Breaking the Seismic Wall: How to Improve Gravitational Wave Detectors at Low Frequency

Samuel James Cooper

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

DOCTOR OF PHILOSOPHY

Institute for Gravitational Wave Astronomy

School of Physics and Astronomy

College of Engineering and Physical Sciences

University of Birmingham

July 2019

Abstract

The era of gravitational-wave astronomy was enabled by the incredible sensitivity of the LIGO and VIRGO detectors. However, they are still plagued by technical noises at frequencies below 30 Hz [1, 2], driven in part by the limitations of the seismic isolation of the detector [3, 4]. To detect gravitational waves at low frequency, the isolation performance must be improved to reduce these technical noises.

To improve the performance of seismic isolation systems, I have developed HoQI a new interferometrically sensor [5], that can be applied to both the isolation tables and suspensions. HoQI has a resolution a factor 1000 higher than sensors currently used in LIGO and I have quantified the level of non-linearity present in the sensor and shown this to not being a limiting factor.

HoQIs impact on the performance of the seismic isolation system have also been quantified, through the use of a accurate model of an Advanced LIGO isolation platform that I have developed. Using the model I have shown that using using HoQI the expected isolation platform motion can be reduced by a factor of 70 at 0.1 Hz and a factor of 10 at 2 Hz. I have shown that the control filters used in this model can be improved by up to 70% by designing them using particle swarm optimisation.

Acknowledgments

First, I must thank my supervisor Conor for guiding me through the duration of my PhD and for giving me confidence in my work, giving me the chance to travel to conferences and to work as a fellow at Hanford. Thank you for organising Pub Club and many outings to the cinema and for making the group a fun place to work. Thank you to Andreas, my secondary supervisor for encouraging me to get involved with outreach and public engagement projects, they've not only maintained my passion for science but have also made me a better scientist. Thank you for teaching me about interferometry and for encouraging me to apply for this PhD, it has been one of the best decisions I've made in my life.

Thanks to Alberto for creating and maintaining a fantastic group of people which continues to be a pleasure to be a part of. Thanks to my examiners Jon and Giles for a enjoyable viva and constructive comments and discussions that have improved the overall quality of my thesis.

Thanks to Chris for your help in the lab and with complicated mathematics. To Dave and John for your help fixing electronics, writing new software and for help in debugging experiments. Thank you all for answering my many questions and requests for editing various lab software and electronics. Thank you to David for fixing many computer problems and for your company at morning coffee.

Thanks go to everyone in the ifolab group for your input on my projects, for useful conversations and advice and for your friendship throughout, Anna, Hayou, Dan, Miguel, Daniel, Aaron, Phil, Sam, Joe, Chiara, Riccardo, Amit, George, Chris, Artemiy, and Leo. To all my friends in the ASR group thank you for making the past four years so enjoyable. Thanks to Janna, Callum and Jake for many distractions and for helping me not take life too seriously.

A special mention needs to go to Anna, Hannah and those who worked on the Thinktank museum exhibit and to Aaron and Conner for adapting it into the Royal Society project, you all made it a fantastic experience to work on. Thanks to Aaron and Sam for helping make Chirp into something we can be proud of, I could not have done it without you.

Thanks to Josh and Sam for your friendship through all these years, for being there to talk to, for reminding me that there's other things in life than work. To Iain, Dan, Oliver, Will, Max, and many others, thank you for the games and distractions.

To my family, thank you for supporting me all my life, for your unwaivering love and guidance, thank you for helping me realise a childhood dream. Thanks to Cat, for being a fantastic sister, I'm so proud of you. Lastly to my parents, thank you for taking me to the local observatory, for showing me the stars and for helping me find my love of physics. Thank you for never doubting me, for your continuous belief in me, for pushing me to achieve my potential and for never putting limits on what I could do. This is for you.

Contents

List of Figures

List of Tables

Glossary

ADC Analog-to-Digital Converter.

AOM Acousto Optic Modulator.

ASD Amplitude Spectral Density.

BBH Binary Black Hole.

BLRMS Band Limited Root Mean Square.

BNS Binary Neutron Star.

BOSEM A type of compact sensor used to actuate on stages of the LIGO suspension. Birmingham Optical Sensor and Electro-Magnetic actuator.

BRS A Beam Rotation Sensor used to measure ground rotation.

BS Beamsplitter.

BSC One of the two types of vacuum chambers in LIGO, Beamsplitter Chamber (BSC).

CAD Computer Aided Design.

CDS Control and Data System.

CPS Capacitive Position Sensors.

DARM Differential Arm Length.

DCC LIGOs internal document repository, the Document Control Center first.

ESD Electro Static Drive.

ETM End Test Mass.

EUCLID A compact interferometer developed at the University of Birmingham, Easy to use compact laser interferometric device.

FF Feed Forward.

FFT Fast Fourier Transform.

FPGA Field Programmable Gate Array.

GCN Gamma-ray Coordination Network.

HAM One of the two types of vacuum chambers in LIGO, Horizontal Access Module (HAM).

HEPI Hydraulic External Pre-Isolator.

HoQI The Homodyne Quadrature Interferometer, described in chapter 2.

HP Abbreviation for High Pass.

HWP Half Waveplate.

ILIAD A compact rotation measuring interferometer, innovative laser interferometric angular device.

ISI Internal Seismic Isolation.

ITM Input Test Mass.

LHO The LIGO Hanford Observatory.

LIGO The name of the USA based gravitational wave detector sites..

LLO The LIGO Livingston Observatory.

LP Abbreviation for Low Pass.

MCSS2 Multi-Channel Coherent Subtraction.

NDS Network Data System servers used to access LIGO data..

NPBS Non-Polarising Beamsplitter.

O1 First Observing Run.

O2 Second Observing Run.

O3 Third Observing Run.

PBS Polarising Beamsplitter.

PD Photodiode.

PER Polarisation Extinction Ratio.

PM Polarisation Maintaining.

PRC Power Recycling Cavity.

PRM Power Recycling Mirror.

PSD Power Spectral Density.

QWP Quarter Waveplate.

RMS Root Mean Square.

SC Sensor Correction - a type of feedforward technique.

SEI Abbreviation for seismic.

SRC Signal Recycling Cavity.

SRCL Length degree of freedom of the Signal Recycling Cavity.

SRM Signal Recycling Mirror.

ST0 The zeroth (ground) stage of the ISI.

ST1 The first stage of the ISI.

SUS Abbreviation for suspension.

SVN A type of version control named Subversion.

UGF The frequency at which the gain of a controller is equal to 1.

UIM Upper Intermediate Mass.

VCSEL Vertical-cavity surface-emitting laser.

ZPK Zero-pole-gain system used in controls..

Chapter 1

Introduction

All things are difficult before they are easy.

Dr. Thomas Fuller

1.1 Gravitational Waves

One hundred years after their prediction by Albert Einstein in 1916, gravitational waves were observed for the first time by Laser Interferometer Gravitational-Wave Observatory (LIGO) on the 14th September 2015 [14] as a result of a Binary Black Hole (BBH) collision named GW150914. These colliding black holes had masses of around 36 and 29 solar masses and merged together 410^{+160}_{-180}Mpc from the Earth. This first merger provided tests of General Relativity in the strong field regime which found no evidence of deviations from the predictions made by General Relativity [15]. Moreover the first detection provided an estimate of the likely rate of binary black hole mergers [16] and demonstrated that black holes with a mass of tens of solar masses can form in nature [17]. During the first observing run (O1) a two further BBH events were detected. The second observing run (O2) resulted in a seven more BBH events and the first ever Binary Neutron Star (BNS) event [18, 19], signaling the era of multi-messenger astronomy, bringing the total number of confirmed events to 11 [20]. The BNS event is significant as it was observed not only

using gravitational waves but also in every part of the electromagnetic spectrum [19]. Due to the multi-messenger nature of the detection, GW170817 served as an independent measure of Hubble's constant [21] and proved that gravitational waves travel at the same speed as light [22]. At the time of writing, 15 event candidates have been detected as part of the third observing run (O3) as well as the previous 11 confirmed detections. With the start of O3, low-latency public notifications of potential gravitational wave events [23, 24] using the NASA Gamma-ray Coordination Network (GCN), containing information such as sky localisation and source origin probabilities, have been released.

