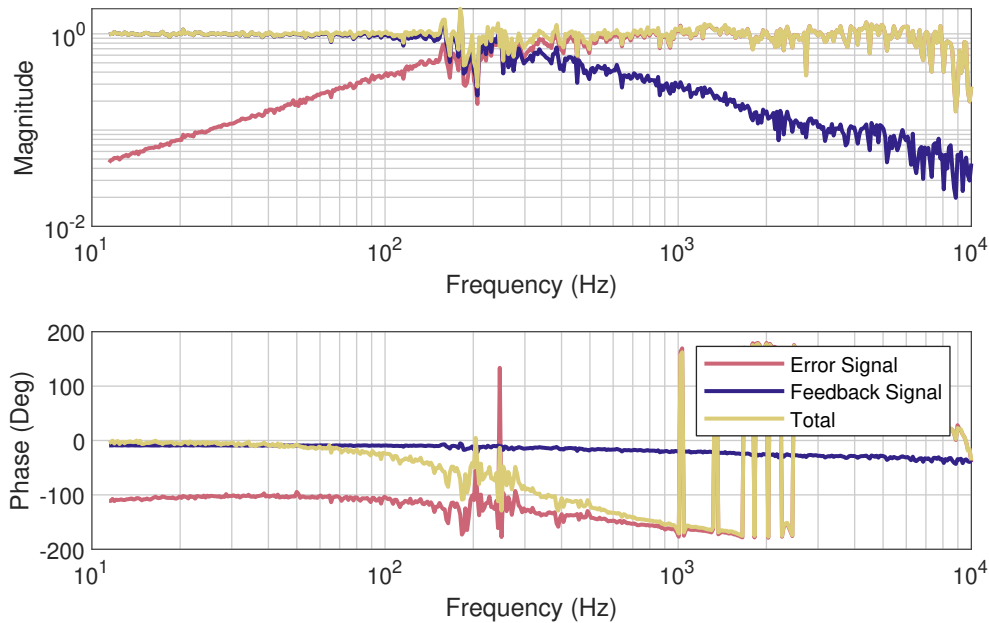
Figure 6.22: Calibrated and normalised frequency response of the mechanical PZT. The gain of the PZT is fairly constant, but there is an increasing phase lag above 100 Hz. The gain is normalised to the gain at 20 Hz, but could also be scaled into nm/V by scaling by 2.4.

effects in the PZT and non-linearity in the response. This was shown by the corollary of the measurement for Figure 6.8, where the laser PZT was used to keep the resonator at its operating point and the mechanical PZT was scanned; this is shown in Figure 6.22. While the total magnitude of the response was fairly flat with respect to frequency, there was a substantial phase lag in the Piezo response from 100 Hz and above. These led to problems with the resonance of the oscillators, and so the control loops had to use low unity gain frequencies to suppress the feedback controls around resonance while still compensating for ground motion. As the resonance had a high Q factor, it was a symmetric and narrow frequency component, which meant that it did not cause longer-term drifts that needed to be compensated.

The DC operating point was measured in two cases when the laser PZT was used to keep the laser resonant with the cavity and when the mechanical PZT was used to control the resonator. The results are compared with when the light is not resonant. The time series is shown in Figure 6.23 and the ASDs in Figure 6.24. The laser PZT control agrees well with the off resonance level suggesting this noise is largely the power noise of the laser. The mechanical PZT control does not sit as well and approaches the limits of the linear response of the PDH signal. When the oscillation is excited (by someone closing the lab door, for example), taking measurements is impossible until the oscillation has died down, which can take some time due to the high Q factors. Therefore, the measurements shown here were taken at night or other quiet times.

The control loops of OIS 1 & 2 are shown in Figures 6.25 and 6.26. They had Unity Gain Frequencys (UGFs) of 20 Hz and 27 Hz respectively, with very limited phase margin for stability. This phase margin was caused by the need for a 2-pole suppression from the servo to achieve sufficient suppression for the oscillator motion close to f_0
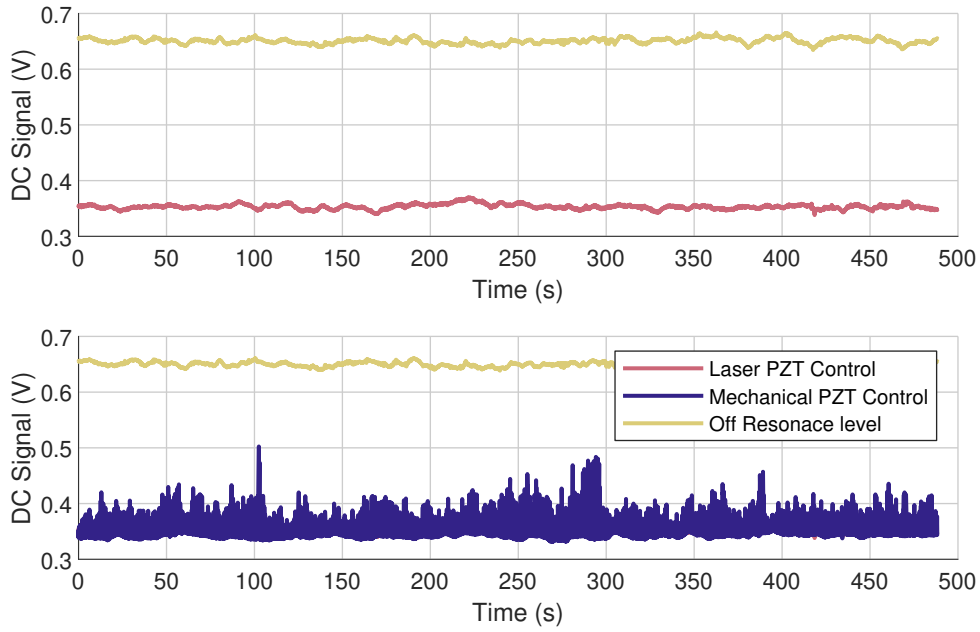
Figure 6.23: A comparison of the DC operating point of OIS 2 when controlled with modulation of laser frequency with PZT (top) and with a mechanical length change using a PZT stack (bottom).
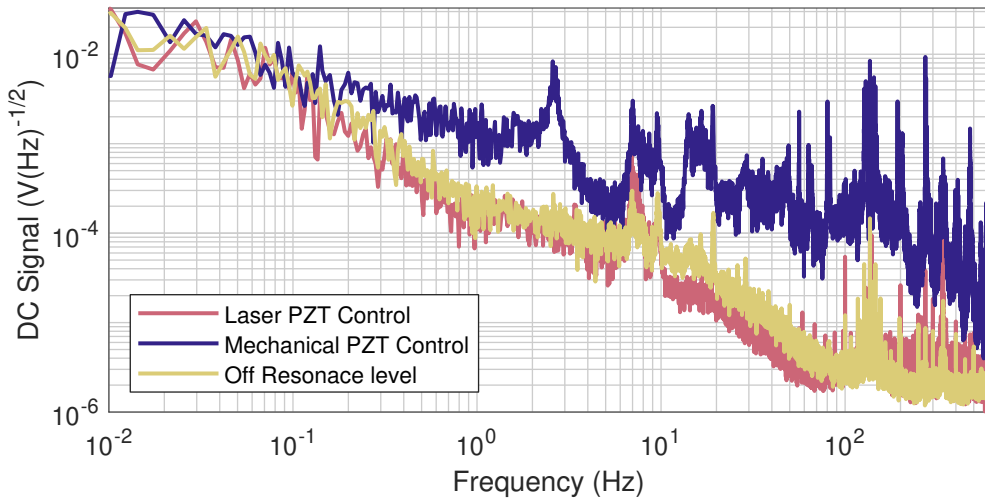

Figure 6.24: The ASD of the DC operating points signals in three cases: when out off resonance, when the light field is the optical resonator condition with modulation on the laser PZT, and when the optical resonator is kept resonant by length change using a mechanical PZT.
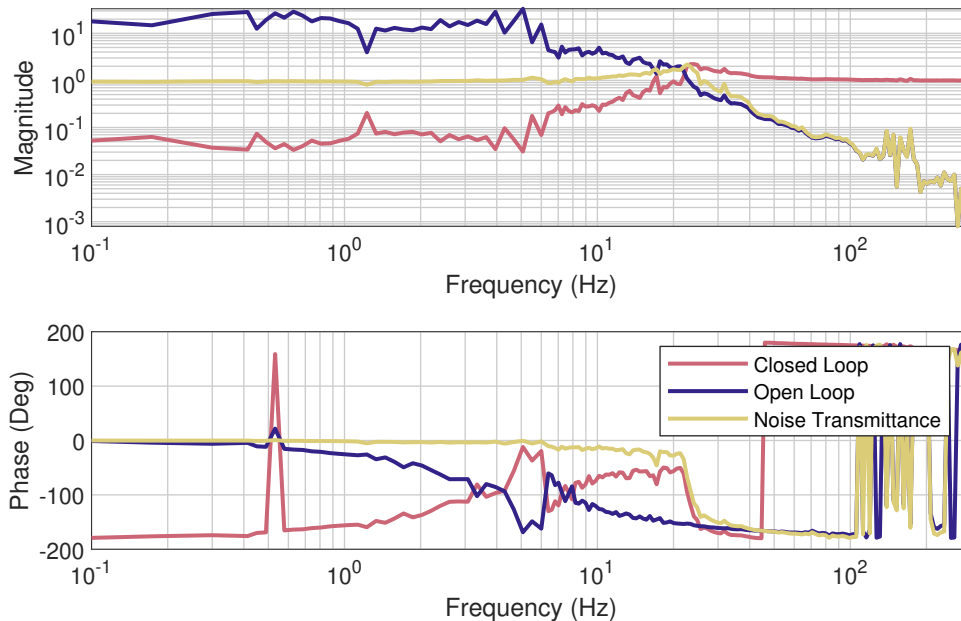
111

Figure 6.25: The transfer functions of the loop controlling OIS1. A UGF of 20 Hz is achieved, but with a very limited phase margin for higher gains.

to stop the loop from becoming unstable. It may be the case that the 50 Hz or lower frequency oscillators may be easier to control with the mechanical PZT as there there is no phase lag on the stack itself. Other control methods could be using the PZT stack as the slow response of a two-servo control, with another means of fast control for suppressing the oscillator motion on f_0.

With this, a comparison can be made of the sensor's noise performance when controlled with the laser PZT and mechanical PZT. The noise performance in both modes is shown in Figure 6.27 along with the unsuppressed laser frequency noise of an NPRO. By using a mechanical PZT to keep the resonance condition, the frequency reference can suppress the laser frequency noise. This is exactly what we see in the graph at low frequencies, where the mechanical PZT control surpasses the frequency noise floor, showing that the frequency noise suppression is working as intended. The trade-off here is extra noise above 20 Hz likely from the insufficiently suppressed system causing disturbances. This means that the servo is suitable for low-frequency control, but not high-frequency noise performance. As 10 Hz and below is the typical bandwidth for gravitational wave detector controls, we aim to achieve the best performance at lower frequency.

6.3.4   Noise Performance of the OIS

In order to assess the noise performance of the OISs, two were set up as close as possible, as seen in Figure 6.18. The geophones were used in Figure 6.10, but the sensors were moved and rotated to sit above and aligned with the new OIS positions. The first test was to check the OIS's calibration against the geophone, resulting in Figure 6.28. It is expected that below the resonance of the isolation feet at 8 Hz the geophones should follow the OIS well, assuming they are measuring ground motion. The sensors do follow
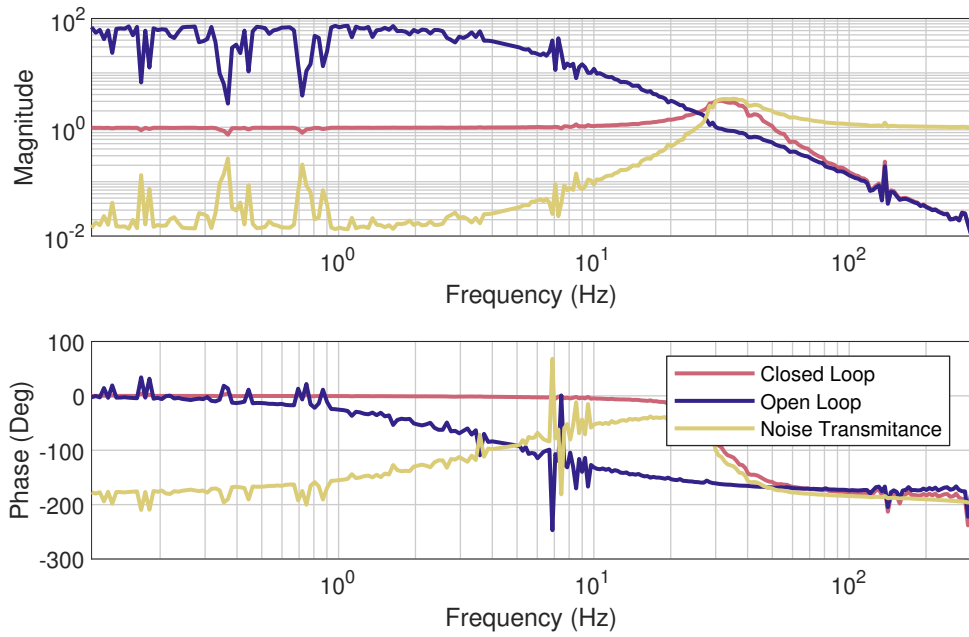
Figure 6.26: The transfer functions of the loop controlling OIS1. A UGF of 27 Hz is achieved, but again with a very limited phase margin for higher gains.
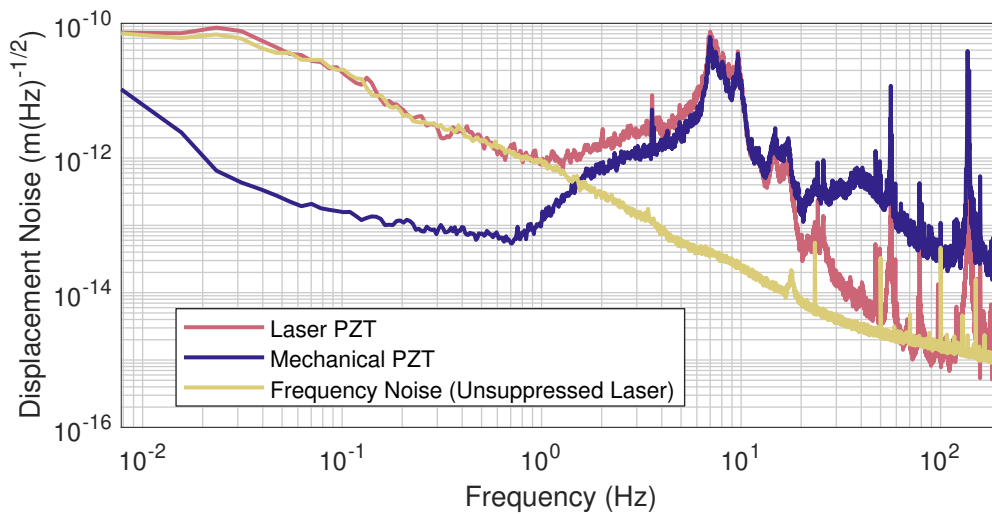

Figure 6.27: The displacement noise of OIS 2 when controlled with the laser modulation using its PZT and with a mechanical PZT stack controlling the resonator length. Using a mechanical PZT on the resonator allows for suppression of laser frequency noise with the stable frequency reference discussed in section 6.2. This is shown below 2 Hz where the measured noise is below the laser frequency noise of an unsuppressed NPRO laser. Above 20 Hz the noise of the mechanical PZT control sits largely above that of the laser controlled line.
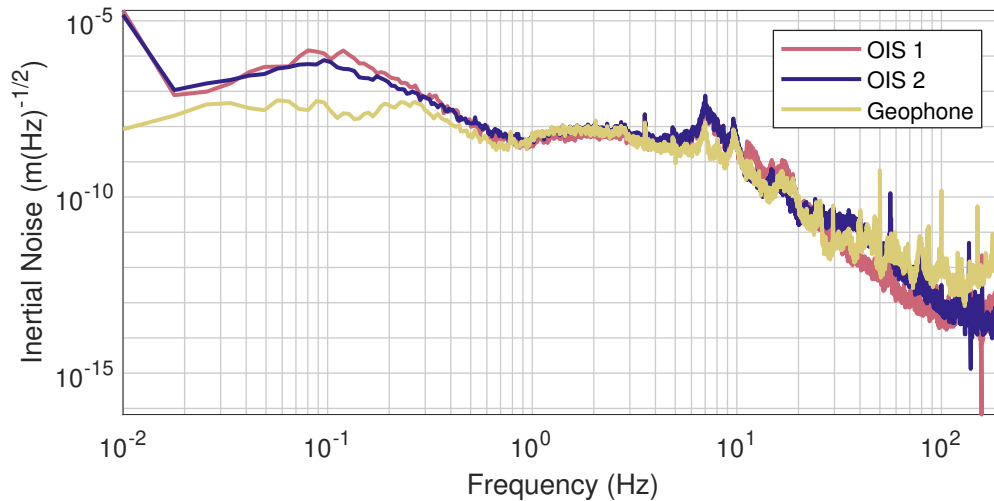
Figure 6.28: A comparison of measured noise from two OIS and the geophone mounted outside the vacuum. The signals are in good agreement in frequencies where this should be expected, below the isolation feet resonance at 8 Hz (acceleration equivalent D.4)

well between 1 Hz and 8 Hz, but diverge below this. The increase is close to a $1/f^2$ slope, strongly indicating extra tilt coupling to the measurement, as seen in Equation 2.11. This is likely because the platform itself was not perfectly balanced, so unequal loads were applied to each suspended foot. This likely needs correcting in future experiment versions, with some space for mounting extra balancing weights.