Gravitational waves are a consequence of the Einstein's theory of General Relativity and are small perturbations, commonly referred to as 'stretching and squashing' of space-time. These waves are incredibly difficult to detect, the "loudest" event observed so far had a peak strain on the order of 10^{-18}, and Einstein himself doubted that these effects could ever be detected [25, 26]. The early gravitational wave detectors in the 1960s used Weber bars [27], pioneered by Joseph Weber as the mechanism to detect gravitational waves; this lead to a detection claim originally published nine years later [28]. In 1975 Hulse and Taylor demonstrated evidence of the existence of gravitational waves [29] by the increase in rotational speed of a binary pulsar system. Around this time, two independent groups at MIT and Caltech started work on prototype interferometric gravitational wave detectors, measuring 1.5 and 40 m in length [30, 31] respectively. These projects were combined and funded in the late 1980s to form LIGO. A nice summary on the history of gravitational waves and their detection is written by Cervantes-Cota et. al. [31]. A global network of detectors is now operational, including both LIGO detectors, the LIGO Hanford Observatory (LHO) and the LIGO Livingston Observatory (LLO) as well as the VIRGO detector in Italy, while the Japanese detector KAGRA is expected to become operational at the end of O3. The GEO-HF detector is a 600 m long detector located in Sarstedt, Germany. While it's sensitivity is comparable to first-generation detectors, it is still operational and remains able to detect sources located in our galaxy [32].

1.2 Detection Method

LIGO is based on a design developed by Albert Michelson and Edward Morley [33] called a Michelson interferometer. The LIGO detector features 11 separate vacuum chambers with the end mirrors separated from the corner station beamsplitter by a 4 km long vacuum tube, necessary to reduce the effect of air on the laser beam, shown in FIG 1.1. Inside each of these vacuum chambers is an Internal Seismic Isolation (ISI) and multiple suspensions systems to control and attenuate the motion of the ground from overwhelming the gravitational wave signal. In a basic Michelson interferometer, the input laser beam is split by a suspended Beamsplitter (BS) into two beams that travel perpendicularly along the 4 km long arms of the interferometer. The light is then reflected from the end mirrors where it recombines an interferes at the central beamsplitter, this interference signal is readout by a photodiode.

LIGO adapts the base design presented by Michelson and Morley to increase the sensitivity of the detector. Cavities have been installed in both arms of the interferometer to circulate the light in the arms, to increase the effective length of the detector, increasing the gravitational wave signal. Two additional cavities have been installed at the input and output of the detector, these are the Power Recycling Cavity (PRC) and Signal Recycling Cavity (SRC) respectively. The purpose of the PRC is to further increase the effective power at the beamsplitter, reducing the relative shot noise. Since LIGO operates near the dark fringe, the majority of the light is reflected from the beamsplitter towards the laser. The Power Recycling Mirror (PRM) reflects the light back into the interferometer, increasing the power on the beamsplitter and improving the shot-noise limited sensitivity. The SRC can be used in two configurations, one to enhance the sensitivity in a narrow frequency region [34] and one to maintain a broadband frequency response by effectively lowering the arm cavity finesse.

For control purposes, the detector is held close to the dark operating point of the interference fringe. This requires the motion of each of the mirrors in the detector to be stabilised to within, typically, picometers of their operating point to remain locked. To ensure the detector remains locked and

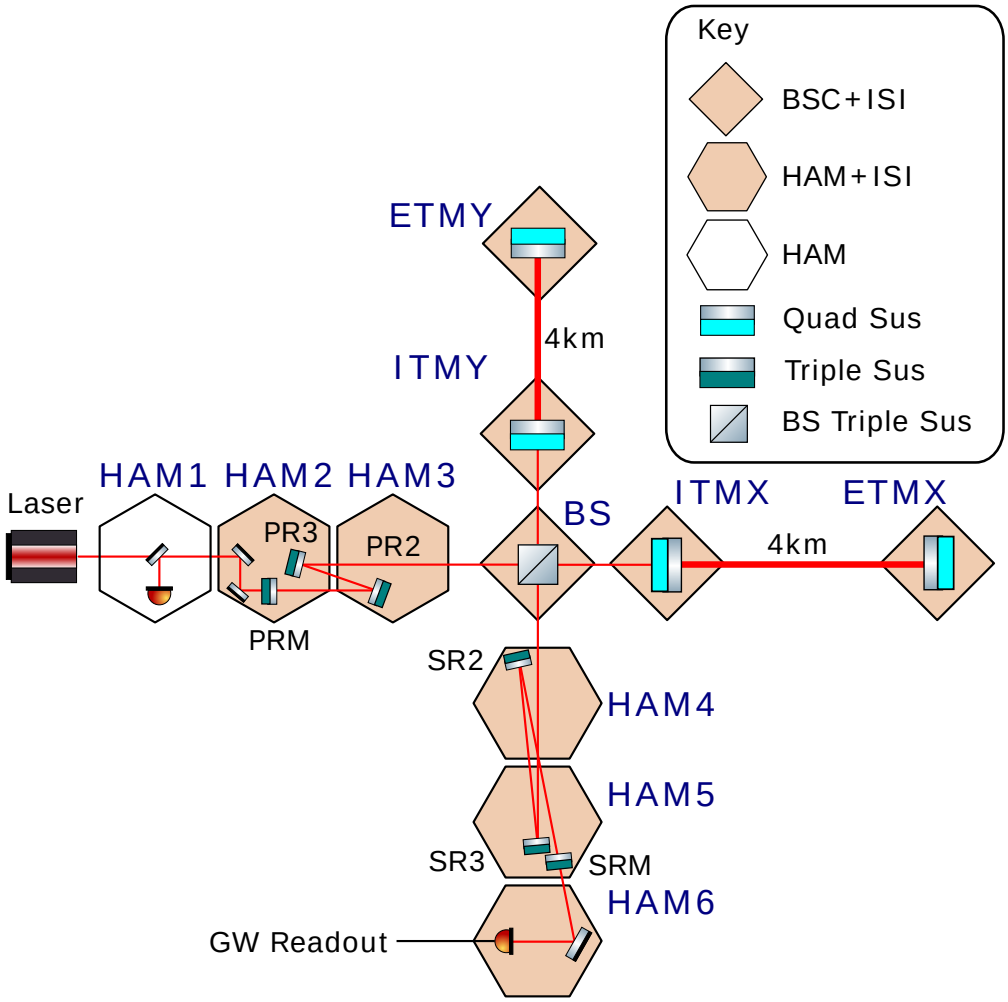

Figure 1.1: A simplified schematic of LIGO, showing key optics and the types of isolation used in the detector. Each of the Horizontal Access Module (HAM) chambers feature an single stage ISI, with the exception of HAM1. The Input Test Mass (ITM), End Test Mass (ETM) and BS are Beamsplitter Chamber (BSC) ISIs and have two stages of passive isolation.

operational, various noise sources must be controlled and suppressed. This was one of the largest technical issues that needed to be overcome to detect gravitational waves. The next subsections will discuss some of the primary noise sources: seismic noise, thermal noise and quantum noise.

1.2.1 Noise Sources

The detector is limited by various noise sources throughout the gravitational wave detection band, which spans from 10 Hz to 5 kHz. The noise budget of advanced LIGO is shown in FIG 1.2. This is a typical plot used to identify limiting noise sources in the detector and is compared against ground motion as measured at one of the detector sites highlighting the scale of ground motion compared to gravitational wave signals. The sensitivity of advanced LIGO is shown by the dashed black curve in FIG 1.2. The shape of the curve describes how the sensitivity of the detector changes with frequency and is sometimes referred to as a 'bucket'. The 'arms' of the detector are 4 km long and the Fabry–Pérot arm cavities have a finesse of approximately 300, making the round trip effective arm length 2400 km. The gravitational wave will have the largest effect on the differential phase when the wavelength of the incoming wave is equal to the effective arm length. As gravitational waves travel at the speed of light, the optimum frequency can be calculated by the equation,

$$f_{\mathrm{opt}} = \frac{c}{\lambda_{\mathrm{GW}}}, \tag{1.1}$$

where λ_{GW} is the wavelength of the gravitational wave and c is the speed of light. Substituting λ_{GW} as 2400 km, we find the frequency for the peak sensitivity, or the lowest point of the 'bucket' is approximately 100 Hz, the bucket itself spans from 60 to 200 Hz.

At high frequencies above 100 Hz the detector is limited by quantum noise, often referred to as shot noise and can be thought of as the uncertainty in the number of detected photons [35]. Techniques to reduce shot noise vary from increasing laser power, thereby reducing the uncertainty in the number of photons to injecting squeezed light to reduce the zero point fluctuations [36].