With the sensors seemingly well calibrated, it was decided to look at the level of noise suppression between the two channels by coherent subtraction. This is shown in Figure 6.29. Only a limited subtraction is seen in this case, largely due to small calibration errors in the setup. This is a very common problem in the field of inertial sensors, where calibration errors are much larger than the sensor's noise performance. This is not a problem when these inertial sensors are used in active control as servo loops, actuators responses, and blending filters are all tuned to achieve the desired suppression, and the absolute value is not so relevant to this.

However, this does present a problem in trying to measure the noise of these sensors, but a script for this was developed and has been used within the community for some time [125]. The script finds any common signal between two sensors and optimises the coherent suppression of the signal. Doing so reveals the true noise floor of a sensor. This script was run on a set of time series data, and the result is shown in Figure 6.30. Here OIS 2 was used as the original signal, and other channels were subtracted from it. Using OIS 1 as the subtraction channel only produced a slight noise reduction, but a much better coherent subtraction was achieved when a ground geophone was used. This indicated that OIS 1 was not performing as well as expected or was measuring something else.

The coherence of all the sensors to each other was measured and is shown in Figure 6.31. OIS 1 and the ground sensor have a much lower coherence to each other than OIS 2 and the ground sensor. A reason for this was sought, and after trying several things, it was found that OIS 1 had a much stronger than expected coherence to the vertical geophone, as seen in Figure 6.32. As uneven vertical ground motion would
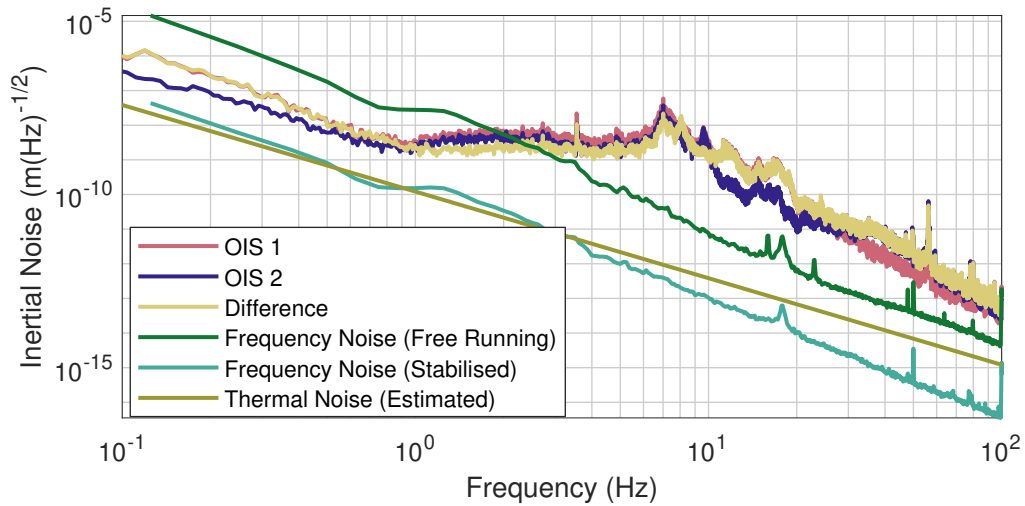
114

Figure 6.29: Measured inertial noise of the OIS and the "naive" differential signal, with the fundamental limiting noise sources of the sensor, the stabilised laser noise and the thermal noise of the oscillator. The subtraction is very limited due to small calibration differences limiting the subtraction of common noise in this setup; hence the MCCS2 routine is necessary. (acceleration equivalent D.5)
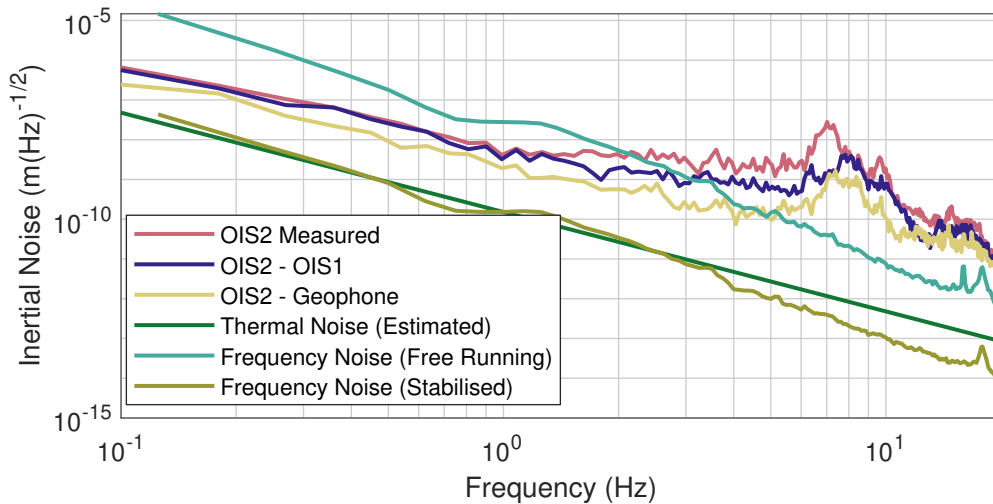

Figure 6.30: The noises of the OIS when coherent subtraction is performed with different sensors using the method described by [125]. This is used to subtract motion and measure the actual noise floor, despite calibration differences between the sensors. OIS 2 is used as the sensor and OIS as the subtraction to produce the blue line, which showed only limited subtraction. When a ground-mounted geophone was used for subtraction, a much greater reduction was measured, but it still sits a factor of 5-10 above the fundamental noise sources. The remaining motion is likely platform motion independent of the geophone motion. (acceleration equivalent D.6)
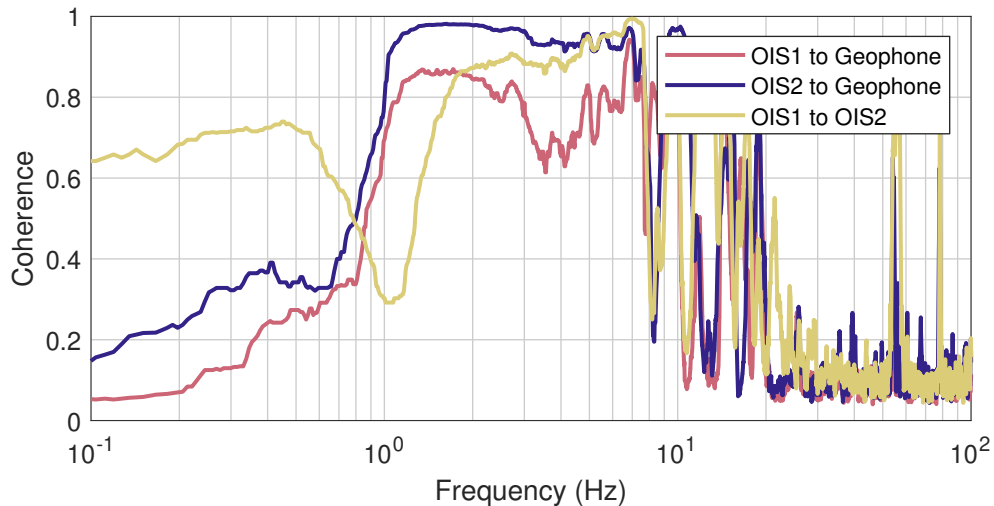
Figure 6.31: The coherence of the two OIS to a geophone with their sensitive axis aligned. OIS 2 shows a strong coherence to the ground sensor, while OIS 1 only shows a limited coherence to the ground sensor. Both OIS show limited coherence to the ground sensor below 1 Hz. This likely corresponds to tilt differences between the isolated platform and the geophone. Above 8 Hz, the OIS is isolated from the ground and would show only limited coherence to the geophone in that region.

drive platform tilt, it is effectively a proxy for tilt motion on the measurement. Hence, OIS 1 seems much more susceptible to ground tilt. This indicated that maybe the mirror was glued to the oscillator off-centre. As the alignment of the mirror was done by eye, it likely is that the mirror was offset and introduced additional tilt coupling. Furthermore, as the mirror overhung the centre flexure region on one side, it would be slightly weighted to that side. As the samples were meant to test the Q effects of bonding more than their actual use in inertial sensors, these were not designed to be perfectly balanced. The 50 Hz oscillator had a space evacuated for the mirror to sit in.

The subtraction from OIS 2 and the ground geophone still does not reach the expected noise performance from the combination of displacement readout floor and thermal noise projected in Figure 6.30 also. This is likely due to the difference in seismic motion between the out-of-tank and in-tank states. This residual seismic motion would prevent us from measuring the readout floor. The only solution would be to bond and install a replacement OIS for OIS 1 in the setup.

6.4 Chapter Summary

This chapter has presented the development of an optical readout scheme with a noise floor capable of doing justice to the low thermal noise of these oscillators. A prototype sensor was first tested, highlighting two key problems that must be overcome. The first was frequency noise, and the second was residual ground motion disturbing the measurements. The first required installing an ultra-stable resonator reference; the second was overcome by testing multiple sensors and subtracting common signals from the measurement.
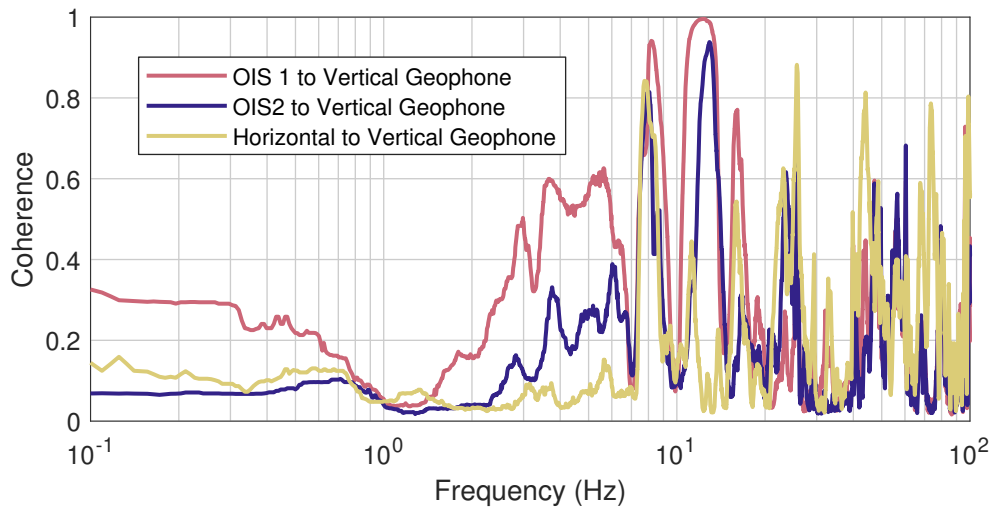
Figure 6.32: The coherence of the various sensors to a vertically mounted geophone. The vertical ground motion would effectively lead to tilt motion on the platform, indicating tilt susceptibility. OIS 1 shows a much greater coherence to this than OIS 2, which suggests a greater tilt susceptibility. This greater tilt contribution leads to the limited coherence measured by this sensor.
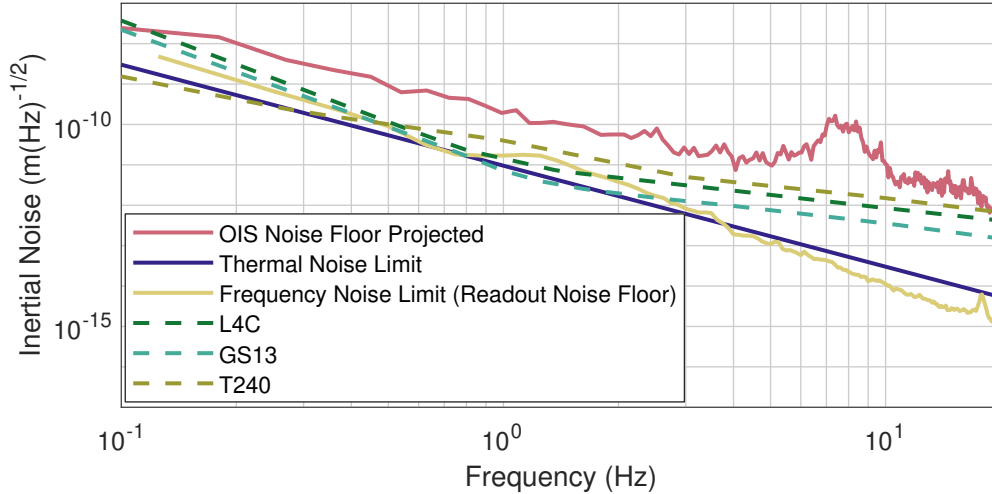

Figure 6.33: A comparison of the projected noise from this readout scheme onto the 50 Hz oscillator, along with its thermal noise and frequency noise against several of the state-of-the-art sensors used in LIGO. The sensor's theoretical noise floor, defined by the readout noise and thermal noise, surpasses that of many of the state-of-the-art sensors. The measured readout noise does not reach that of the theoretical limits, but still is approaching close to the performance of the state-of-the-art sensors. (acceleration equivalent 6.33)

A stable ULE spacer was used to make an optical resonator based frequency reference. Out-of-loop sensors assessed this resonator's performance, and it is seen that the sensor reached the required noise performance such that frequency noise is no longer the dominant noise source in the sensors. The ultra-stable resonator also proves as an estimate for the noise performance of the readout. It shows that the readout noise reaches about the level needed to measure the thermal noise of the oscillator.

In order to utilise the frequency reference, alternative means of actuating on the resonator had to be tested. AOMs would be a simple and easy candidate for extremely quiet environments such as free-floating satellites [127, 128] or in the suspension chains of ground-based gravitational wave detectors [15, 16, 17]. Sadly, the lab environment was too noisy for their direct use, so alternative actuators were used. These were the mechanical PZTs. Due to their slow response, they had trouble tracking the resonance of the oscillators at 150 Hz. Measurements indicate that they may be able to track a 50 Hz resonance better, but this has still to be determined.

A huddle test was performed on two OIS. One sensor performed well, while the other did not reach expected performing likely due to a mismatched load. This is largely because of misalignment with the mirror bonding. Future oscillators will have a cutout to reduce that risk. A ground-mounted geophone was used to subtract ground motion from the signal successfully. The results were still a factor of 5 above the noise floor of the readout, and this difference is largely attributed to a residual difference in seismic motion in and out of the tank.

The measured displacement noise performance from OIS 1 and 2 can be projected onto the case when a 50 Hz oscillator is used. The result of this is shown in Figure 6.33. Here it is compared with the state-of-the-art sensors; while the OIS still sits a factor of 5 above the performance of the L4C, we believe that the difference is largely due to residual seismic noise disturbing our measurements. We could already show that the sensors in this configuration would achieve a sub-ng performance in the range from 0.1 Hz to 10 Hz, and, likely, they are already capable of achieving sub-100 pg in this range, but this must be verified. The clear next step of the project is to put the 50 Hz oscillators into the setup and measure their noise preference directly. Due to time constraints and the late arrival of samples, this was not yet done but is planned for the near future.

While the work on the readout scheme is not finished, excellent noise performance is nonetheless achieved, and several technical limitations have been overcome.

Chapter 7

Processes for Inertial Sensor Design

*There is no optimal inertial sensor design. Every application has unique design require-
ments and would ideally have a sensor uniquely designed for its role. Therefore, optimal
design only makes sense with respect to system requirements. This chapter focuses on
a process for designing inertial sensors to meet a use case. We then look at how this
was used to make inertial sensors within the scope of this project.*

7.1   A General Design Process

The key aim of this thesis has been demonstrating that inertial sensors can achieve high performance in a compact housing, suitable for use in a gravitational wave detector. The next step for the field is producing designs for specific applications within the detectors. To this end, we evaluate how the sensors were designed in this project and use this to develop a process for others to use when designing such sensors.