In the 10 to 100 Hz band the design sensitivity is limited by a mixture of the previously mentioned quantum noise and coating Brownian noise; caused by the Brownian motion of the surface coating of the test masses, as described in [37]. Around 10 Hz the design curve becomes limited by suspension thermal noise, caused by the random motion of atoms in the suspension itself [38, 39]. Both these

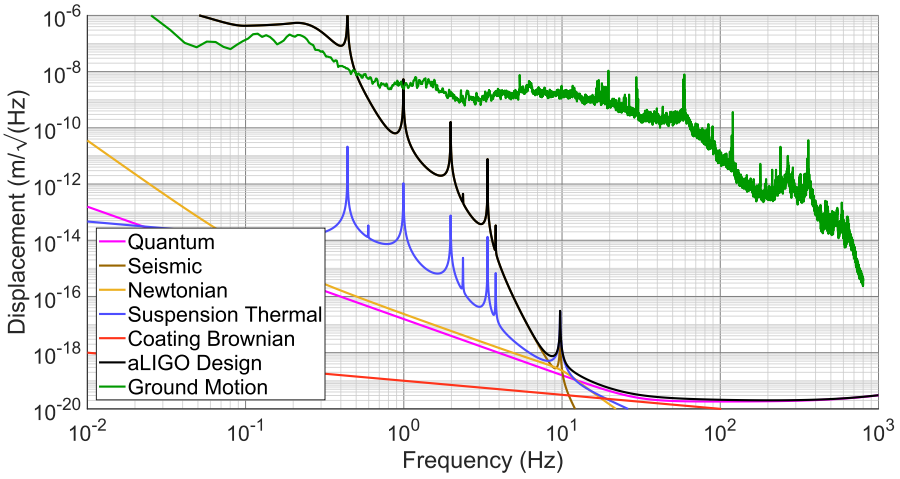

Figure 1.2: A comparison between seismic noise measured on site at LIGO Hanford and the design noise budget of Advanced LIGO, and is calculated using GWINC [6], with the total noise estimate of aLIGO shown in black. Note that the seismic noise is over ten orders of magnitude higher than the design sensitivity at the start of the detection band.

.

effects can be reduced by cryogenically cooling the mirrors, which is something being actively developed for 3rd generation detectors [40]. Improvements in reducing the thermal noise at room temperature are also being explored using materials with ultralow mechanical loss [41].

Newtonian noise affects the low frequency part of LIGO's spectrum, below 20 Hz. This is caused by changes in the local gravitational field by changes in the local density [42]. These density changes can be estimated by sensor arrays, comprised of accelerometers and geophones, modeled and subtracted acausally from the gravitational wave readout [43]. At the time of writing, Newtonian noise is not a limiting noise source [44].

Below 10 Hz the design sensitivity is limited by seismic noise, caused by residual ground motion that couples through the ISI and is transmitted through the suspensions. This motion is attenuated by a factor of 10^3 at the start of the detection band [45, 46] by the ISIs and is further attenuated by the suspensions. The suspensions have resonance frequencies between 0.5 and 4 Hz and attenuates the residual platform motion proportionally to $\frac{1}{f^8}$ above all the resonance frequencies, giving rise

to a seismic 'wall'. This results in the residual seismic motion only being limiting at frequencies below 10 Hz. As noted by [47] this residual seismic motion must be reduced in order to detect gravitational waves at 10 Hz and below.

Technical noise, or control noise, is noise in the detector that is imposed in order to maintain the operation of the interferometer. This noise source is a result of the various control loops that must be activated to keep the detector within its operational window. As such, technical noise does not appear on any design curve, however in practice it remains a limiting noise source below approximately 30 Hz and must be reduced to detect gravitational waves in this frequency band [48, 49]. The coupling pathway to how this technical noise couples into the Differential ARM (DARM) length, otherwise known as the gravitational wave readout, is complicated and there is likely many different coupling paths relating to how each auxiliary degree of freedom is causing an increase in noise. However, this technical noise, in the most part, can be traced back to residual ISI motion, an example of this is shown in chapter 5. As such, this thesis focuses on developing sensors and improved control systems, with the aim of reducing this residual motion of the ISIs and suspensions.

1.3 Thesis Overview

Chapter 2 details the design, development and testing of HOmodyne Quadrature Interferometer (HoQI) a compact, long range, high resolution interferometer for use as a standalone sensor, or as a readout mechanism for inertial sensors and is published in Classical and Quantum Gravity [5]. Such sensors can be used to measure the motion of the suspended masses on the suspensions of LIGO.

Chapter 3 quantifies how imperfections in the optics or alignment of homodyne phasemeters leads to the reduction of the sensors sensitivity and describes methods of how to reduce the amplitude of these effects. The chapter also details simulation work undertaken to quantify the effect that non-linearities will have on applications of HoQI, such as a suspension sensor or as a readout

mechanism for a high Q seismometer.

Chapter 4 details the development and testing of an optical inertial sensor that uses HoQI to increase the resolution of a L-4C geophone by a factor of 60 at 10 mHz. Such sensors can then be used to reduce the residual motion of ISI tables.

Chapter 5 is based on a technical note describing the development of a predictive model of a HAM ISI that is used to evaluate new sensors described in this thesis. The model can be used to compare the ISI performance using new control filters and sensors with those currently installed in LIGO to estimate the performance benefit they would bring to the isolation systems.

Chapter 6 is based on a technical note describing a new way of designing sensor correction control filters through the use of particle swarm optimisation. The filters designed with this technique reduce the velocity RMS of the injected ground motion to the isolation tables by up to 70% when compared against the current filters using the same input motion. The filters designed by this method not only reduce the ground motion injection but also reduce the gain peaking present, which in turn should improve the controllability of the detector.

Chapter 2

Interferometry

The hard we'll do immediately, the impossible might take some time.

Alan Brown

This chapter describes the construction of a compact interferometer for use as either an independent displacement sensor or for a readout mechanism for inertial sensors, such as those covered in chapter 4. Text and figures of this chapter are copied from the paper, 'A compact, large-range interferometer for precision measurement and inertial sensing [5]', of which I was the principal author, the text and figures have been expanded upon from this paper. I designed, built and tested the prototype sensor. Some of the elements used in this chapter were taken from [50] of which I was the second author, however figures have been remade and the description rewritten to be more thorough.

2.1 Simple Interferometers

Optical interferometers are devices that measure the interference of two or more beams of light. Such devices are commonly used to measure the distance between two reflecting surfaces. Interferometers can be split into two distinct types, homodyne and heterodyne interferometers. In the

case of homodyne interferometry, a single frequency of light is used to perform the measurement, whereas in heterodyne interferometry two laser beams with two different frequencies, separated typically by a few MHz, are used. These are either generated by adding an Acousto Optic Modulator (AOM) or by using two separate lasers. The work covered in this chapter is on the former of these two types, a description of both interferometer types can be found in a review on compact interferometers by Watchi et.al [50] and a description of the conventions used for calculating the outputs of the interferometers detailed in this chapter can be found in Bond et al. [51].

2.1.1 Michelson Interferometer

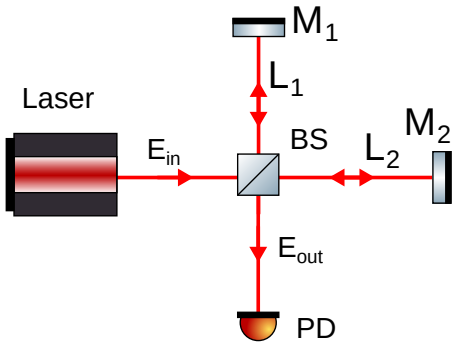

Figure 2.1: Figure shows the basic schematic for a Michelson interferometer. The input and output electric fields are denoted by E_{in} and E_{out}. The power measured by the Photodiode (PD) is given by equation 2.3

.

One of the simplest interferometers is the Michelson interferometer, a sketch of which is shown in FIG 2.1. Here a laser beam is split into two by a central beamsplitter. The two beams travel perpendicular to each other along the two interferometer 'arms' $L1$ and $L2$, defined as the distance between the beamsplitter and the end mirror. The light is reflected by the two end mirrors back to the central beamsplitter where it interferes and recombines. The interference pattern, and therefore the differential distance between the two arms is measured by a photodiode on the output port.

To calculate the expected form of the output signal for a simple Michelson, the input electric field

must be propagated through FIG 2.1 the electric field is defined as being a complex number of the form $E_{\rm in} = Ae^{i\phi}$, where A is the amplitude of the electric field, and ϕ is the initial phase. The initial phase is usually ignored as it cancels out when calculating the output power by taking the complex conjugate of it.