Figure 7.1 shows a flow diagram of this process. The diagram breaks down the process into four phases:

- a concept study,

- a flexure design phase,

- a geometric oscillator design phase,

- and a sensor performance evaluation phase;

which we discuss in the next four sections.

7.1.1   Concept Study

The initial concept study aims to define the style of inertial sensor needed. The sensor's overall goals, such as performance, operation environment, size and practical limitations must be considered. These initial criteria are also relevant to other decisions that should be made early, in particular, the readout method. In the diagram, this is highlighted as both a design input and an area of study. Often, experience or knowledge of a particular readout scheme within a group will make the choice of readout already fixed by external factors. If this is not the case, then literature research and study on different readout schemes should be undertaken to choose a suitable method, as done in Section 3.4 for the state-of-the-art compact optical readout techniques. We also use our design requirements and experience to determine what manufacturing techniques are feasible and available and what materials that allows us to use. We can then choose appropriate materials to work with from what can be manufactured and what has the desired material properties, such as shear stress and bulk loss. We can also evaluate a cost metric for our design. Combining all these, an initial concept can be envisioned. The key point here is that the fundamental frequency, f_0, must be chosen. f_0 will now define the rest of the process, so ideally, we set this as early as possible.

Before moving on to more intricate design steps, we should evaluate the chosen parameters. Does the expected readout noise floor meet the target noise? Assuming values for the bulk loss as a Q, how heavy must the test mass be to reach a sufficient thermal noise floor? Is this amount of suspended mass feasible with the design methods? Does the readout and housing fit in the space limitations? Until these questions are answered, there is no sense in doing further detailed design work.

7.1.2   Flexure Design

Once we have the design parameters we wish to explore, we must determine what flexures are needed to achieve these parameters. An in-depth study of flexure response
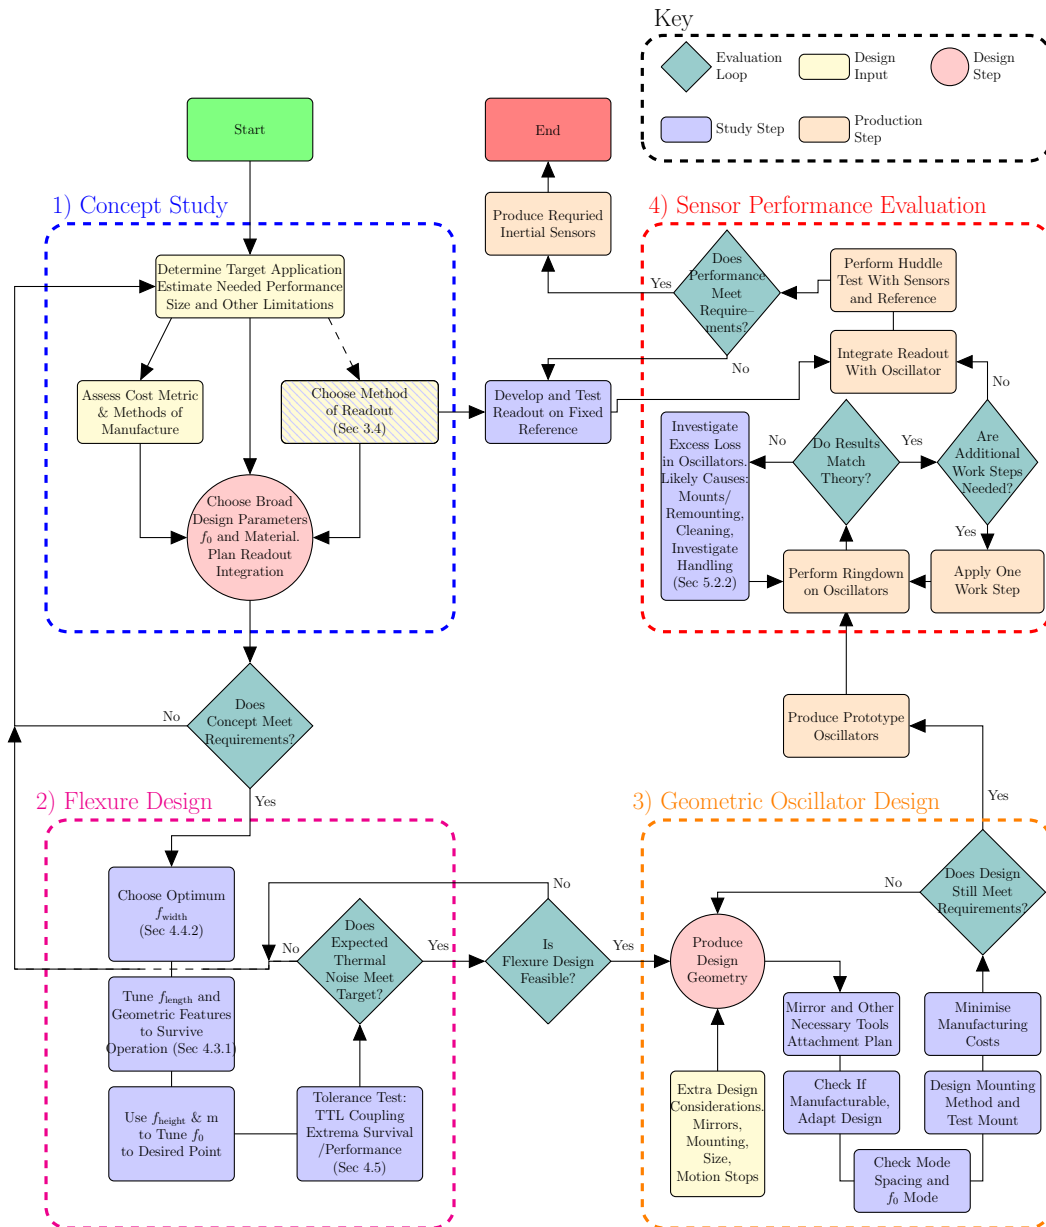
Figure 7.1: A flow chart showing the design process for inertial sensors developed in this thesis. The dashed boxes highlight each phase. Each of these phases is broken down further into individual steps. The colour and shape of the boxes indicate the nature of the step. Yellow boxes are design inputs from the user, typically from the use cases of the sensors. Blue boxes indicate a step where a detailed study is needed to get a definitive design answer, which will depend again on the use case. Sections where the relevant studies have been conducted in this thesis are referenced. Orange boxes indicate a production step where a specific action must be taken. Teal diamonds create a self-evaluation loop in conjunction with the red circles. While often presented as a Yes/No answer to aid the flow of the chart, it would be better when a No is a result to evaluate why this was the case and use that to help with the iterative design steps for the next loop. Red circles are broader design steps. The previous iterations and their inputs usually inform them, but typically are meant to be open to allow for more out-of-the-box thinking.

to design parameters was conducted throughout Chapter 4. We can now use this to define our oscillating system. The first parameter to define is the flexure width, as other flexure parameters do not define its optimum width. The flexure width and f_0 can then be used to determine the length of flexure needed to survive its operational environment and handling, which is discussed in Section 4.3.1. Additional features such as motion stops, radius of curvatures and varying thicknesses can be added to distribute stress over the flexure more evenly if the manufacturing methods allow it. From here, flexure height and test mass must be tuned. Typically this means tuning to achieve sufficient thermal noise performance and the correct resonance frequency. Overshooting the thermal noise target will often come with extra manufacturing costs, so this should only be pushed as far as necessary, whether that means meeting the design target or pushing thermal noise below readout noise. When doing so, it is advisable to leave some safety margin on Q factor estimates because, as shown throughout Chapter 5, Q factor do not always reach expected values.

After the idealised flexure geometry has been established, we must consider the imperfections of the real-world application. Section 4.5 explores how best to define tolerances. Typically, this means looking at effects on the mode frequencies (both f_0 and higher order modes), checking that the design can still survive in its operational environment, and studying angular effects and TTL coupling from mismatches. When testing tolerances, we could find the extremes where the design still works and see if that tolerance is manufacturable. Alternatively, we can take the tolerances from the manufacturing methods and see if the resultant designs are acceptable across their whole range. When working with third-party companies, the former of these is often better, as companies are often unwilling or unable to state the exact tolerances their methods can achieve. Furthermore, it guarantees a minimisation of manufacturing cost as, by definition, we have the broadest acceptable tolerance. The latter is a better approach when we have some control over the manufacturing, as then it becomes immediately apparent if the manufacturing method can work to such precision or if refinement or another approach is needed.

We should then check that the flexure meets the thermal noise requirement. It may be the case that either the requirement of the concept design cannot be met. If this is the case, we could explore pushing f_0 to lower frequencies, but this may lead to further changes in the readout method and send the process back to the start. It may also be the case that the feasibility of the design target must be checked. If the thermal noise from the flexures is sufficient, we should still evaluate if the flexures are manufacturable with the available design techniques. If the flexures have been pushed to an extreme position on any parameter, we should evaluate how this was caused and if we can remove the restriction causing it.

7.1.3   Geometric Oscillator Design

With the flexure and mechanical parameters chosen, the next question is how to design the oscillator geometry. The geometry again invites infinite possibilities, but the decision should focus on the five M's: modes, mirrors, manufacturability, mounting, and minimisation.

We start by considering the modes. We want f_0 to have a linear motion parallel

to the sensitive axis. Other modes should be checked and pushed out of the sensitive bandwidth and as far away as possible from f_0. Breaking the flexure height between different sets and mounting from different sides can have this effect, leading to the drum geometry studied throughout this project. Other approaches, such as using two oscillators with different f_0, can also broaden the sensitive range [129].

How the necessary extra tools will attach to the device is also key. For optical readout, a highly reflective coating is often needed, the design may need motion stops, and it will also need to be integrated with the readout. An important step is working out the order they will attach to the oscillator. A compatible order of work should be established. This neatly moves into the mounting. It is advisable to design already mounting and handling tools at this stage. Using physical models of oscillators, it is possible to get a feel for where mistakes can be made and how to add ease-of-use features to the design. We should minimise design costs. The design features that add to the cost and do not change performance should be removed. In the context of the FEMTOprint manufacturing method, this typically means minimising etched material. These steps should be iterated upon to produce the necessary geometry. When we have a design that can be made and meets the design requirements, we can move on to manufacturing.

Manufacturing Prototypes

Figure 7.1 does not do justice to the manufacturing process. Typically, for such complicated geometries, this step will be a very involved process (often with its own design and workflow charts). As I do not make the oscillators myself, I leave this to be better documented by those who do [63, 102] and an upcoming thesis from Pascal Birckigt, who designed the process for the 15 Hz oscillator manufacture. Producing a small prototype batch before committing to larger orders is typically advisable. For testing purposes, this small batch may be the entire sample run. At least three samples are generally a good target for the prototype batches, as breaking samples is not uncommon, and sometimes samples do not ring like the rest of the batch (sample 09 of the coating testers, for example). Ideally, two working prototypes must be produced if a huddle test is needed.

7.1.4   Sensor Performance Evaluation

When prototypes are produced, their performance should be evaluated. We use a two-step process. Using a ringdown method, the mechanical behaviour of the oscillators can be verified against the theory and models. The huddle test measurements can assess the performance of the overall sensors when integrated with their readouts.

The ringdown measurement process is presented in Section 5.2, where the Q factor and f_0 can be measured. Although conceptually simple, the implementation has many nuances. It can be difficult to distinguish between measurement limitations and oscillator Q factor when the theory does not match the measurements. Typically, if a common Q factor is seen through all measurements, a common factor limits them, usually hinting at the experiment rather than the oscillator. In these cases, the principal suspect is always the mount. Changing mounts should show this, even if the new mount does not make the Q factor higher, as a change indicates the experiment is mount-sensitive.

Even a good mounting plan can lose an order of magnitude in measured Q factor from something as simple as an oscillator not being centred. When remounting does not affect the Q factor, evaluating the handling for contamination opportunities is also a good idea. If a very low Q factor is consistently measured from a contamination event, a HF acid bath has been shown to restore some ring by etching away the top layer, but this never reached the expected Q factor. If the oscillators are not reaching the theoretical Q factor, we can ask if the performance is still sufficient to meet the goal; if not, we go back to the flexure design phase and add a larger safety margin on the Q factor estimates.

If the oscillators ring sufficiently, we then begin the add to extra work steps, such as adding coatings or readouts. These steps should be done one at a time, and ideally on one oscillator at a time. Checking after each stage allows the identification of Q factor loss actions, which can be avoided in future designs.

The specific readout implementation should be tested in parallel to the oscillator testing. A fixed mirror in place of the oscillator can be used to estimate the readout noise floor and show that it works as intended.

The final stage of evaluation is then a huddle test between two sensors as described in 6.3. Often the additional feature of the test mass motion means that the readout will require further work, but ultimately when the measurements from the huddle test show the correct expected results, a repeatably producible inertial sensor suitable for the application should be ready for deployment.

7.2   Applied Design Process

With the design process defined, we now look at its implementation and how it led to two of the inertial sensors made throughout this project. First, we will show how we reached the design geometry of the 50 Hz oscillator shown in Figure 5.20 and discussed in Section 5.4.

7.2.1   50 Hz Oscillator

The 50 Hz oscillator had a very simple design goal: achieve performance matching an L4C on a fused silica resonator. The readout choice was not fixed, but the cavity readout was the most sensitive available and was working in the lab, so this was chosen as the target. As FEMTOprint made the oscillators, fused silica was the only available material. From prior experience, we determined that a Q factor of 10,000 was a safe lower bound as an estimate for calculations. Combining this together, a 50 Hz f_0 was needed with the readout method and a test mass of 3 g to meet the noise requirement. The flexures could fit within a 1" optic, so this was the target design space, as we had experience handling such geometries. Together, this gave the initial design concept.

The flexure width was chosen at 180 μm as it gave the optimum TED on resonance while providing a good noise floor across the whole control bandwidth of 0.1 Hz to 100 Hz. It was then calculated that 18 mm long flexures were needed to achieve a safety factor of 5 under 1 g of load. The large safety factor allowed the device to work as a vertical inertial sensor. It was calculated then that approximately 2 mm of flexure height would be needed to give the correct resonance frequency. The tolerances were fixed by

what FEMTOprint could achieve. The criterion of maximum deviation of thickness was used to define the tolerance on the flexures, as discussed in Section 4.5. A tolerance of +/-10 μm could be achieved in this definition. It was determined that even if all flexures were as thick as the definition allowed, a f_0 of 57 Hz would be achieved, which was deemed acceptable.

The geometry was designed to fit on a 1" diameter, 6.35 mm thick wafer. We already had experience of how best to mount these samples, so we wanted to stick with this geometry. At the time of design, only bonding mirrors to the optic had been shown to work to achieve a high Q factor. Therefore, mirror bonding was chosen as the method to integrate a high reflectivity on the sample. Small, high-reflectivity mirrors were still left over from the coating test, so these could be used. As the samples were bathed with HF acids, the mirrors would likely not survive production, so they would have to be attached after manufacture. There had been previous attempts to protect coatings in manufacturing, but there were high failure rates where the acid pentrated the protection and often residuals of the protective layer were left after removing. As attaching mirrors post-manufacture had worked in the coating test, it was deemed a lower risk. Due to balancing concerns on the test mass, a small inlet was made for the mirror to sit, to keep the centre of mass towards the centre, and to help with alignment.

When trying to fit the flexures in the 1" geometry, it was found that the three pairs used in the prototypes simply did not fit. Therefore, an alternating flexure design was used similar to that used in the coating tester. With two sets of alternating flexures, the height was distributed evenly between them, leaving 1 mm height on each. When studying the stress, it was found that the stress was not evenly distributed across the flexure because of the curvature of the flexure. The length was therefore increased to 22 mm on the shorter edge to meet our goal of a safety factor of 5 under 1 g of load.

The alternating flexures had a much lower mode spacing than the parallelogram pairs. This caused a tip-tilt mode at 209 Hz. A similar mode had been seen in the coating tester and had not caused too many issues. The f_0 mode was linear and parallel to the input force.

A key element of this design was the minimisation of material etch. As the flexures were so long, this became critical, both from a cost perspective and preventing the etch time from being so long that it damaged the flexures. This was done using the inlets shown in Figure 5.20. Doing so made the overall design cost about the same as the prototype oscillators.

With the design finished, three demonstrators were produced. These were tested in a ringdown setup and showed a Q factor higher than expected, leading us to re-evaluate our estimates on flexure roughness. The designs now await the attachment of mirrors and, once done, will be tested for their noise floor in the setup described in Chapter 6. These measurements will be done in the near future.