The electric field, after reflecting off the end mirror and accumulating phase along the arms of the interferometer is given by the equations,

$$E_{L1} = E_{\rm in} r e^{i\phi_1}, \tag{2.1}$$

$$E_{L2} = E_{\rm in} i t e^{i\phi_2}. \tag{2.2}$$

Where ϕ_1 and ϕ_2 are the round trip phase accumulated in each of the two arms, and r and t are defined as $r = \sqrt{R}$ and $t = \sqrt{T}$ are the reflectivity and transmissibility coefficients respectively. To conserve power, we shall use the convention that the transmitted beam from a beamsplitter will gain an additional factor of i. At the beamsplitter, the two beams pick up factors of it and r respectively, and can be summed to give the total output electric field, which has the form,

$$E_{\rm out} = E_{\rm in} i r t \left(e^{i\phi_1} + e^{i\phi_2}\right). \tag{2.3}$$

These phases can be defined in terms of the common ϕ_c and differential ϕ_d phases, given by

$$\phi_1 = \frac{\phi_c + \phi_d}{2} \qquad \text{and} \qquad \phi_2 = \frac{\phi_c - \phi_d}{2}. \tag{2.4}$$

Substituting in ϕ_c and ϕ_d for the individual arms transforms the output electric field into the form,

$$E_{\rm out} = E_{\rm in} i r t e^{\frac{i\phi_c}{2}} \left(e^{\frac{i\phi_d}{2}} + e^{-\frac{i\phi_d}{2}}\right). \tag{2.5}$$

To calculate the output power, P_{out}, we take the modulus squared of the output electric field. The common parts of the optical phase cancel to unity, and the output power is a function of the input power $P_{\text{in}} = |E_{in}|^2$ [52] and the differential optical phase is

$$P_{\text{out}} = \frac{P_{\text{in}}}{2}\left(1 + \cos(\phi_d)\right). \tag{2.6}$$

Here ϕ_d is the differential optical phase and is given by,

$$\phi_d = 2k(L_1 - L_2), \tag{2.7}$$

where k is the wavenumber and is equal to $k = \frac{2\pi}{\lambda}$. Due to the periodic nature of the measured output power, Michelson interferometers are typically only able to measure the optical phase over a fraction of a wavelength. This is caused by the ambiguity that arises at the turning points of the output signal. This is further limited as, close to the turning points, the change in power will not be linear with the change in distance. Such devices are therefore unsuitable for use in applications that require large range measurement. An example of these are the optical seismometers described in chapter 4, where the primary microseismic peak ground motion can regularly exceed a micrometer, which is comparable to the laser wavelength of devices described in this chapter. However, it is possible to use the principles of this device to create one with similar resolution but with a much larger working range and belong to a group of interferometers called fringe counting interferometers.

2.2 Fringe Counting Interferometers

A standard two-beam interferometer has an operating range that is typically less than a quarter of a wavelength of path-length difference. To increase both the dynamic range and the operating range, without using actuators or modulation schemes, we employ a HoQI that can measure two nearly orthogonal quadratures of the interferometer output. In this case, we use a Mach-Zehnder interferometer with two independent recombination beamsplitters. A polarisation scheme is employed to generate the required differential phase shift [53].

The optical path of HoQI is shown schematically in FIG 2.2. Compared with the Easy to Use Compact Laser Interferometric Device (EUCLID) and the Innovative Laser Interferometric Angular Device (ILIAD) [7, 54], HoQI is simpler than these previous devices. The number of birefringent elements, which increase noise and non-linearity, has been reduced, and there are no longer waveplates in the arms of the interferometer. The tilt-compensation system, developed to reduce tilt-to-length coupling and increase the angular operating range of the instrument [55], was also removed. The double-pass nature of the 'cat's eye' system used in EUCLID, which resulted in parasitic interferometers with a relatively large arm-length mismatch, was also removed. This, in turn, reduces the frequency noise coupling into the measurement, reducing noise at low frequency.

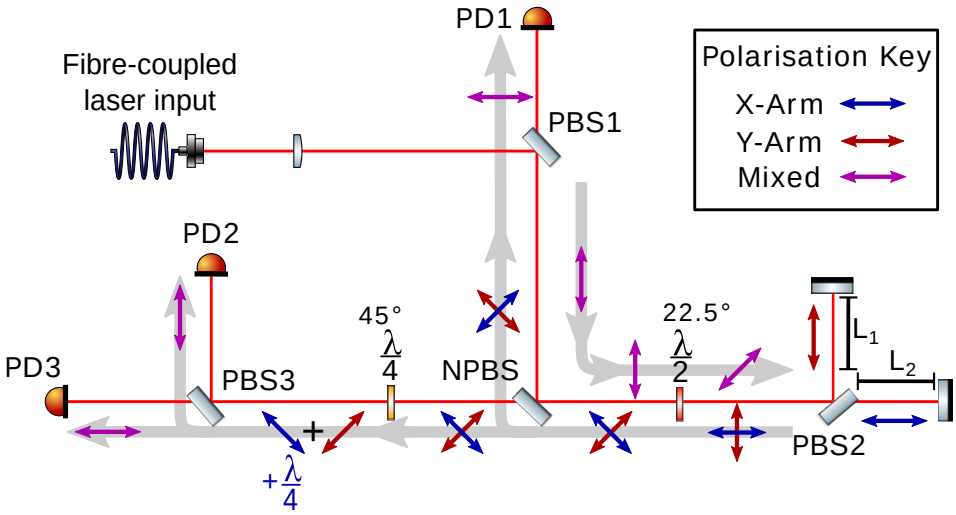

Figure 2.2: The optical layout of HoQI. Orthogonal polarisation states are used to track the length difference between L_x and L_y over multiple optical fringes. The input beam is split at polarising beamsplitter Polarising beamsplitter (PBS)2 and interferometrically recombined at PBS1 and PBS3, producing signals proportional to the sine, cosine, and minus cosine of the differential optical phase. Grey arrows indicate the direction of propagation of the beam.

To further reduce frequency noise coupling, we use a narrow-linewidth 1064 nm solid-state Innolight Mephisto 500NE laser (1 kHz linewidth for 0.1 s averaging period) in place of the Vertical-Cavity Surface-Emitting Laser (VCSEL) diode laser and we carefully match the arm lengths.

During our first tests of HoQI, we use carefully-aligned high-stability steering mirrors. To interrogate external targets, such as an inertial sensor reference mass, we will need to increase the angular operating range. At present, we intend to use a double-pass lens, placing a small focus on the target mirror. This will increase the divergence angle of the beam (up to ∼10 mrad), making us less susceptible to misalignment. For cases where larger operating ranges are required, a corner cube can be placed on the remote optic, and a large beam-size (∼2 mm) used within the interferometer, allowing operation over (at least) several degrees.

Assuming the target mirror remains aligned, the operating range of HoQI is only limited by the fringe visibility degradation due to spot-size changes, and is more than 10 mm for this configuration.

The laser light is fiber-coupled to the interferometer by a 2 m long single-mode polarisation maintaining fiber with an input power of 10 mW. The first Polarising BeamSplitter, PBS1, ensures there is a clean input polarisation state. With a sufficiently high Polarisation Extinction Ratio (PER) the first PBS effectively converts any polarisation fluctuations into input power fluctuations. These intensity fluctuations can be removed to first order, by subtracting pairs of photodiodes, shown in equations 2.22 and 2.23. PBS2 splits the input beam into two orthogonally polarised beams, one for each arm. These beams are recombined at PBS2 and co-propagate without interfering. The beam is divided, again without interference, at the Non-Polarising beamsplitter (NPBS). The quarter-wave plate before PBS3 then adds an additional phase shift of 90 degrees to the light from one of the arms such that when the beams interfere at PBS1 and PBS3, the resulting intensity fluctuations are 90 degrees out of phase.

To determine the output electric fields of HoQI we can use the output electric fields given by the Michelson interferometer shown in equation 2.3 as a starting basis for the electric field on the output of PBS2 in FIG 2.2. In this derivation, the beam incident on the Half Waveplate (HWP) is assumed to be composed of a single polarisation state, and both half and Quarter waveplate (QWP)s are rotated to their ideal angle of 22.5 and 45 degrees respectively. The PBSs are defined to reflect s-polarised light and transmit p-polarised light, though the extinction ratio of these components are deemed to be infinite in these calculations. Chapter 3 deals with imperfections in

these components and assumptions.

The input beam to the second PBS is defined as,

$$E_{\text{in}}^{\text{PBS2}} \begin{pmatrix} \text{L}_1 \\ \\ \text{L}_2 \end{pmatrix} = \frac{E_{\text{in}} e^{\frac{-i\pi}{2}}}{2} \begin{pmatrix} 1 \\ \\ 1 \end{pmatrix}, \tag{2.8}$$

where E_{in} is the input electric field, and $e^{\frac{-i\pi}{2}}$ is the common phase gained by both polarisation states from the half wave plate and $\binom{\text{L}_1}{\text{L}_2}$ represents the two polarisation states present in the interferometer in the basis of the second PBS. The half waveplate is tuned to an angle of 22.5° to ensure that equal powers are present in both polarisation states, shown by equation 2.8. The beams are split according to their individual polarisation state and accumulate an optical phase proportional to the path length of each arm of the interferometer. The output beam from PBS2 is then,

$$E_{\text{out}}^{\text{PBS2}} \begin{pmatrix} \text{L}_1 \\ \\ \text{L}_2 \end{pmatrix} = \frac{E_{\text{in}} e^{\frac{-i\pi}{2}}}{2} \begin{pmatrix} e^{i\phi_1} \\ \\ -e^{i\phi_2} \end{pmatrix}, \tag{2.9}$$

where ϕ_1 and ϕ_2 are the optical phases along the two arms of the interferometer. After a second pass through the HWP, the two copies of the beam on output of the NPBS are therefore,

$$E_{\text{R}}^{\text{NPBS}} \begin{pmatrix} \text{L}_1 \\ \\ \text{L}_2 \end{pmatrix} = -\frac{E_{\text{in}} e^{-i\pi}}{4} \begin{pmatrix} e^{i\phi_1} + e^{i\phi_2} \\ \\ e^{i\phi_1} - e^{i\phi_2} \end{pmatrix}, \tag{2.10}$$

$$E_{\text{T}}^{\text{NPBS}} \begin{pmatrix} \text{L}_1 \\ \\ \text{L}_2 \end{pmatrix} = -i\frac{E_{\text{in}} e^{-i\pi}}{4} \begin{pmatrix} e^{i\phi_1} + e^{i\phi_2} \\ \\ e^{i\phi_1} - e^{i\phi_2} \end{pmatrix}. \tag{2.11}$$

The Jones Matrices for transmission (PBS_T) and reflection (PBS_R) from an ideal PBS are defined as,

$$PBS_T = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \mathrm{PBS_R} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \tag{2.12}$$

The electric field and therefore the power, using $P_{PD1} = |E_{PD1}E^*_{PD1}|$, are then,

$$E_{\mathrm{PD1}} = \mathrm{PBS_T E_T^{NPBS}}, \tag{2.13}$$

$$E_{\mathrm{PD1}} = \frac{-E_0 e^{-i\pi}}{4}(e^{i\phi_1} + e^{i\phi_2}), \tag{2.14}$$

$$P_{\mathrm{PD1}} = \frac{P_{\mathrm{in}}}{8}(1 + \cos(\phi_1 - \phi_2)). \tag{2.15}$$

The beam reflected by the NPBS passes through a quarter wave plate which has its fast axis aligned with the polarisation of the transmitted beam at PBS2, retarding the beam from $L2$ by 90 degrees. The beam that is incident on PBS3 is therefore,

$$E_{\mathrm{PBS3}} \begin{pmatrix} \mathrm{PD2} \\ \mathrm{PD3} \end{pmatrix} = -i\frac{E_{\mathrm{in}} e^{\frac{-5i\pi}{4}}}{4} \begin{pmatrix} e^{i\phi_1} + ie^{i\phi_2} \\ e^{i\phi_1} - ie^{i\phi_2} \end{pmatrix}. \tag{2.16}$$

Using equation 2.12 we can then calculate the the powers measured by PD2 and PD3, using the same method as was used to calculate equation 2.15. These are,

$$P_{\mathrm{PD2}} = \frac{P_{\mathrm{in}}}{8}(1 + \sin(\phi_1 - \phi_2)), \tag{2.17}$$

$$P_{\mathrm{PD3}} = \frac{P_{\mathrm{in}}}{8}(1 - \sin(\phi_1 - \phi_2)). \tag{2.18}$$

For convenience, we phase shift each of the signals by 90°. The power measured on the photodiodes

is now given by the following equations,

$$\text{PD1} = \frac{P_{in}}{8}(1 + a\sin(\phi_{opt})), \tag{2.19}$$

$$\text{PD2} = \frac{P_{in}}{8}(1 + a\cos(\phi_{opt})), \tag{2.20}$$

$$\text{PD3} = \frac{P_{in}}{8}(1 - a\cos(\phi_{opt})), \tag{2.21}$$

$$\text{PD1} - \text{PD2} = \frac{\sqrt{2}aP_{in}}{8}\sin(\phi - \frac{\pi}{4}), \tag{2.22}$$

$$\text{PD1} - \text{PD3} = \frac{\sqrt{2}aP_{in}}{8}\sin(\phi + \frac{\pi}{4}), \tag{2.23}$$

where P_{in} represents the input power, a is the fringe visibility and ϕ_{opt} represents the differential optical phase and is defined as $\phi_{opt} = \frac{4\pi(L_x - L_y)}{\lambda}$. Equations 2.22 and 2.23 show how these signals can be combined to provide substantial common-mode rejection of laser intensity noise by reducing the dependence on both the input power and the fringe visibility. Equations 2.19 is plotted against equation 2.20 in FIG 2.3 creating a Lissajous figure.

Unwrapping the 4-quadrant arctangent of equations 2.22 and 2.23 returns the optical phase. To achieve high resolution, each photodiode signal is digitised with a high dynamic range, 18-bit Analog-to-Digital Converter (ADC) and the arctangent is performed using a cordic engine implemented on an Field Programmable Gate Arrays (FPGA). The analogue front-end and digital processing use an electronics module developed for the EUCLID and ILIAD interferometers [54], which have exceptionally low input-referred noise at low-frequencies and a proven signal processing chain. The displacement-equivalent noise of the readout electronics is shown in FIG 2.9, and it is what enables the high precision reported here.

2.3 Sensor Requirements

On the 14th September 2015 Advanced LIGO made the first direct detection of gravitational waves [14, 56]. To achieve the extraordinary sensitivity required for this discovery, LIGO uses a complex configuration of suspended mirrors to enhance the signal-to-noise performance of the detector. The mirrors are held at a precise operating point via closed-loop feedback systems to ensure that the

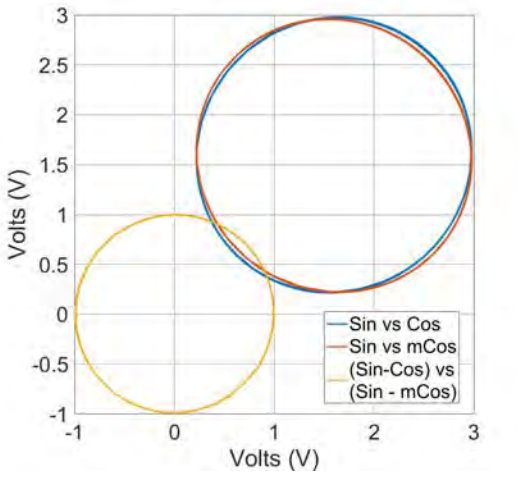

Figure 2.3: Figure showing the typical signals produced by HoQI when measuring distances more than one optical fringe. Here the three photodiode signals are plotted against each other and are used during alignment. The yellow trace, showing the result of equations 2.22 and 2.23 is gain balanced and used to calculate the distance measured.

laser light is resonant in the various optical cavities in the interferometer.

In order to reduce the required feedback forces, and associated noise, all core interferometer components are placed on ISI systems to reduce their inertial and relative motion. The ISIs employ many high-precision inertial and position sensors to reduce the transmission of ground motion [45]. Additionally, the core optics are mounted inside multi-stage suspension systems that are actively damped using local position sensors [8, 57].

Motivated by the goal of improving local sensing in gravitational-wave detectors, we present a compact interferometer based on the EUCLID and ILIAD sensors developed at Birmingham [7, 8, 54]. There are two specific applications within LIGO where such a device could be readily employed: as a replacement for the local position sensors in the suspensions, currently shadow-sensors called Birmingham Optical Sensor and Electro-Magnetic actuator (BOSEM)s [58]; and as a replacement for the coil-magnet readout of Geotech GS-13 geophones. With a focus on the second application, we develop sensitivity requirements to be of interest for LIGO and make an estimate of the potential impact on the observatories.

There exist, however, a large range of other possible applications. Within our narrow focus, we include a comparison with past compact interferometers and other LIGO position sensors, and an analysis of the performance of a Watt's Balance [59] with interferometric readout. Wider applications include, but are not limited to, atom interferometers [60, 61, 62], particle accelerators [63], and drag-free control of satellites [54].