7.2.2   15 Hz Oscillator

The second design study presented here is the case for the 15 Hz inertial sensor discussed in Section 5.5. It was known that the FEMTOprint approach would struggle with such a large material etch, so different manufacturing techniques were needed. The technique discussed in Section 5.5 was developed from the experience of working

with glass wafers at the IOF. Again f_0 was determined by the readout method and target noise performance. As a readout noise of $10^{-13}\mathrm{m(Hz)}^{-1/2}$ was expected from the compact state-of-the-art readout methods studied in Section 5.5, a f_0 of 15 Hz was needed to match the target noise performance. Using the same estimates on thermal noise as for the 50 Hz case, a mass of at least 1 g was needed. Ultimately, as a 6.35 mm thick wafer was used and the test mass needed to be large enough to host both bonding regions and coatings with several millimetres of separation, a final test mass of 3 g was used.

Flexure design was focused on survivability. Because the manufacturing method changed, the optimum flexure design ended up being to make them as thin as possible. Due to manufacturing limitations, the thickness target was decided as 100 µm. A length of at least 40 mm for flat flexures is required for a safety factor of 5 at 1 g load. The length requirement was reduced by studying how best to distribute stress across the flexure length. As shown in Figure 5.25, the stress is more evenly distributed across the flexure by having varying flexure height across their length. The overall height was then tuned to achieve the correct f_0. The key tolerances studied were flexure width and alignment. The tolerance on the flexure width was estimated for the manufacturing method as +/-10 µm. The target frequency was tuned to below 15 Hz to mean that even in its highest thickness, the f_0 was less than 15 Hz, and that it could survive with safety factor 4.5 under load on the minimum thickness allowed by the tolerance. The alignment was much more difficult to define, but the geometry was configured such that even when misaligned to their maximum extent they would fit together.

As the flexure design was already very mature when it came to geometrical implementation, only a little was changed at this step. It was decided to use metallic coatings on the test mass, as an extremely high reflectivity coating was unnecessary for the proposed readout schemes. Because no acid etch was used, this could be applied to the test mass before bonding with the flexures. A coating diameter of 5 mm was specified, which left 3 mm to any bonding regions for the flexures. The test mass was designed to allow each flexure's maximum possible bonding area. This is because the success chance of the direct bonding method is related to the area available to bond. Furthermore, stringent requirements surface qualities on the surfaces were defined as less than 0.5 nm rms. The test mass and outer-frame were made from the same wafer to minimise differences in wafer thickness. A flat side was added to the outer frame to allow the device to stand freely. This meant that clamps were no longer needed. Figure 5.24 shows the produced geometry.

The first prototype has been produced, shown in Figure 5.27. This highlighted some issues regarding alignment and cleaning that must be overcome, but in the near future it is expected that a full production run of 3 samples will be undertaken.

7.3   Chapter Summary

This chapter has shown how to approach the design of inertial sensors. A robust method of designing such sensors is presented, which, when followed, should produce something that meets design criteria. Other starting limitations can also be used, and these methods can be used as guides for them, but different optimal design points will be reached.

Designing such sensors is challenging, as so many factors effectively cross couple. The iterative design processes discussed here allow a linear exploration through the field of possibility. Following such processes is time-consuming. These kinds of cross-coupled, complex problems are exactly what large parameter space optimisation routines excel at solving [130], and the development of such a tool for this purpose is something I believe would be well worth approaching within the field.

A key difficulty with these design processes is the combination of parametrically optimal outcomes fighting with more undefined limits such as manufacturability. The more loose rules I have followed to define these steps are listed below.

- Consult with the manufacturers early, and often, what looks easy to make is frustratingly uncorrelated with what is easy to make (and sometimes this is a pleasant surprise).

- There is no point in making gains in thermal noise when readout noise is limited, and vice versa.

- Exploring flexure behaviour on simplified models can allow us to get very close to a suitable design before considering how it fits into the rest of the geometry. This often makes calculations and modelling easier.

- Define the frequency of the fundamental mode as early as possible and try to stick to it. It defines how the system reacts to readout and the flexure geometry needed. Changing it means going back to design step one.

- 3D printing designs is a cheap and easy way to test mounting solutions. If you cannot get the plastic version in and out of the mount without touching the test mass or applying a large force, you cannot get the glass samples out, either.

We then explore how this design process led to different sensors' designs throughout this project. The two sensors studied have similar inputs, the difference being the expected readout noise floor. Very quickly, the designs diverge, requiring entirely different geometries, manufacturing methods, and areas of concern throughout the design phases. The fact that the design method can accommodate this shows its robustness and can be applied to make all kinds of inertial sensors. There are some limitations to the approach. As the method is meant to produce sensors for research applications where the number of sensors is few, cost per sample is not a stringent criteria of the method. It is only assessed if the design is financially feasible within a given budget. Financial cost has a much greater impact for commercial use, which this design approach cannot optimise. In addition, the tool is not so useful for developing new techniques of manufacture and readout, which often have a more hands-on approach. Instead, it is intended to use known techniques to produce suitable sensors for a given application. Still, as a tool for designing inertial sensors for research applications, following the process can fill many niches in the field of inertial sensing in scientific experiments.

Chapter 8

Conclusions

Throughout this thesis, I have presented my ongoing work to make high-precision, compact inertial sensors. Developing devices like this will be integral to the next generation of gravitational wave detectors. Compact sensors can be used to better understand the responses of the seismic isolation stages and allow better active control. The next generation of gravitational wave detectors will rely on developments like these compact sensors to reach sufficient sensitivity to detect gravitational waves in the lower limit of their frequency band. Achieving this goal will ultimately lead to a better understanding of some of the most violent phenomena in the universe and so test our understanding of physics in some of the most extreme limits.

Inertial sensors should be custom designed for specific applications, with considerations on size, environment, sensitivity, and readout method. To this end, a set of design rules for compact sensors have been presented. These rules can make an inertial sensor suited for various applications.

Throughout this project, drum-style oscillators have been designed and tested. An initial batch of oscillators offered a promising start. Although the initial batch never reached its predicted Q factors limits, there was still high enough Q factors to produce low thermal noise designs. Coating samples with IBS proved to be highly detrimental to the device's Q factor. Alternative means of applying a highly reflective coating through mirror bonding and metallic coatings have been shown not to reduce Q factors. Furthermore, later oscillator designs have been produced with thermal noises that would allow the compact devices to match several state-of-the-art designs when combined with a sufficiently precise readout. The readout of these sensors currently requires some state-of-the-art techniques to be comparable to other high-precision seismometers.

A design has been produced with a lower resonance frequency that could achieve excellent performance when combined with several readout techniques within the community. Producing such oscillators requires a different manufacturing technique due to scaling problems with the etching method. A proposed process has been discussed. If the pieces can be made this way, inertial sensors can be made in a compact form with simple to use readout. Such devices would see applications in wide-ranging fields concerned about seismic disturbances.

A precision readout technique has been developed to read out the motion of the high-frequency oscillators with sufficient sensitivity. The scheme was initially limited by laser frequency noise at low frequencies. This limit has been overcome by using

a frequency reference cavity. However, this required alternative means of keeping the light field resonant with the cavity, which proved problematic. In a seismically quieter environment, AOMs could shift the frequency to keep sensors at their operating point. In an environment with typical seismic disturbances, mechanical PZT stacks had to be used to change cavity length to keep the field resonant. These PZT stacks struggled to keep up with the large movements of the oscillators. The method was able to suppress frequency noise at frequencies below 20 Hz, but did introduce a lot of extra control noises above this frequency. Still, the method is promising for producing seismic sensors capable of high performance in the control bandwidth of gravitational wave detectors.

A huddle test was performed on some initial prototype inertial sensors. So far, the sensors have yet to use the best performing oscillators, but they will be tested soon. A few samples used to test coating methods have been assembled into inertial sensors. Sadly one of these sensors did not perform as expected due to extra tilt coupling through an unbalanced test mass. Using a mounted geophone, coherent subtraction could be achieved between the OIS and the geophone. The results did not hit the expected noise floor of the readout or the thermal noise floor. Therefore, it is likely that some residual motion differences between the sensors prevented a true assessment of its noise floor. Regardless a promising noise measurement has been made, which, even with the above limitations, was only a factor of 5 worse performing than some of the state-of-the-art sensors available today.

Experiences developed in this project lead to a design process for inertial sensors. The design process typically requires iterative steps, bringing the device towards a targeted goal with each iteration. This problem is because of the large parameter space where changes to one parameter force every other parameter to change. The complex and cross-coupled nature of the design parameter space hints at the need for a large-scale optimisation tool within the community. Even a crude version of such a tool could save a lot of design effort running through FEA simulations and may even lead to better inertial sensor configurations being found.

Overall, compact sensors of this kind will be an integral tool in the future of gravitational wave astronomy. The working principle of the design has been demonstrated here and by other groups in the field [101]. The next key target will be to produce sensors useful for current gravitational wave detectors. There is a desire to measure the motion in the suspension chain. Current plans are to use displacement sensors. Combining this with compact inertial sensors would be an excellent opportunity to demonstrate their potential for use in gravitational wave observatories. Concurrently, as the plan for seismic isolation of next-generation gravitational wave detectors is developed, design specifications for inertial sensors should be explored. With specifications in mind and the tools shown here, a suite of seismic sensors could be incorporated to allow the control systems to meet their stringent requirements on transmitted motion.

Appendices

Appendix A

A Displacement Damped Harmonic Oscillator

Many of the oscillators in this work are damped with proportion to their displacement rather than velocity. To model the behaviour of such an oscillator is derived here.

For a displacement damped oscillator the equation of motion is given by

$$f_{\text{ext}} = mx_{\text{m}} + k(1 - i\phi)x_{\text{m}}. \tag{A.1}$$

where f_{ext} is the force of the spring compressing under motion from ground

$$f_{\text{ext}} = k(1 - i\phi)x_{\text{g}}, \tag{A.2}$$

and ϕ is the fractional energy dissipation per radian of oscillation, or the inverse of the Q factor. Following through as with the displacement case it is found

$$\omega^2 X_{\text{g}} = -\omega^2 \Delta X + \omega_0{}^2 \Delta X + \frac{i\omega_0^2}{Q}\Delta X, \tag{A.3}$$

where the capital denotes the function is Fourier transformed and therefore a function of ω. Rearranging for the ratio of the two displacements results in the transfer function

$$\frac{\Delta X}{X_{\text{g}}} = \frac{-\omega^2}{\left(\omega^2 - \omega_0{}^2 - \frac{i\omega_0^2}{Q}\right)}. \tag{A.4}$$

The difference between the displacement damped and velocity damped case is subtle. For high Q oscillation The change effects only the linear term in the denominator which shapes near resonance. As in this limit $\omega \approx \omega_0$ the effect is very minimal, even with a Q factor of 10 and resonance frequency of $1\,\text{Hz}$, as can be seen in Figure A.1
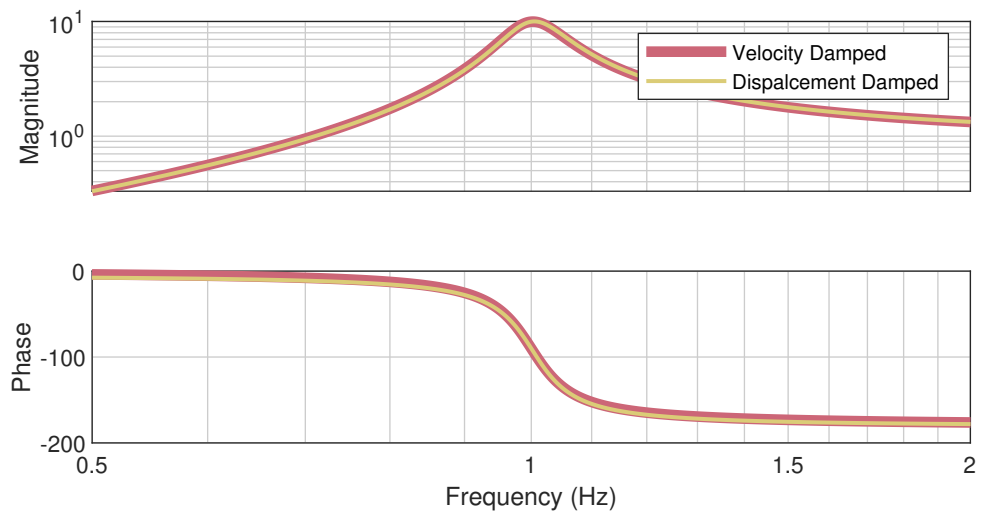
Figure A.1: A comparison of the transfer functions of a displacement and velocity damped inertial sensor. The difference between the two is minimal, and only when extreme sensitivity of peak shape is relevant should it be considered, the line of the velocity damped curve has to be broadened so that it can even be seen.

Appendix B

Thermoelastic Damping of Thin Plates Under Parallelogram Mode of Oscillation

Here a derivation for the loss due to thermoelastic damping during oscillation in a mode like that shown in Figure 1.4 (a) is presented. A general solution for the thermoelastic damping of an arbitrary isotropic body was first derived by Chadwick [76], who wrote the specific solutions for thin rods and plates under flexural motion, where the edges are fixed and the centre oscillates [131]. Later Norris and Photiadis [75] presented a robust method for predicting Quality factors from thermoelastic damping in MEMS resonators. They then presented solutions for an infinite plane and a thin beam undergoing flexural motion. This derivation uses the results they derive in Section 3 and follows the method they use in Section 4, but with the geometry adapted to the parallelogram mode of oscillation.

To begin, consider the definition of Q factor from Equation 2.7. The two necessary quantities are therefore; the energy dissipated per cycle and the energy stored in the system. For this deviation the axis of oscillation from Figure 1.4 is z, the x direction parallel to flex height from Figure 4.4 and the y axis parallel to flex length. It is assumed that an isotropic material such as fused is used and that the plate is significantly taller and longer than the flexure width, such that it can be treated as an infinite plan, with no stresses normal to the plane.

From Equation 4.5 in Norris [75], the instantaneous heat dissipation rate per unit length is given by

$$\dot{\mathcal{E}}(x) = -d_{\text{height}} \sum_{n=0}^{\infty} \frac{\omega^2 \tau_{\text{n}}}{\left(1 - \omega^2 \tau_{\text{r}} \tau_{\text{n}}\right)^2 + \omega^2 \tau_{\text{n}}^2} \frac{T_{\text{a}}}{C_{\text{p}}} \overline{\langle \psi_{\text{n}}(z), \alpha \cdot \sigma(\mathbf{x}, z) \rangle}^2 \qquad \text{(B.1)}$$

where $\mathcal{E}$ is the energy density per unit length on the plane, C_{p} is the volumetric heat capacity at constant pressure, T_{a} is the ambient temperature, α is the thermal expansion coefficient tensor, σ is the stress tensor, and the brackets indicate the inner product $\langle f, g \rangle = \int dV f(\mathbf{x}) g(\mathbf{x})$, with $\mathbf{x}$ being the position vector, τ_{r} is subtlety different from the relaxation time of the thermal excitation τ_{n}. τ_{r} is a measure of how long after excitation a steady state solution is applicable as defined in [132]. In the case that the flexure width is much large than the mean phonon free path this becomes negligible.

ψ_n refers to the eigenmode of heat equation which for a thin plate are given by

$$\psi_n = \sqrt{\frac{2}{d_{\text{fwidth}}}} \sin\left((2n+1)\frac{\pi z}{d_{\text{fwidth}}}\right), \tag{B.2}$$

in the case that the heat only depends on z displacement and that the modes are asymmetric around the middle of the plane [75]. τ_n refers to the relaxation time of the heat across the beam corresponding to this mode, this is given by

$$\tau_n = \frac{d_{\text{fwidth}}^2}{(2n+1)^2\pi^2\chi}. \tag{B.3}$$

Following Norris's method, as a the oscillation is treated as a thin plate, the elements of the stress tensor normal to the z axis are all 0. Equation 4.1 from Norris then gives

$$\alpha \cdot \sigma = \frac{z}{1-\nu}\left[E\alpha\kappa_{xx} + E\alpha\kappa_{yy}\right], \tag{B.4}$$

where the isotropic qualities of fused silica have been used to simplify the formula, ν is the Poisson's ratio and E is the Young's Modulus. κ is the curvature tensor of the plane defined by

$$\kappa_{ij} = -\frac{\partial^2 w(x,y)}{\partial x_i \partial x_j} \tag{B.5}$$

where $w(x,y)$ is the transverse deflection of the centre plane of the flexure. Here a significant simplification over the full thin plane solution can be made as both the derivatives in the x and z direction are zero, leaving the only none zero element of the tensor as κ_{yy}.