2.3.1 Inertial Sensor Readout

The inertial sensors employed by LIGO have internal noises that are substantially higher than the suspension thermal noise limit of their proof masses [64, 65]. The readout mechanisms used in high precision inertial sensors are generally either inductive, capacitive, or optical. Capacitance based readouts can achieve high precision (e.g. [66]) but the sensor electrodes must be positioned very close to the target object, limiting their operating range. They also apply significant forces to the object, as well as having a large spatial force derivative (i.e. stiffness), which may be problematic for a suspended mass. This can be partially alleviated by use of multiple electrodes whose contribution to the force and stiffness can be made to cancel [67], but the residual effect may still be too great for some applications. Additionally, because the electrodes generally comprise extended plates, the capacitance will depend on some combination of displacement and attitude, directly coupling tilt to displacement. The drive signal of capacitance measurements may also pollute their environment with audio frequency electric fields, which are undesirable in sensitive experiments such as GW detectors [59].

Inductive sensors suffer from many of the same technical issues as capacitive sensors, including the trade-off between sensitivity and range and cross-coupling, but they are even more sensitive to EM interference [68, 69]. A final class of readout scheme, which should be considered separately from classical coil-driver inductance measurements, employs superconducting inductance measurements such as SQUIDs [70]. These can achieve very high sensitivity, but their cryogenic nature clearly makes them expensive and impractical for many applications.

Using interferometers to measure the proof mass position has the potential to remove some of the existing limitations in readout and actuation noise, while circumventing the technical challenges inherent in capacitative, inductive, and superconducting sensors. Other groups [71, 72, 73, 74, 75] have had success in improving the performance of inertial sensors using optical readout to both commercial and custom mechanics. We propose to extend the state-of-the-art by combining interferometric readout with commercial inertial sensor mechanics, improving sensitivity below the noise floor of the best force-feedback seismometers, such as the Trillium T240 [76].

2.3.2 Sensitivity Requirements

At the LIGO detector sites the ground motion at 10 Hz is approximately 10 orders of magnitude larger than measured gravitational-wave signals. The use of complex, multi-stage, passive and active isolation systems attenuates input motion below other noise sources at frequencies above 10 Hz [34]. Seismic noise at frequencies below 1 Hz lies outside the sensitive band of the interferometer. Nevertheless, ground motion at these frequencies, where active feedback provides most of the isolation, can still increase the root mean square (RMS) motion of the interferometer mirrors enough to prevent operation. The primary contributions to residual motion between the optics (excluding earthquakes) comes from the secondary micro-seismic peak (typically between 0.15 and 0.35 Hz) and the coupling between tilt and translation (typically below 0.1 Hz) [77]. The secondary micro-seismic peak is caused by reflected waves from the shoreline that collide with incoming waves creating a downward pressure on the sea floor.

It is difficult to predict the effect of new instruments on LIGO; the control systems and behavior of the instrument is extremely complex. However, during the first observation runs, Advanced LIGO was unable to operate for approximately 18% of the time due to elevated wind and microseismic motion [2]. By reducing the RMS motion of the isolation platforms, the interferometer should be able to operate during a wider range of environmental conditions. Moreover, due to the implementation of phasemeter readout (sometimes called fringe counting), our interferometric sensors have a larger working range than both the GS-13's and T240's employed at LIGO. The extra range and

improved low-frequency sensitivity may improve the detectors' ability to stay 'locked' during small or remote earthquakes by suppressing only the differential inertial motion.

The control band for LIGO's active inertial isolation for the ISIs is approximately 100 mHz to 30 Hz [69]. At low frequencies the noise on the inertial sensing output increases as $\frac{1}{f^2}$. As such, the inertial signal is substituted with displacement sensors, effectively locking the isolated platforms to the ground below approximately 30 mHz. However, due to the constraints of causal filtering, the inertial sensors must perform well down to 10 mHz to avoid injecting sensor-noise or tilt-coupling. Performance requirements between 1 and 10 Hz mean that the unity gain frequency must be about 30 Hz and, as such, good inertial sensor performance (in both sensitivity and phase response) is needed up to 100 Hz. Beyond 100 Hz it is possible to rapidly reduce the loop gain and the sensitivity requirements are subsequently relaxed. For these reasons, to be of interest for Advanced LIGO (and other gravitational-wave detectors), any new inertial sensor should have sensitivity at least equal to state-of-the-art inertial sensors between 10 mHz and 100 Hz.

Further improvements to the detector's performance can be made by increasing the sensitivity of the BOSEM displacement sensors placed on LIGO's quadruple suspensions [58]. Due to the noise of the BOSEMs, local feedback forces can only be applied to the uppermost suspended mass, and even then the control filters have strict requirements imposed by the need to prevent sensor noise from spoiling the detector sensitivity at 10 Hz [57]. Interferometric displacement sensors would allow for improved damping of the top mass of the quad suspension system, as well as opening the possibility for local damping on lower stages, reducing both vibration transmission and settling time. To apply significant damping using a sensor at the Upper Intermediate Mass (UIM) [78] (the second stage of the 'quad' suspension from the ground), the noise of the sensor at 10 Hz should be of order 100 times smaller to exceed the increase in mechanical transmissibility at this frequency, and our measurements here more than satisfy this criterion [79, 80].

2.4 Development of the Sensor

Other sensors, such as EUCLID [54] and ILIAD [7] had previously been developed, though the limiting noise couplings were not well understood. An initial, bench-top version of HoQI was constructed, as shown in FIG 2.4, to investigate these noise couplings, such as polarisation fluctuations caused by the waveplates and beamsplitters. It was at this point that the cats eye retro-reflector, used in both EUCLID and ILIAD was dropped in favor of a pair of mirrors to simplify the setup, and eliminate parasitic interference that may have been present in the previous devices. Removing retro-reflectors removes the interferometer's ability to double pass the beam, reducing the potential resolution of the sensor by a factor of two. Considerable space was left between components, so the polarisation state and power after each component could be measured to investigate these issues. Tests on the effect of waveplate rotation to induce non-linearities were also conducted at this stage, but are covered in more detail in chapter 3.

The interferometer was constructed with one inch optics and posts with a 10 cm gap in between each of the components. This was done to allow for a power meter to be placed between each of the optics to map the polarisation state throughout the interferometer. The large area of the device highlighted some of the key noise sources that would need to be addressed in future versions of the interferometer, namely: frequency noise, thermal effects, air currents and acoustic effects. Birefringence noise, caused by imperfections in the polarising ability of optics were reduced through the use of high quality optics, with high extinction ratios and easy tunability of the device.

The frequency noise coupling is dependent on the path length difference between the two arms. In the initial version the arm lengths were not matched, resulting in large frequency noise coupling. To reduce this, the arms of the interferometer were fixed in length and the differential arm length was tuned by injecting a signal at 80 Hz and minimising the injected line, by adjusting one of the mirror mounts, thereby reducing frequency noise coupling into the interferometer. This noise coupling into the interferometer could then be estimated and plotted to ensure it was not responsible for any sensor noise.

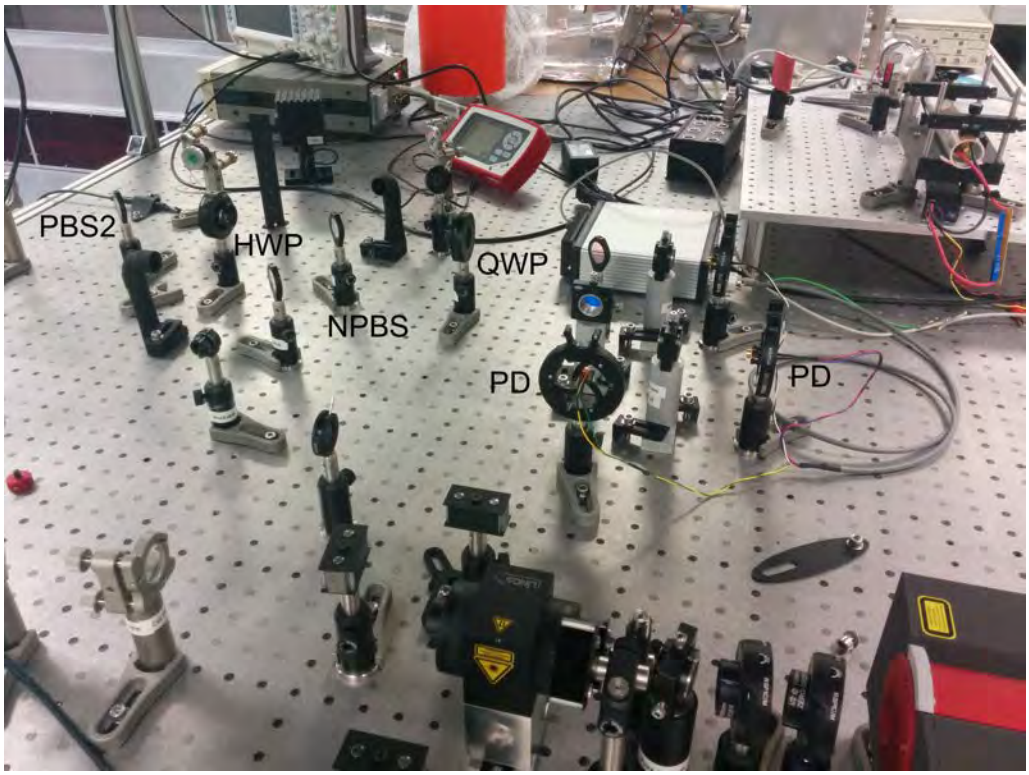

Figure 2.4: Figure showing the initial bench top version of HoQI, with key components labeled. The purpose of this version was to investigate noise couplings in the interferometer, rather than to evaluate the resolution of the sensor.