The inner product from Equation B.1 must now be solved. The only z dependence in the inner product is

$$\langle \psi_n(z), \alpha \cdot \sigma(\mathbf{x}, z)\rangle = \frac{E\alpha\kappa_{yy}}{1-\nu}\langle \psi_n(z), z\rangle. \tag{B.6}$$

A new function is defined here to make the solution more tidy

$$f_n \equiv \langle \psi_n, z\rangle^2 / \langle z, z\rangle = 96/[(2n+1)\pi]^4 \tag{B.7}$$

combining all this back into equation B.1 gives

$$\dot{\mathcal{E}}_{\text{loss}}(\boldsymbol{x}) = -\overbrace{\frac{d_{\text{height}}\langle z, z\rangle E\kappa_{yy}^2}{(1-\nu)^2}}^{\mathcal{E}_{\text{Diss}}}\frac{ET_a\alpha^2}{C_p}\sum_{n=0}^{\infty}f_n\frac{\omega^2\tau_n}{1+\omega^2\tau_n^2} \tag{B.8}$$

where the equation has been arranged such that the 3 parts contain the frequency response, a dimensionless rate constant $ET_a\alpha^2/C_p$, and a term which has units of energy per length, $\mathcal{E}_{\text{Diss}}$, which is effectively the measure of the dissipating energy. All of the time dependence now lies in the curvature tensor and so this can be integrated over one oscillation to give the energy loss per cycle.

The total potential energy per unit length is also given by Norris [75] Equation (4.3)

$$\mathcal{E}_{\text{PE}} = d_{\text{height}} \frac{\langle z, z \rangle}{2} \frac{1}{1 - \nu^2} \left(E \kappa_{yy}^2 \right) \tag{B.9}$$

and using the viral theorem

$$\bar{\mathcal{E}}_{\text{PE}} = -\bar{\mathcal{E}}_{\text{KE}} \tag{B.10}$$

where KE denotes the kinetic energy. Hence,

$$\mathcal{E} = \bar{\mathcal{E}}_{\text{KE}} - \bar{\mathcal{E}}_{\text{PE}} = 2\bar{\mathcal{E}}_{\text{PE}} = \frac{\omega}{\pi} \int_0^{2\pi/\omega} \mathcal{E}_{\text{PE}} dt. \tag{B.11}$$

Looping back to the original definition of Q factor, all terms are now known so

$$Q_{\text{TED}}^{-1} = \frac{\int_0^{d_{\text{flength}}} \int_0^{2\pi/\omega} \dot{\mathcal{E}} \, dt \, dy}{2\omega \int_0^{d_{\text{flength}}} \int_0^{2\pi/\omega} \mathcal{E}_{\text{PE}} \, dt \, dy}. \tag{B.12}$$

The the only spacial or temporal variation is in κ_{yy} and therefore the integrals cancel once factors are taken out. This leaves

$$Q_{\text{TED}}^{-1} = \frac{1 + \nu}{1 - \nu} \frac{E T_a \alpha^2}{C_p} \sum_{n=0}^{\infty} f_n \frac{\omega \tau_n}{1 + \omega^2 \tau_n^2}. \tag{B.13}$$

The infinite sum decays rapidly as order n^6 for $\omega << \tau$ and so can well be approximated as 0^{th} order leaving

$$Q_{\text{TED}}^{-1} = \frac{96}{\pi^4} \frac{1 + \nu}{1 - \nu} \frac{E T_a \alpha^2}{C_p} \frac{\omega \tau_0}{1 + \omega^2 \tau_0^2}. \tag{B.14}$$

As the equation is more an approximation for Q factor, the first two fraction in this equation can be approximated to 1, to give the estimate

$$Q_{\text{TED}}^{-1} = \frac{E T_a \alpha^2}{C_p} \frac{\omega \tau_0}{1 + \omega^2 \tau_0^2}. \tag{B.15}$$

Appendix C

Ringdown Plots

This Chapter include the raw plots used measured in the ringdown section. They are split into their three batches. The first set in Figure C.1 is the sample initial testing was done on. The second set in Figure C.2 is the batch of oscillators after a treatment in HF bath after initial measurements showed they were not ringing. The set in Figure C.3 are the results from the coating test samples before coatings were applied. Figure C.4 shows the results after UV glueing of mirrors to the samples. Figure C.6 shows the rings form the Hz resonators.
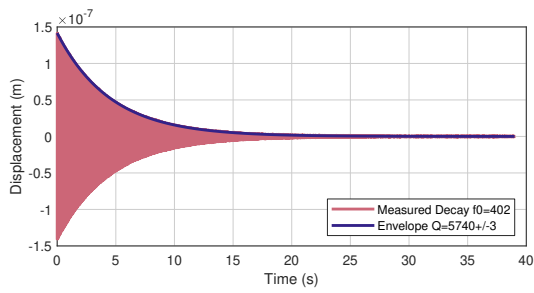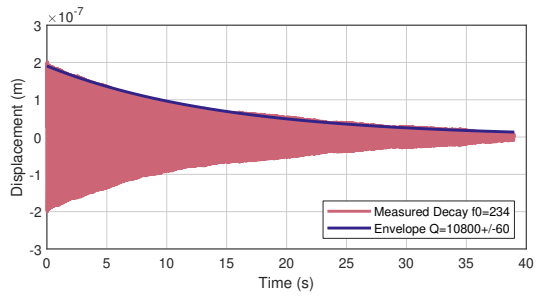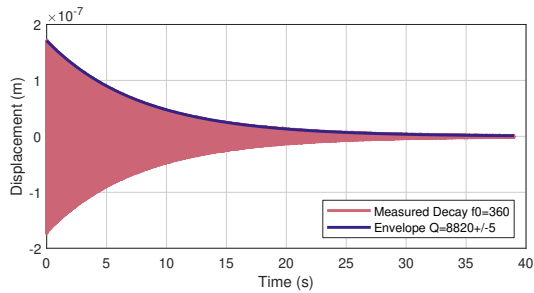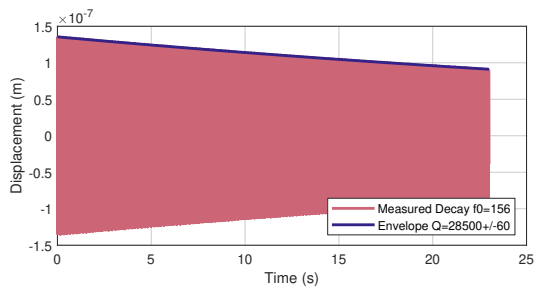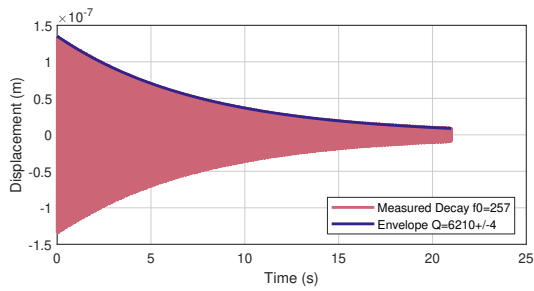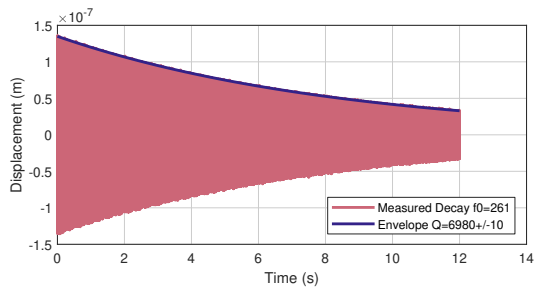
Figure C.1: The Ringdowns form the first batch of oscillators produced.
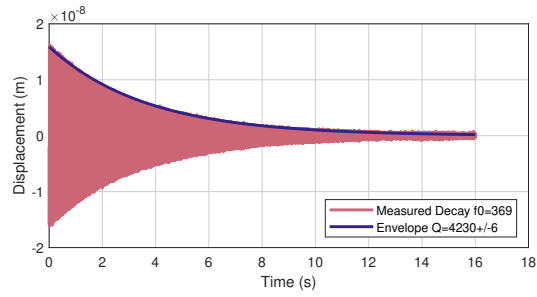
(a) C2S01
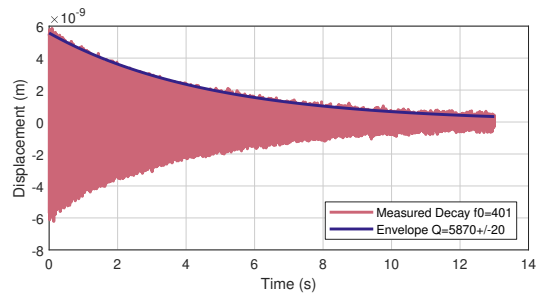
(b) C2S04
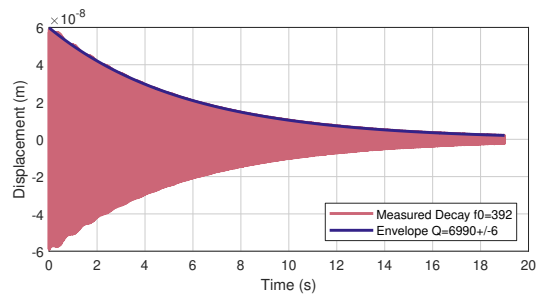
(c) C2S05

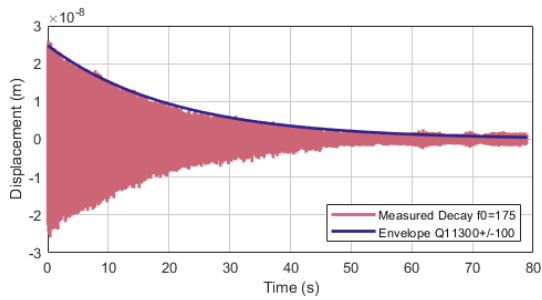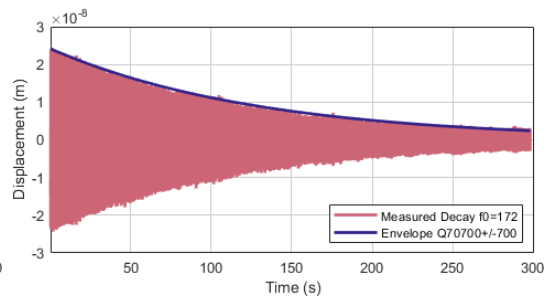Figure C.2: The Samples from the second batch after their return form femtoprint and treatment in the HF bath.
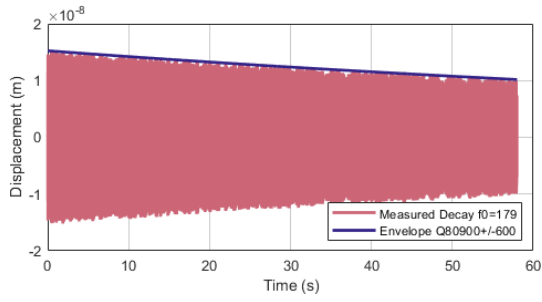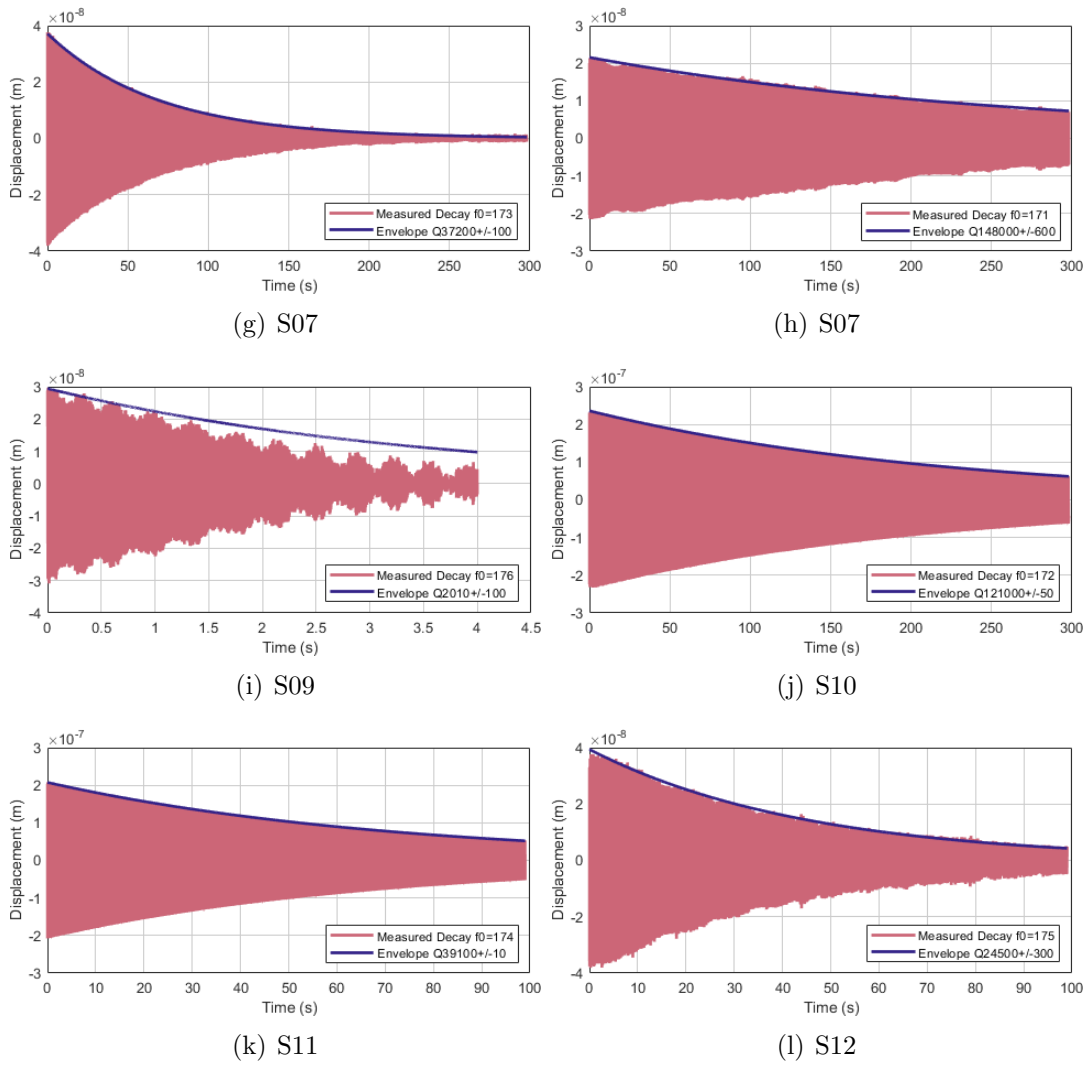
(a) S01

(b) S02

(c) S03

(d) S04

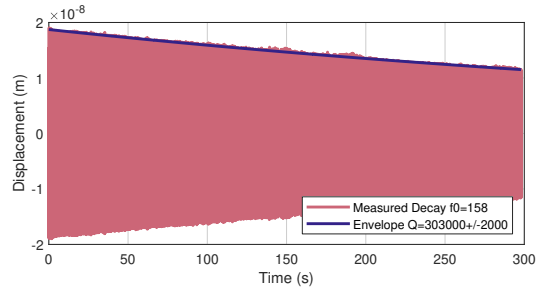(e) S05

(f) S06
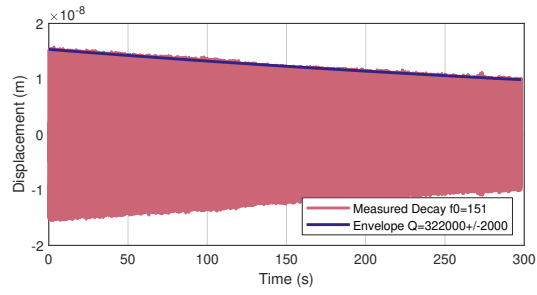
(g) S07

(h) S07

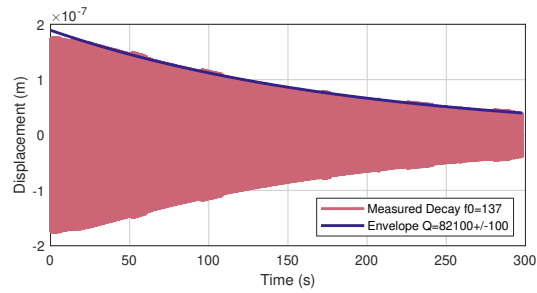(i) S09

(j) S10

(k) S11

(l) S12
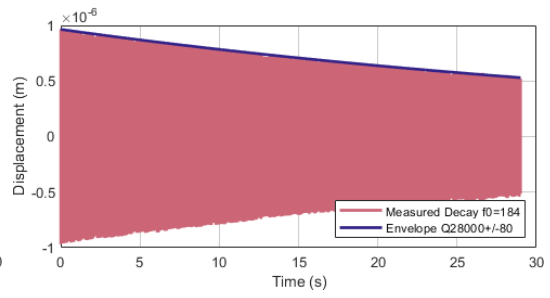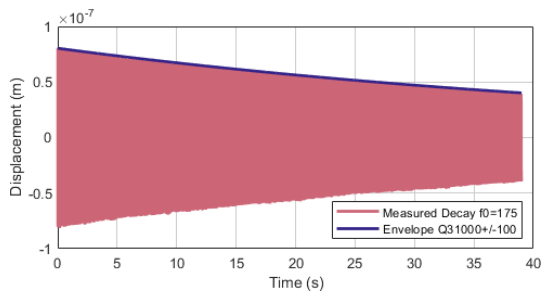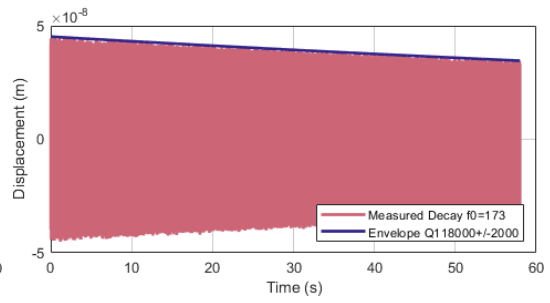
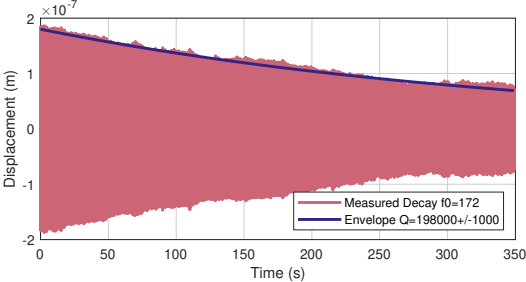Figure C.3: The Rings from coating testers upon arrival
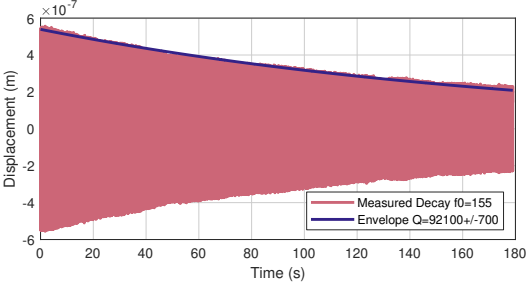
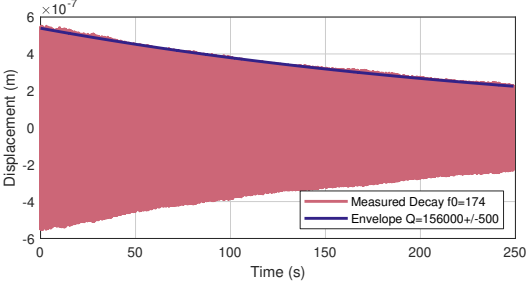(a) S01

(b) S04

(c) S07

Figure C.4: The Coating test samples after they were bonded with UV glue to mirrors.

(a) S06

(b) S08

(c) S12

Figure C.5: The Coating test samples after applying a metallic coating. The specific metal coating was a silver coating with a protective layer.
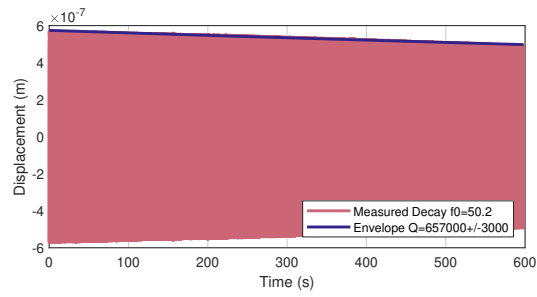
(a) B01
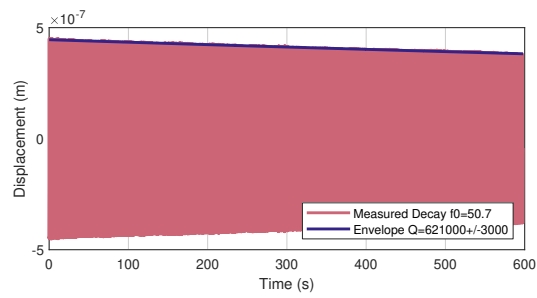
(b) B02
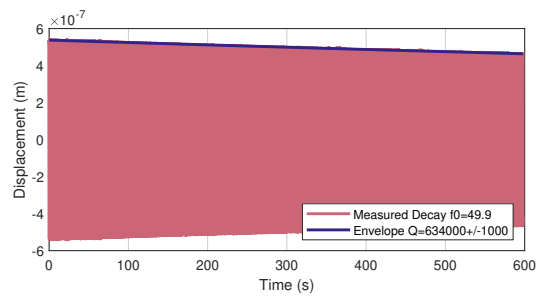
(c) B03
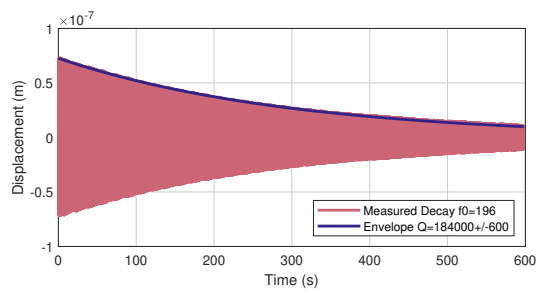
(d) Higher Order of B03

Figure C.6: The rings of the 50Hz oscillators.

Appendix D

Acceleration Equivalents of Inertial Noise Plots

This section contains the acceleration equivalents of most graphs in acceleration units, they are linked to their inertial equivalents in the main text.
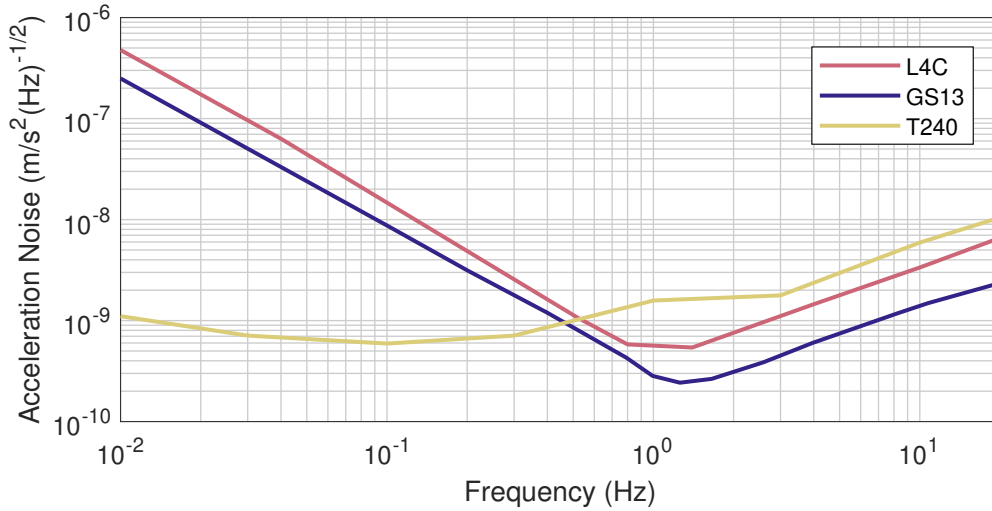
Figure D.1: (acceleration equivalent of D.1) The noise performance of several of the inertial ground sensors used at LIGO.
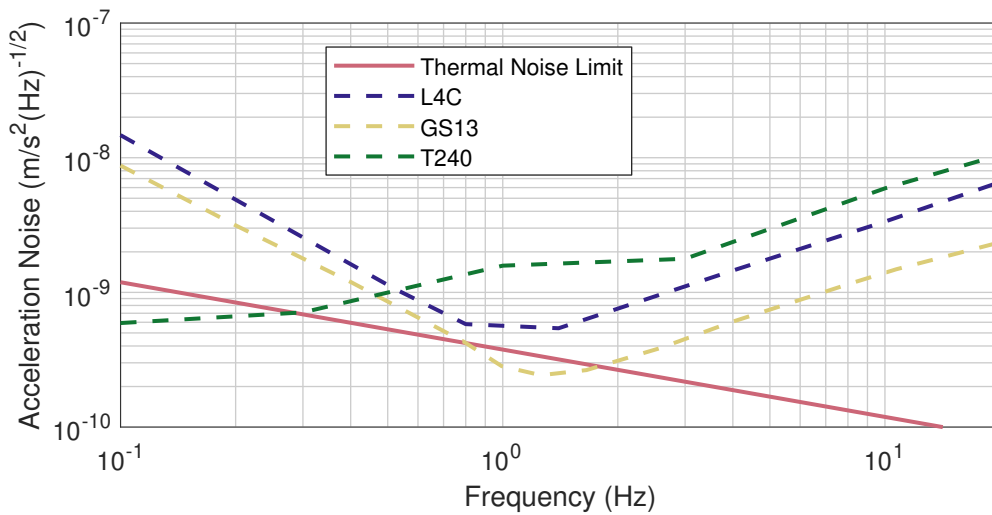
Figure D.2: (Acceleration Equivalent of 5.23) A comparison of the projected thermal noise of the 50 Hz oscillator with several state-of-the-art sensors used at LIGO. The thermal noise of the oscillator surpasses or matches all these sensors across most of the relevant control bandwidth for gravitational wave detectors.
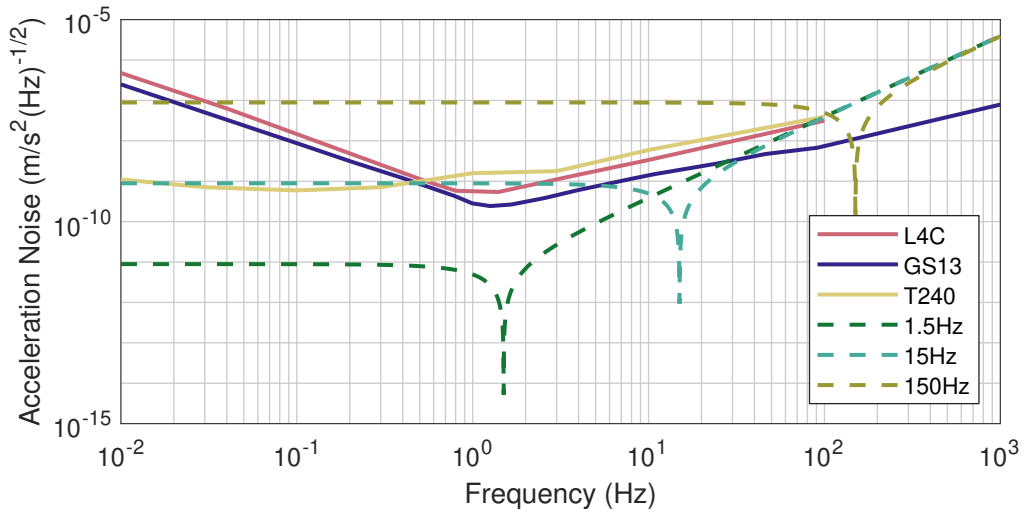
Figure D.3: A comparison of the acceleration noise floors of oscillators with differing resonance when a Moku:Lab Phasemeter (or a 16 bit CDS phasemeter) self noise is the limiting noise source. We see that to achieve comparable noise performance to other high-performance sensors, we would need a resonance frequency below 15Hz.
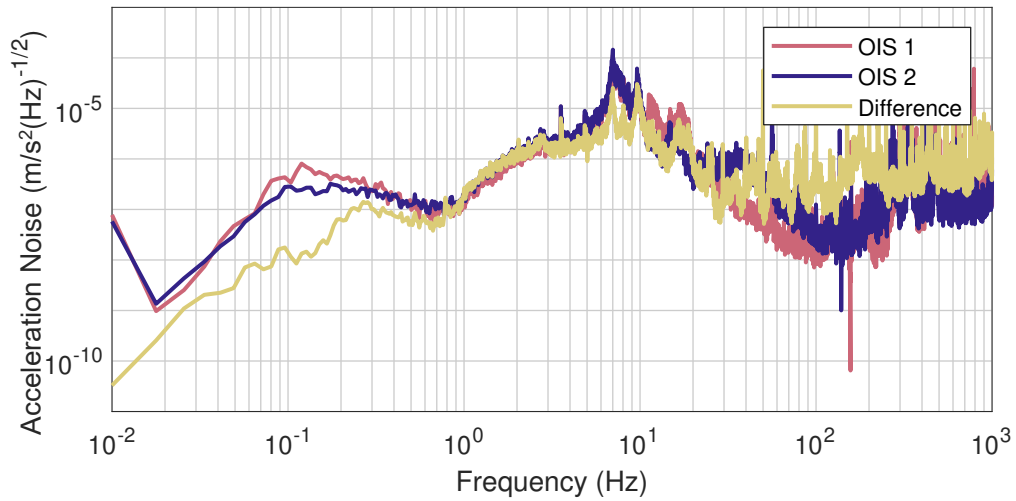
Figure D.4: (Acceleration equivalent of 6.28) A comparison of measured noise from two OIS and the geophone mounted outside the vacuum. The signals are in good agreement in frequencies where this should be expected, below the isolation feet resonance at 8 Hz.
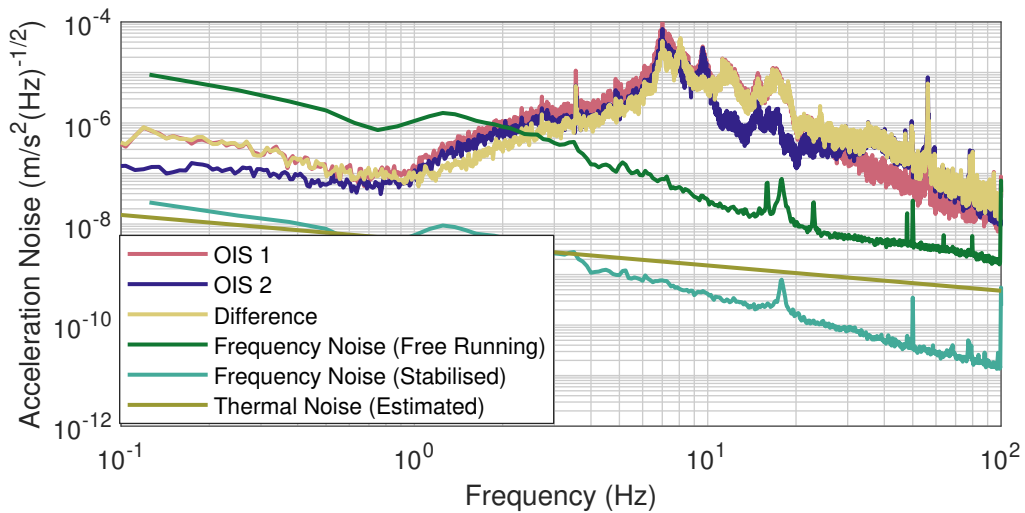
Figure D.5: (Acceleration equivalent of 6.29) Measured noise Inertial noise of the OIS and the "naive" differential signal, with the fundamental limiting noise sources of the sensor, the stabilised laser noise and the thermal noise of the oscillator. The subtraction is very limited due as small calibration differences would limit the subtraction of common noise in this setup, hence the MCCS2 routine is necessary.
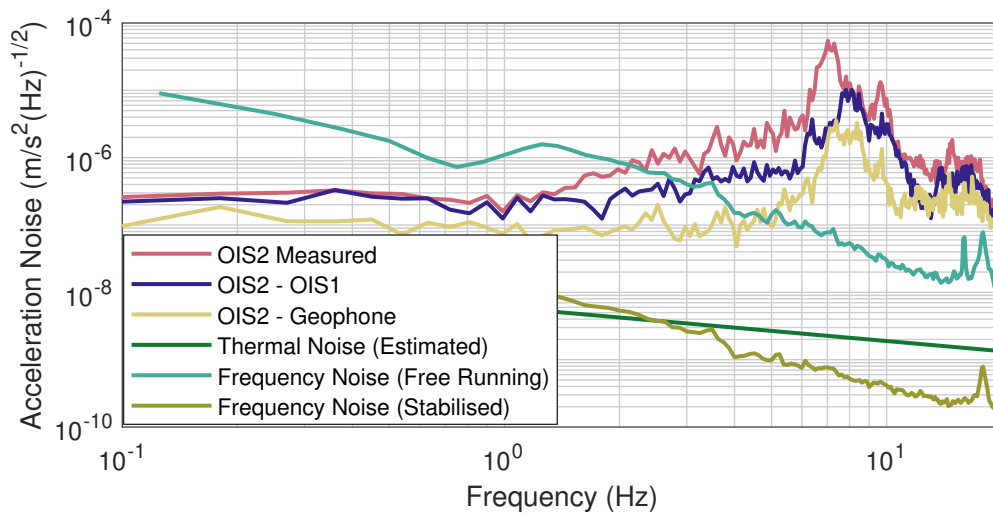
Figure D.6: (acceleration equivalent of 6.30) The noises of the OIS when coherent subtraction is performed with different sensors using the method described by [125]. This is used to subtract motion and measure the actual noise floor, despite calibration differences between the sensors. OIS 2 is used as the sensor and OIS as the subtraction to produce the blue line, which showed only limited subtraction. When a ground mounted geophone was used for subtraction much greater reduction was measured, but still sits a factor of 5-10 above the fundamental noise sources. The remaining motion is likely platform motion that is independent of the geophone motion.
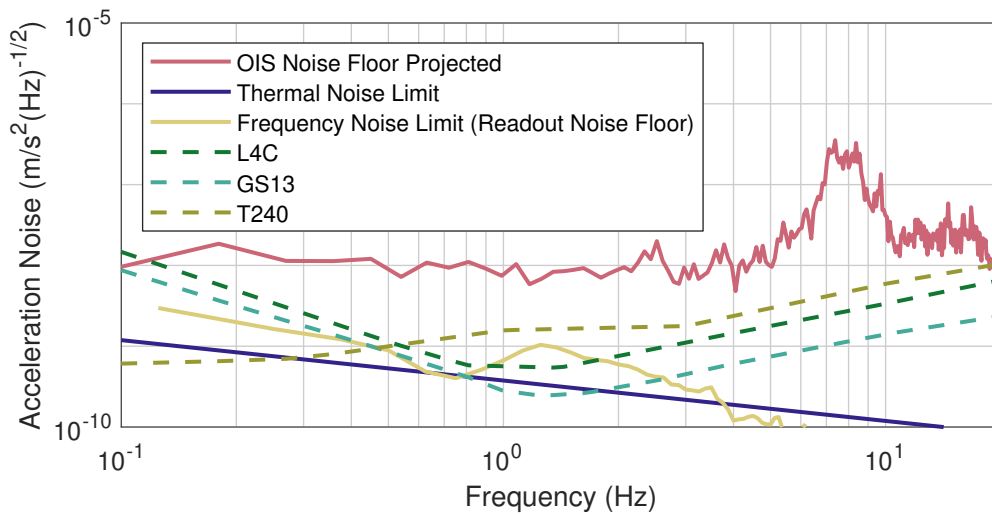
Figure D.7: (acceleration equivalent of D.7) A comparison of the projected noise form this readout scheme onto the 50 Hz oscillator, along with its thermal noise and frequency noise against several of the state of the art sensors used in LIGO. The sensors theoretical noise floor defined by the readout noise and thermal noise surpasses that of many of the state of the art sensors. The measured readout noise does not reach that of the theoretical limits, but still is approaching close to the performance of the state of the art sensors.