The noise coupling caused by thermal effects, air currents and acoustic effects can be reduced in two ways, the first is to reduce the size of the interferometer, increasing the proportion of the noise that will be common between the optics and therefore not appear in the differential signal. The amplitude of these noise sources can be reduced further by shielding the instrument in an acoustic insulated box or by placing the interferometer in a vacuum tank. FIG 2.5 shows the resolution of the first test of the HoQI sensor when in its diagnostic configuration, highlighting the areas that needed to be improved with the design. Future versions would need to be made more compact, to reduce the differential motion seen by the device and to enable shielding of the device from thermal effects and acoustic couplings.

Mentioned previously, in section 2.3, the end goal for these devices was to be employed as a compact displacement sensor or as a readout for inertial sensors for use in LIGO. To achieve this goal the

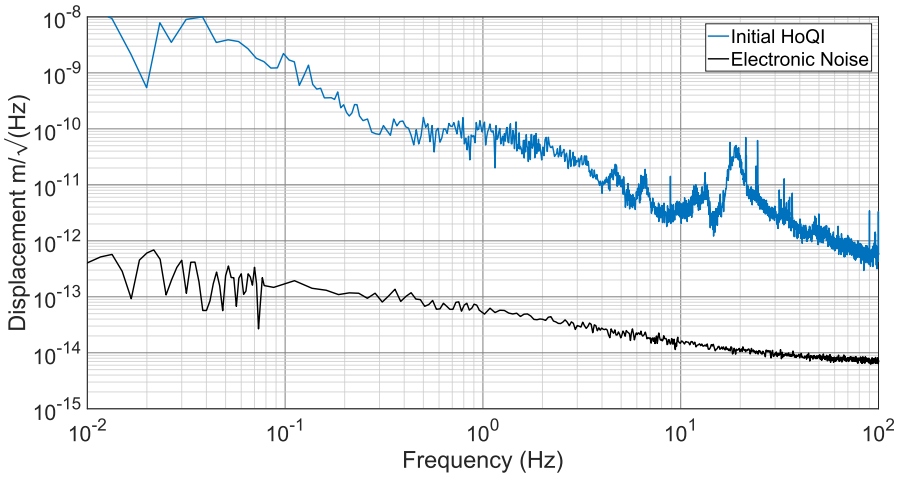

Figure 2.5: Figure showing the measured resolution of the initial version of HoQI (blue) compared to the noise of the measurement electronics (black). At low frequency, air current, thermal effects and frequency noise are likely the largest causes of noise. At high frequency the interferometer is likely limited by acoustic noise coupling.

device needed to shrink in size. The initial version used solely free-space optics, which meant, while the device was easy to diagnose, it was not very portable. Using the same one inch optics, the interferometer was placed onto a 30 cm by 30 cm breadboard using the same layout as presented in FIG 2.2.

This device was placed in low vacuum at a pressure of 10^{-2} mbar to reduce temperature fluctuations and eliminate air pressure fluctuations. The laser light was launched through a polarisation maintaining optical fibre through a window into the vacuum system. Fibre feed-throughs were found to induce large polarisation noise in the interferometer a low frequencies, and subsequently remain a problem to be solved. FIG 2.6 shows the interferometer when placed in a vacuum chamber.

FIG 2.7 shows the performance of the in-vacuum version of HoQI (red) when compared with the initial version of HoQI (blue). Reducing the arm length mismatch between the two arms and placing the device inside vacuum is thought to be the cause of the improvement between 10 mHz and 1 Hz. At 10 Hz the motion measured by the two versions of the interferometer converges, indicating that the source of the motion observed at high frequency is likely not due to direct

Figure 2.6: Image showing the interferometer in vacuum. Here, laser light was launched from an optical fibre through a window to the right of the image, to investigate the use of fibres as a method of delivering input light and to remove low frequency noise couplings such as temperature and air pressure fluctuations.

acoustic coupling, as the vacuum tank should have reduced its coupling factor, however, indirect coupling through the base of the vacuum chamber is still possible. There is a factor 10 suppression in the observed motion with the in vacuum measurement compared to the in-air measurement. Further tests revealed that the peak at 20 Hz is caused by a table resonance.

Coupling of the input light through a window of the vacuum tank, while getting around the polarisation noise induced by the fibre did result in the alignment of the interferometer moving when pumping down the vacuum tank. This complicated the measurement of the sensor's resolution as often the alignment would drift out of the range of the external tuning. This then meant venting the vacuum tank, re-aligning the optics and pumping the tank down again, resulting in a time intensive process.

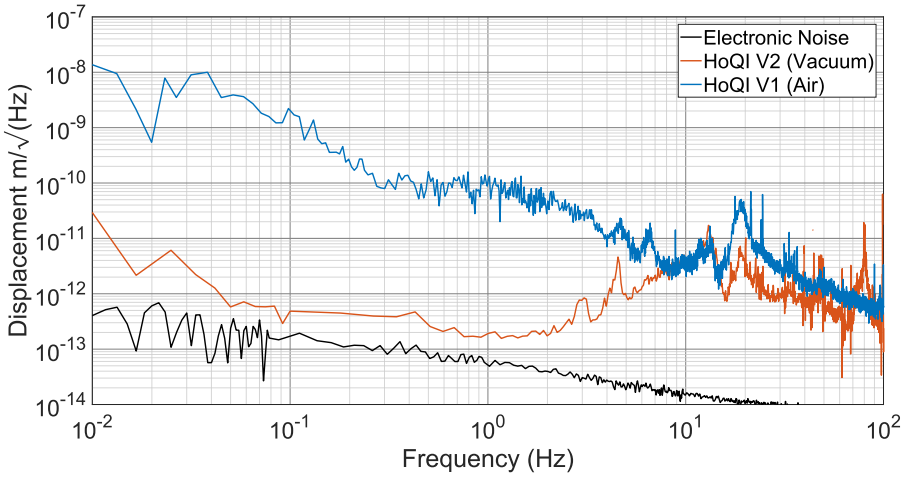

Figure 2.7: Figure showing the comparison between the resolution of the initial (blue) and in vacuum (red) versions of HoQI. The electronic noise (black) highlighting the best possible resolution of the sensor.

2.4.1 Small Prototypes

To further improve the resolution, the sensor needed to be reduced in size. The size of optics was reduced from one inch versions to half inch versions, with the tunability of the placements of the optics retained to ensure the operation of the device remained linear. Furthermore, tests would be conducted in-air while the in-vacuum fibre feedthrough issue was worked on in parallel to improve the repeatability of the testing procedure.

To investigate the sensitivity limits of HoQI, we reduced optical and mechanical noise where possible. The largest anticipated sources of noise were: mechanical vibration, thermal expansion and gradients, birefringence noise, frequency noise, and electronic noise. All optics were rigidly mounted close together on an aluminum baseplate with a relatively large thermal mass, seen in FIG 2.8, resulting in large common-mode rejection of mechanical noise and reducing thermal gradients.

Birefringence fluctuations between the non-polarising beamsplitter and the recombination polarising beamsplitters are indistinguishable from arm-length changes. Since the beams are well aligned and co-propagate, the dominant effect is expected to come from quarter-wave plate, and a high qual-

Figure 2.8: The prototype version of HoQI, the base plate is 170 × 100 mm with 10 mm gaps between components.

ity zero-order waveplate was used to reduce this. Alignment fluctuations on the photodiodes cause uncorrelated fluctuations in the photocurrent due to inhomogeneities in the quantum efficiency across the surface of the photodiode [81]. The single-mode fibre strips away pointing fluctuations, and the output mode is mechanically fixed to the baseplate by the fibre output collimator.