Bibliography

[1] B.P. Abbott et al. "GW150914: The Advanced LIGO Detectors in the Era of First Discoveries". In: *Physical Review Letters* 116.13 (Mar. 2016), p. 131103. DOI: `10.1103/physrevlett.116.131103`.

[2] J Aasi et al. "Advanced LIGO". In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001. DOI: `10.1088/0264-9381/32/7/074001`.

[3] B.P. Abbott et al. "GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence". In: *Physical Review Letters* 116.24 (June 2016), p. 241103. DOI: `10.1103/physrevlett.116.241103`.

[4] B. 0.167emP. Abbott et al. "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs". In: *Physical Review X* 9.3 (Sept. 2019), p. 031040. DOI: `10.1103/physrevx.9.031040`.

[5] B. P. Abbott et al. "Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A". In: *The Astrophysical Journal* 848.2 (Oct. 2017), p. L13. DOI: `10.3847/2041-8213/aa920c`.

[6] F Acernese et al. "Advanced Virgo: a second-generation interferometric gravitational wave detector". In: *Classical and Quantum Gravity* 32.2 (Dec. 2014), p. 024001. DOI: `10.1088/0264-9381/32/2/024001`.

[7] E. Troja et al. "The X-ray counterpart to the gravitational-wave event GW170817". In: *Nature* 551.7678 (Oct. 2017), pp. 71–74. DOI: `10.1038/nature24290`.

[8] Vishal Baibhav et al. "Gravitational-wave detection rates for compact binaries formed in isolation: LIGO/Virgo O3 and beyond". In: *Physical Review D* 100.6 (July 2019), p. 064060. DOI: `10.1103/physrevd.100.064060`.

[9] T Akutsu et al. "KAGRA: 2.5 generation interferometric gravitational wave detector". In: *Nature Astronomy* 3.1 (Jan. 2019), pp. 35–40. DOI: `10.1038/s41550-018-0658-y`.

[10] P. Kwee et al. "Stabilized high-power laser system for the gravitational wave detector advanced LIGO". In: *Optics Express* 20.10 (Apr. 2012), p. 10617. DOI: `10.1364/oe.20.010617`.

[11] S M Aston et al. "Update on quadruple suspension design for Advanced LIGO". In: *Classical and Quantum Gravity* 29.23 (Oct. 2012), p. 235004. DOI: `10.1088/0264-9381/29/23/235004`.

[12]   N A Robertson et al. "Quadruple suspension design for Advanced LIGO". In: *Classical and Quantum Gravity* 19.15 (July 2002), pp. 4043–4058. DOI: `10.1088/0264-9381/19/15/311`.

[13]   A V Cumming et al. "Design and development of the advanced LIGO monolithic fused silica suspension". In: *Classical and Quantum Gravity* 29.3 (Jan. 2012), p. 035003. DOI: `10.1088/0264-9381/29/3/035003`.

[14]   Alberto Stochino et al. "The Seismic Attenuation System (SAS) for the Advanced LIGO gravitational wave interferometric detectors". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 598.3 (Jan. 2009), pp. 737–753. DOI: `10.1016/j.nima.2008.10.023`.

[15]   F Matichard et al. "Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance". In: *Classical and Quantum Gravity* 32.18 (Aug. 2015), p. 185003. DOI: `10.1088/0264-9381/32/18/185003`.

[16]   F. Matichard et al. "Advanced LIGO two-stage twelve-axis vibration isolation and positioning platform. Part 1: Design and production overview". In: *Precision Engineering* 40 (Apr. 2015), pp. 273–286. DOI: `10.1016/j.precisioneng.2014.09.010`.

[17]   F. Matichard et al. "Advanced LIGO two-stage twelve-axis vibration isolation and positioning platform. Part 2: Experimental investigation and tests results". In: *Precision Engineering* 40 (Apr. 2015), pp. 287–297. DOI: `10.1016/j.precisioneng.2014.11.010`.

[18]   M G Beker, J F J van den Brand, and D S Rabeling. "Subterranean ground motion studies for the Einstein Telescope". In: *Classical and Quantum Gravity* 32.2 (Dec. 2014), p. 025002. DOI: `10.1088/0264-9381/32/2/025002`.

[19]   L Naticchioni et al. "Characterization of the Sos Enattos site for the Einstein Telescope". In: *Journal of Physics: Conference Series* 1468 (Feb. 2020), p. 012242. DOI: `10.1088/1742-6596/1468/1/012242`.

[20]   Evan D. Hall et al. "Gravitational-wave physics with Cosmic Explorer: limits to low-frequency sensitivity". In: (Dec. 2020). arXiv: `2012.03608 [gr-qc]`.

[21]   S. Braccini et al. "An improvement in the VIRGO Super Attenuator for interferometric detection of gravitational waves: The use of a magnetic antispring". In: *Review of Scientific Instruments* 64.2 (Feb. 1993), pp. 310–313. DOI: `10.1063/1.1144249`.

[22]   Jonathan J Carter et al. "Particle swarming of sensor correction filters". In: *Classical and Quantum Gravity* 37.20 (Sept. 2020), p. 205009. DOI: `10.1088/1361-6382/abb32a`.

[23]   J. Kissel. "Calibrating and Improving the Sensitivity of LIGO". PhD thesis. Jan. 2010.

[24]   L Barsotti, M Evans, and P Fritschel. "Alignment sensing and control in advanced LIGO". In: *Classical and Quantum Gravity* 27.8 (Apr. 2010), p. 084026. DOI: `10.1088/0264-9381/27/8/084026`.

[25]   A Staley et al. "Achieving resonance in the Advanced LIGO gravitational-wave interferometer". In: *Classical and Quantum Gravity* 31.24 (Nov. 2014), p. 245010. DOI: 10.1088/0264-9381/31/24/245010.

[26]   Denis V Martynov. "Lock acquisition and sensitivity analysis of advanced LIGO interferometers". PhD thesis. 2015. URL: https://thesis.library.caltech.edu/8899/.

[27]   Laurence Datrier and Sheila Dwyer. "Impact of wind speed on BNS range and duty cycles at LHO for the first two months of O3 DCC: T1900433". In: *https://dcc.ligo.org/LIGO-T1900433* (July 2019).

[28]   Michael Ross. *Updated Lock vs Wind.* https://alog.ligo-la.caltech.edu/SEI/index.php?callRep=1544. URL: https://alog.ligo-la.caltech.edu/SEI/index.php?callRep=1544.

[29]   M. Evans et al. "Lock acquisition of a gravitational-wave interferometer". In: *Optics Letters* 27.8 (Apr. 2002), p. 598. DOI: 10.1364/ol.27.000598.

[30]   Adam J. Mullavey et al. "Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers". In: *Optics Express* 20.1 (Dec. 2011), p. 81. DOI: 10.1364/oe.20.000081.

[31]   F Acernese et al. "The variable finesse locking technique". In: *Classical and Quantum Gravity* 23.8 (Mar. 2006), S85–S89. DOI: 10.1088/0264-9381/23/8/s12.

[32]   A. Buikema andothers. "Sensitivity and performance of the Advanced LIGO detectors in the third observing run". In: *Physical Review D* 102.6 (Sept. 2020), p. 062003. DOI: 10.1103/physrevd.102.062003.

[33]   Kazuo MATSUMURA. "On regional characteristics of seasonal variation of shallow earthquake activities in the world". In: *Bulletin of the Disaster Prevention Research Institute* 36.2 (1986), pp. 43–98.

[34]   G. Hillers et al. "Seasonal variations of seismic velocities in the San Jacinto fault area observed with ambient seismic noise". In: *Geophysical Journal International* 202.2 (June 2015), pp. 920–932. DOI: 10.1093/gji/ggv151.

[35]   Krishna Venkateswara et al. *Windproofing LIGO: Improving low-frequency active seismic isolation using rotation sensors.*

[36]   Krishna Venkateswara et al. "Subtracting Tilt from a Horizontal Seismometer Using a Ground-Rotation Sensor". In: *Bulletin of the Seismological Society of America* 107.2 (Mar. 2017), pp. 709–717. DOI: 10.1785/0120160310.

[37]   M Punturo et al. "The Einstein Telescope: a third-generation gravitational wave observatory". In: *Classical and Quantum Gravity* 27.19 (Sept. 2010), p. 194002. DOI: 10.1088/0264-9381/27/19/194002.

[38]   David Reitze et al. "Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO". In: *2019 BAAS 51(7) 035* (July 2019). arXiv: 1907.04833 [astro-ph.IM].

[39]   Michele Maggiore et al. "Science case for the Einstein telescope". In: *Journal of Cosmology and Astroparticle Physics* 2020.03 (Mar. 2020), pp. 050–050. DOI: `10.1088/1475-7516/2020/03/050`.

[40]   Philip Jones et al. "Implications of the quantum noise target for the Einstein Telescope infrastructure design". In: *Physical Review D* 101.8 (Apr. 2020), p. 082002. DOI: `10.1103/physrevd.101.082002`.

[41]   Florian Amann et al. "Site-selection criteria for the Einstein Telescope". In: *Review of Scientific Instruments* 91.9 (Sept. 2020), p. 094504. DOI: `10.1063/5.0018414`.

[42]   Soumen Koley et al. "Surface and underground seismic characterization at Terziet in Limburg—the Euregio Meuse–Rhine candidate site for Einstein Telescope". In: *Classical and Quantum Gravity* 39.2 (Jan. 2022), p. 025008. DOI: `10.1088/1361-6382/ac2b08`.

[43]   M G Beker et al. "Newtonian noise and ambient ground motion for gravitational wave detectors". In: *Journal of Physics: Conference Series* 363 (June 2012), p. 012004. DOI: `10.1088/1742-6596/363/1/012004`.

[44]   Pau Amaro-Seoane and Lucıa Santamarıa. "DETECTION OF IMBHs WITH GROUND-BASED GRAVITATIONAL WAVE OBSERVATORIES: A BIOGRAPHY OF A BINARY OF BLACK HOLES, FROM BIRTH TO DEATH". In: *The Astrophysical Journal* 722.2 (Sept. 2010), pp. 1197–1206. DOI: `10.1088/0004-637x/722/2/1197`.

[45]   Pau Amaro-Seoane et al. "LION: laser interferometer on the moon". In: *Classical and Quantum Gravity* 38.12 (May 2021), p. 125008. DOI: `10.1088/1361-6382/abf441`.

[46]   B Sathyaprakash et al. "Scientific objectives of Einstein Telescope". In: *Classical and Quantum Gravity* 29.12 (June 2012), p. 124013. DOI: `10.1088/0264-9381/29/12/124013`.

[47]   Katharina-Sophie Isleif et al. "Experimental demonstration of deep frequency modulation interferometry". In: *Optics Express* 24.2 (Jan. 2016), p. 1676. DOI: `10.1364/oe.24.001676`.

[48]   C M Mow-Lowry and D Martynov. "A 6D interferometric inertial isolation system". In: *Classical and Quantum Gravity* 36.24 (Nov. 2019), p. 245006. DOI: `10.1088/1361-6382/ab4e01`.

[49]   F. Matichard et al. "Modeling and experiment of the suspended seismometer concept for attenuating the contribution of tilt motion in horizontal measurements". In: *Review of Scientific Instruments* 87.6 (June 2016), p. 065002. DOI: `10.1063/1.4953110`.

[50]   Amit Singh Ubhi et al. "Active platform stabilization with a 6D seismometer". In: *Applied Physics Letters* 121.17 (Oct. 2022), p. 174101. DOI: `10.1063/5.0118606`.

[51]   Nanometrics Seismoligcal Services. *NMX Trillium 240 Broadbrand Spectromoter*. Ed. by Nanometrics Seismoligcal Services.

[52] LLC Geotech Instruments. *Short-period Seismometer*. Ed. by LLC Geotech Instruments. `http://www.ifg.tu-clausthal.de/java/seis/sdem_app-e.html`. 2017.

[53] Robin Kirchhoff. *Sensor noise of L-22D and L-4C Geophones*. URL: `%5Curl%7Bhttps://dcc.ligo.org/public/0138/T1600438/001/L-4C%20huddle%20test%20at%20the%20AEI.pdf%7D`.

[54] N. Yazdi, F. Ayazi, and K. Najafi. "Micromachined inertial sensors". In: 86.8 (1998), pp. 1640–1659. DOI: `10.1109/5.704269`.

[55] Grantham Pang and Hugh Liu. "Evaluation of a Low-cost MEMS Accelerometer for Distance Measurement". In: 30.3 (2001), pp. 249–265. DOI: `10.1023/a:1008113324758`.

[56] Xudong Zou, Pradyumna Thiruvenkatanathan, and Ashwin A. Seshia. "A Seismic-Grade Resonant MEMS Accelerometer". In: 23.4 (Aug. 2014), pp. 768–770. DOI: `10.1109/jmems.2014.2319196`.

[57] Xiaoji Niu et al. "A universal approach for processing any MEMS inertial sensor configuration for land-vehicle navigation". In: *The Journal of Navigation* 60.2 (2007), pp. 233–245.

[58] Alexander A Trusov et al. "Silicon accelerometer with differential frequency modulation and continuous self-calibration". In: *2013 IEEE 26th International Conference on Micro Electro Mechanical Systems (MEMS)*. IEEE. 2013, pp. 29–32.

[59] Sergei A. Zotov et al. "High Quality Factor Resonant MEMS Accelerometer With Continuous Thermal Compensation". In: *IEEE Sensors Journal* 15.9 (2015), pp. 5045–5052. DOI: `10.1109/JSEN.2015.2432021`.

[60] Felipe Guzmán Cervantes et al. "High sensitivity optomechanical reference accelerometer over 10 kHz". In: *Applied Physics Letters* 104.22 (June 2014), p. 221111. DOI: `10.1063/1.4881936`.

[61] O Gerberding et al. "Optomechanical reference accelerometer". In: *Metrologia* 52.5 (Sept. 2015), pp. 654–665. DOI: `10.1088/0026-1394/52/5/654`.

[62] Kenji Numata et al. "Intrinsic losses in various kinds of fused silica". In: *Classical and Quantum Gravity* 19.7 (Mar. 2002), pp. 1697–1702. DOI: `10.1088/0264-9381/19/7/363`.

[63] Y. Bellouard et al. "Femtoprint: A femtosecond laser printer for micro- and nano-scale systems". In: *Conference on Lasers and Electro-Optics 2012*. OSA, 2012. DOI: `10.1364/cleo_at.2012.atu3l.3`.

[64] Felipe Guzmán Cervantes et al. "High sensitivity optomechanical reference accelerometer over 10 kHz". In: *Applied Physics Letters* 104.22 (June 2014), p. 221111. DOI: `10.1063/1.4881936`.

[65] B Merchant and G Slad. "Next Generation Qualification: Nanometrics T120PH Seismometer Evaluation". In: *Sandia Report SABD2917011186* (2017).

[66] Michael P Ross et al. "Towards windproofing LIGO: reducing the effect of wind-driven floor tilt by using rotation sensors in active seismic isolation". In: *Classical and Quantum Gravity* 37.18 (Aug. 2020), p. 185018. DOI: `10.1088/1361-6382/ab9d5c`.

[67] D. C. Champeney. *A handbook of Fourier theorems*. Cambridge University Press, 1987.

[68] Herbert B. Callen and Theodore A. Welton. "Irreversibility and Generalized Noise". In: *Physical Review* 83.1 (June 1951), pp. 34–40. DOI: `10.1103/physrev.83.34`.

[69] J. Weber. "Fluctuation Dissipation Theorem". In: *Physical Review* 101.6 (Mar. 1956), pp. 1620–1626. DOI: `10.1103/physrev.101.1620`.

[70] Kazuhiro Agatsuma. "Study of pendulum thermal noise in gravitational wave detectors". PhD thesis. Department of Physics, Faculty of Science, University of Tokyo, Apr. 2010.

[71] Peter R. Saulson. "Thermal noise in mechanical experiments". In: *Physical Review D* 42.8 (Oct. 1990), pp. 2437–2445. DOI: `10.1103/physrevd.42.2437`.

[72] Peter R Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, 1994.

[73] Clarence M. Zener and Sidney Siegel. "Elasticity and Anelasticity of Metals." In: *The Journal of Physical and Colloid Chemistry* 53.9 (Sept. 1949), pp. 1468–1468. DOI: `10.1021/j150474a017`.

[74] Ron Lifshitz and Michael L Roukes. "Thermoelastic damping in micro-and nanomechanical systems". In: *Physical review B* 61.8 (2000), p. 5600.

[75] A. N. Norris and D. M. Photiadis. "Thermoelastic relaxation in elastic structures, with applications to thin plates". In: *The Quarterly Journal of Mechanics and Applied Mathematics* 58.1 (Feb. 2005), pp. 143–163. DOI: `10.1093/qjmamj/hbi002`.

[76] P. Chadwick. "Thermal damping of a vibrating elastic body". In: *Mathematika* 9.1 (June 1962), pp. 38–48. DOI: `10.1112/s0025579300003077`.

[77] Amy Duwel et al. "Engineering MEMS Resonators With Low Thermoelastic Damping". In: *Journal of Microelectromechanical Systems* 15.6 (Dec. 2006), pp. 1437–1445. DOI: `10.1109/jmems.2006.883573`.

[78] X. Guo, Y.-B. Yi, and S. Pourkamali. "A finite element analysis of thermoelastic damping in vented MEMS beam resonators". In: *International Journal of Mechanical Sciences* 74 (July 2013), pp. 73–82. DOI: `10.1016/j.ijmecsci.2013.04.013`.

[79] *Thermoelastic Damping in a MEMS Resonator*. `https://www.comsol.com/model/thermoelastic-damping-in-a-mems-resonator-1439`. Accessed: 6/1/2021.

[80] Andri M. Gretarsson and Gregory M. Harry. "Dissipation of mechanical energy in fused silica fibers". In: *Review of Scientific Instruments* 70.10 (Oct. 1999), pp. 4081–4087. DOI: `10.1063/1.1150040`.

[81] Zhili Hao, Ahmet Erbil, and Farrokh Ayazi. "An analytical model for support loss in micromachined beam resonators with in-plane flexural vibrations". In: *Sensors and Actuators A: Physical* 109.1-2 (Dec. 2003), pp. 156–164. DOI: `10.1016/j.sna.2003.09.037`.

[82] Jin Hwan Ko et al. "Quality factor in clamping loss of nanocantilever resonators". In: *Applied Physics Letters* 98.17 (Apr. 2011), p. 171909. DOI: `10.1063/1.3575560`.

[83] G. Cagnoli et al. "Suspension losses in low-frequency mechanical pendulums". In: *Physics Letters A* 213.5-6 (Apr. 1996), pp. 245–252. DOI: `10.1016/0375-9601(96)00146-6`.

[84] I. Wilson-Rae. "Intrinsic dissipation in nanomechanical resonators due to phonon tunneling". In: *Physical Review B* 77.24 (June 2008), p. 245418. DOI: `10.1103/physrevb.77.245418`.

[85] I. Wilson-Rae et al. "High-Q Nanomechanics via Destructive Interference of Elastic Waves". In: *Physical Review Letters* 106.4 (Jan. 2011), p. 047205. DOI: `10.1103/physrevlett.106.047205`.

[86] S J Cooper et al. "A compact, large-range interferometer for precision measurement and inertial sensing". In: 35.9 (Mar. 2018), p. 095007. DOI: `10.1088/1361-6382/aab2e9`. URL: `https://doi.org/10.1088/1361-6382/aab2e9`.

[87] Feng Zhou et al. "Broadband thermomechanically limited sensing with an optomechanical accelerometer". In: *Optica* 8.3 (Mar. 2021), pp. 350–356. DOI: `10.1364/OPTICA.413117`. URL: `http://www.osapublishing.org/optica/abstract.cfm?URI=optica-8-3-350`.

[88] Alexander G Krause et al. "A high-resolution microchip optomechanical accelerometer". In: *Nature Photonics* 6.11 (2012), pp. 768–772.

[89] R. W. P. Drever et al. "Laser phase and frequency stabilization using an optical resonator". In: *Applied Physics B Photophysics and Laser Chemistry* 31.2 (June 1983), pp. 97–105. DOI: `10.1007/bf00702605`.

[90] Oliver Gerberding. "Deep frequency modulation interferometry". In: *Optics Express* 23.11 (May 2015), p. 14753. DOI: `10.1364/oe.23.014753`.

[91] Katharina-Sophie Isleif et al. "Compact Multifringe Interferometry with Subpicometer Precision". In: *Physical Review Applied* 12.3 (Sept. 2019), p. 034025. DOI: `10.1103/physrevapplied.12.034025`.

[92] Jiri Smetana et al. "Compact Michelson Interferometers with Subpicometer Sensitivity". In: *Physical Review Applied* 18.3 (Sept. 2022), p. 034040. DOI: `10.1103/physrevapplied.18.034040`.

[93] Yichao Yang et al. "Single-Element Dual-Interferometer for Precision Inertial Sensing". In: *Sensors* 20.17 (Sept. 2020), p. 4986. DOI: `10.3390/s20174986`.

[94] D. T. Smith, J. R. Pratt, and L. P. Howard. "A fiber-optic interferometer with subpicometer resolution for dc and low-frequency displacement measurement". In: *Review of Scientific Instruments* 80.3 (Mar. 2009), p. 035105. DOI: `10.1063/1.3097187`.

[95]   Yanqi Zhang and Felipe Guzman. "Fiber-based two-wavelength heterodyne laser interferometer". In: *Optics Express* 30.21 (Sept. 2022), p. 37993. DOI: 10.1364/oe.466332.

[96]   Kenji Numata et al. "Measurement of the intrinsic mechanical loss of low-loss samples using a nodal support". In: *Physics Letters A* 276.1-4 (Oct. 2000), pp. 37–46. DOI: 10.1016/s0375-9601(00)00646-0.

[97]   Jae Yoong Cho et al. "Fused-Silica Micro Birdbath Resonator Gyroscope". In: *Journal of Microelectromechanical Systems* 23.1 (Feb. 2014), pp. 66–77. DOI: 10.1109/jmems.2013.2291534.

[98]   Doruk Senkal et al. "MEMS Micro-glassblowing Paradigm for Wafer-level Fabrication of Fused Silica Wineglass Gyroscopes". In: *Procedia Engineering* 87 (2014), pp. 1489–1492. DOI: 10.1016/j.proeng.2014.11.580.

[99]   Wan-Thai Hsu. "Recent progress in silicon MEMS oscillators". In: *Proceedings of the 40th Annual Precise Time and Time Interval Systems and Applications Meeting*. 2008, pp. 135–146.

[100]   Adam Hines et al. "Optomechanical inertial sensors". In: *Applied Optics* 59.22 (June 2020), G167. DOI: 10.1364/ao.393061.

[101]   A. Hines et al. "Compact optomechanical accelerometers for use in gravitational wave detectors". In: *Applied Physics Letters* 122.9 (Feb. 2023), p. 094101. DOI: 10.1063/5.0142108.

[102]   Andrea Mary-Helen Nelson and Felipe Guzman. "Compact optomechanical inertial sensors with fused silicia and si-based resonators". In: *Optical and Quantum Sensing and Precision Metrology II*. Ed. by Selim M. Shahriar and Jacob Scheuer. SPIE, Mar. 2022. DOI: 10.1117/12.2609952.

[103]   Sanjay Kumar et al. "Additive manufacturing as an emerging technology for fabrication of microelectromechanical systems (MEMS)". In: *Journal of Micromanufacturing* 2.2 (June 2019), pp. 175–197. DOI: 10.1177/2516598419843688.

[104]   E. Wielandt and G. Streckeisen. "The leaf-spring seismometer: Design and performance". In: *Bulletin of the Seismological Society of America* 72.6A (Dec. 1982), pp. 2349–2367. ISSN: 0037-1106. DOI: 10.1785/BSSA07206A2349. eprint: https://pubs.geoscienceworld.org/ssa/bssa/article-pdf/72/6A/2349/5330641/bssa07206a2349.pdf. URL: https://doi.org/10.1785/BSSA07206A2349.

[105]   J. N. Reddy. *Theory and analysis of elastic plates and shells*. Boca Raton, Fla. London: CRC Taylor & Francis distributor, 2006. ISBN: 9780849384158.

[106]   R J. Melosh. "A Stiffness Matrix for the Analysis of Thin Plates in Bending". In: *Journal of the Aerospace Sciences* 28.1 (Jan. 1961), pp. 34–42. DOI: 10.2514/8.8850.

[107]   J.V. van Heijningen et al. "A cryogenic inertial sensor for terrestrial and lunar gravitational-wave detection". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (July 2022), p. 167231. DOI: 10.1016/j.nima.2022.167231.

[108] A. Bertolini et al. "Monolithic folded pendulum accelerometers for seismic monitoring and active isolation systems". In: *IEEE Transactions on Geoscience and Remote Sensing* 44.2 (Feb. 2006), pp. 273–276. DOI: `10.1109/tgrs.2005.861006`.

[109] Jan Harms and Conor M Mow-Lowry. "Suspension-thermal noise in spring–antispring systems for future gravitational-wave detectors". In: *Classical and Quantum Gravity* 35.2 (Dec. 2017), p. 025008. DOI: `10.1088/1361-6382/aa9e28`.

[110] W.J. Alford, R.D. VanderNeut, and V.J. Zaleckas. "Laser scanning microscopy". In: *Proceedings of the IEEE* 70.6 (1982), pp. 641–651. DOI: `10.1109/proc.1982.12362`.

[111] Stephen W. Paddock. "Confocal Laser Scanning Microscopy". In: *BioTechniques* 27.5 (Jan. 1999), pp. 992–1004. DOI: `10.2144/99275ov01`.

[112] Jonathan J Carter et al. "An Experiment to Test the Mechanical Losses of Different Bonding Techniques in Fused Silica". In: *arXiv preprint arXiv:2201.09909* (2022). URL: `https://arxiv.org/search/?query=carter+mechanical+loss&searchtype=all&source=header`.

[113] Marina Trad Nery. "Laser power stabilization via radiation pressure". en. In: (2021). DOI: `10.15488/11012`.

[114] L Giacoppo et al. "Towards ponderomotive squeezing with SIPS experiment". In: *Physica Scripta* 96.11 (July 2021), p. 114007. DOI: `10.1088/1402-4896/ac1637`.

[115] Anja Schroeter et al. *On the mechanical quality factors of cryogenic test masses from fused silica and crystalline quartz*. 2007. DOI: `10.48550/ARXIV.0709.4359`.

[116] Marie-Sophie Hartig, Sönke Schuster, and Gudrun Wanner. "Geometric tilt-to-length coupling in precision interferometry: mechanisms and analytical descriptions". In: *Journal of Optics* 24.6 (May 2022), p. 065601. DOI: `10.1088/2040-8986/ac675e`.

[117] Zhongping Wang and Zengming Zhang. "Electron Beam Evaporation Deposition". In: *Advanced Nano Deposition Methods*. Wiley-VCH Verlag GmbH & Co. KGaA, Sept. 2016, pp. 33–58. DOI: `10.1002/9783527696406.ch2`.

[118] Pascal Birckigt et al. "Plasma-activated direct bonding of coated optical glasses". In: *Japanese Journal of Applied Physics* 59.SB (2019), SBBD01.

[119] F. M. Phelps. "Airy Points of a Meter Bar". In: *American Journal of Physics* 34.5 (Sept. 1966), pp. 419–422. DOI: `10.1119/1.1973011`.

[120] Miguel Dovale Álvarez. *Optical Cavities for Optical Atomic Clocks, Atom Interferometry and Gravitational-Wave Detection*. Springer International Publishing, 2019. DOI: `10.1007/978-3-030-20863-9`.

[121] John Melcher et al. "A self-calibrating optomechanical force sensor with femtonewton resolution". In: *Applied Physics Letters* 105.23 (Dec. 2014), p. 233109. DOI: `10.1063/1.4903801`.

[122] Oliver Gerberding et al. "Laser-Frequency Stabilization via a Quasimonolithic Mach-Zehnder Interferometer with Arms of Unequal Length and Balanced dc Readout". In: *Physical Review Applied* 7.2 (Feb. 2017), p. 024027. DOI: `10.1103/physrevapplied.7.024027`.

[123] W. G. Schweitzer et al. "Description, Performance, and Wavelengths of Iodine Stabilized Lasers". In: *Applied Optics* 12.12 (Dec. 1973), p. 2927. DOI: `10.1364/ao.12.002927`.

[124] Eric D. Black. "An introduction to Pound–Drever–Hall laser frequency stabilization". In: *American Journal of Physics* 69.1 (Jan. 2001), pp. 79–87. DOI: `10.1119/1.1286663`.

[125] R. Kirchhoff et al. "Huddle test measurement of a near Johnson noise limited geophone". In: *Review of Scientific Instruments* 88.11 (Nov. 2017), p. 115008. DOI: `10.1063/1.5000592`.

[126] S Goßler et al. "The AEI 10 m prototype interferometer". In: *Classical and Quantum Gravity* 27.8 (Apr. 2010), p. 084023. DOI: `10.1088/0264-9381/27/8/084023`.

[127] Karsten Danzmann and the LISA study team. "LISA: laser interferometer space antenna for gravitational wave measurements". In: *Classical and Quantum Gravity* 13.11A (Nov. 1996), A247–A250. DOI: `10.1088/0264-9381/13/11a/033`.

[128] Frank Flechtner et al. "Status of the GRACE Follow-On Mission". In: *Gravity, Geoid and Height Systems*. Springer International Publishing, 2014, pp. 117–121. DOI: `10.1007/978-3-319-10837-7_15`.

[129] Lee A. Capistran et al. "Wide-band dual optomechanical resonator for inertial navigation". In: *Quantum Sensing, Imaging, and Precision Metrology*. Ed. by Selim M. Shahriar and Jacob Scheuer. SPIE, Mar. 2023. DOI: `10.1117/12.2655504`.

[130] Jonathan Carter et al. "A High Q, Quasi-Monolithic Optomechanical Inertial Sensor". In: *2020 IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*. IEEE, Mar. 2020. DOI: `10.1109/inertial48129.2020.9090085`.

[131] P. Chadwick. "On the propagation of thermoelastic disturbances in thin plates and rods". In: *Journal of the Mechanics and Physics of Solids* 10.2 (Apr. 1962), pp. 99–109. DOI: `10.1016/0022-5096(62)90013-3`.

[132] H.W. Lord and Y. Shulman. "A generalized dynamical theory of thermoelasticity". In: *Journal of the Mechanics and Physics of Solids* 15.5 (Sept. 1967), pp. 299–309. DOI: `10.1016/0022-5096(67)90024-5`.