Frequency noise coupling was measured and minimised by modulating the laser frequency and adjusting the macroscopic arm-length difference to minimise the coupling to differential optical phase. The length was precisely tuned using the alignment screws on the 'end' mirrors, with a resolution of a few microns, but the coupling was much larger than predicted. This is attributed to interference from stray light. The residual coupling can be quantified by an effective arm-length mismatch of 0.7 mm. Using the specified frequency noise of the Innolight Mephisto 500NE laser, which has a frequency noise of $10^4 \times [\frac{1}{f}]\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$, we predict the red curve shown in FIG. 2.9.

2.4.2 Resolution of the Compact Sensor

The electronic noise (the black curve in FIG 2.9) is measured by replacing the photodiode inputs with a constant current using a resistor connected to a stable bias voltage. The resistor values are such that the 3 input currents simulate a specific optical phase for the three photodiodes.

The baseplate was placed on rubber 'feet' on an optical bench and sampled at 20 kHz over a 10 hour period. FIG.2.9 shows the amplitude spectral density of the measurement over a ten minute segment of this data. A time series plot of this data is shown in FIG 2.10 highlighting the drift of the sensor. The interferometer reaches a peak sensitivity of $2 \times 10^{-14}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 70 Hz. At 10 mHz a sensitivity of $7 \times 10^{-11}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ is achieved.

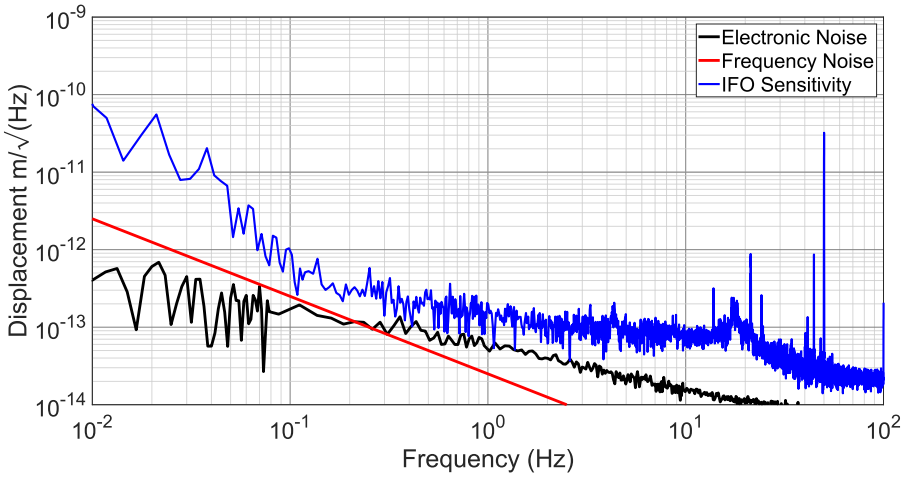

Figure 2.9: Sensitivity of the fibre-coupled prototype HoQI showing the interferometer signal (blue), the measured readout noise (black), and an estimate of the frequency noise that couples into the interferometer (red)

The total sensitivity is probably limited by electronic noise at frequencies near 0.5 Hz. Below this, the limiting factor is assumed to be a combination of air currents, temperature fluctuations, and frequency noise. Above 1 Hz, the sources of noise are less well understood except for the peak near 18 Hz, which is caused by mechanical vibration of the optical table, and the large peak at 50 Hz, caused by pickup in the unshielded photodiode cables.

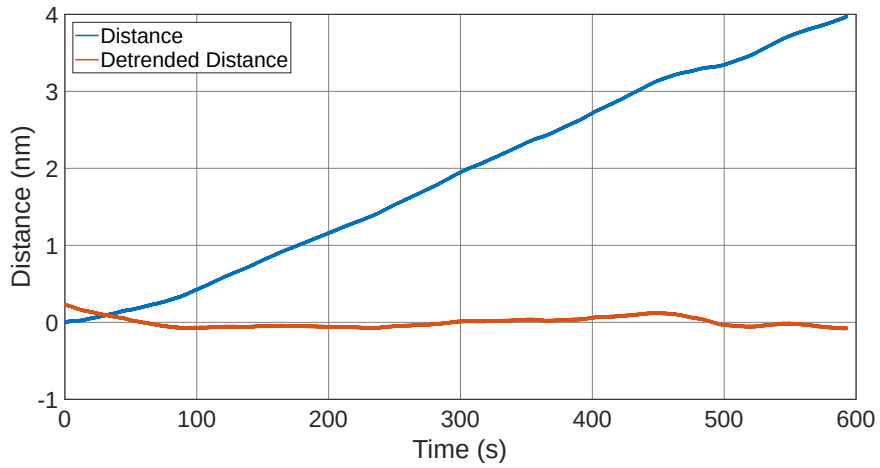

Figure 2.10: Figure showing the raw time series (blue) and detrended time series data (red) of FIG 2.9, highlighting the low frequency drift of the sensor.

FIG 2.11 compares the sensitivity of HoQI with the Capacitive Position Sensors (CPS), which are employed on the first stage of LIGO's ISI. In the frequency band of interest, HoQI has 250 times lower noise at 100 mHz and 1000 times lower noise at 10 Hz. When compared with the BOSEMs, the improvement is more substantial: HoQI has a factor of 500 lower noise at 100 mHz and 1000 times lower noise at 10 Hz.

In order to compare HoQI's readout noise with existing inertial sensors, we multiply the interferometer sensitivity curve by the inertial-sensing transfer function of both a GS-13 and a Watt's linkage similar to those employed at the Virgo gravitational-wave detector [59]. This can be done by multiplying the measured Amplitude Spectral Density (ASD) shown in 2.9 by the frequency response of the Zero-Pole-Gain (ZPK) system of a GS-13 and a Watts linkage. The result of this is shown in FIG 2.12. This readout-noise is then summed in quadrature with the estimated suspension thermal noise for each sensor. The mechanical thermal noise is given by,

$$F_{\mathrm{th}}(\omega) = \sqrt{4k_B T R(\omega)}, \tag{2.24}$$

where $F_{\mathrm{th}}(\omega)$ is the amplitude spectral density of the force due to thermal noise, T is the temperature, Q is the quality factor and $R(\omega)$ is the mechanical resistance (the real part of the mechanical impedance) [64]. For a simple mass-spring system with mass m, resonant frequency ω_0, the me-

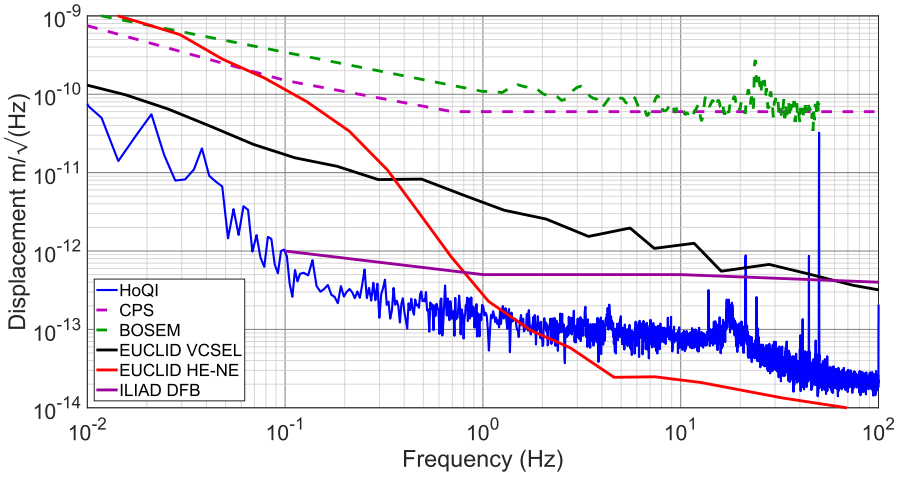

Figure 2.11: HoQI (blue) compared with other precision displacement sensors including: previous interferometers developed at Birmingham, ILIAD (purple) [7] and EUCLID with both an external HE-NE laser (red) and with its integrated VCSEL laser (black) [8]; and with devices used at LIGO, the 0.25 mm range CPS dashed purple), BOSEM (dashed green). The CPS and low-frequency BOSEM curves are stick-figure fits to noise spectra from multiple devices.

chanical resistance is given by,

$$R(\omega) = \frac{m\omega_0^2}{Q\omega} \tag{2.25}$$

The GS-13 is assumed to have a 5 kg proof-mass, a resonant frequency of 1 Hz, and a (structural-damping) quality factor of 40. The Watt's linkage, with its low mechanical-dissipation and resonant frequency, has lower thermal noise (everywhere) and lower readout noise below 1 Hz. For the suspension thermal noise calculation we assume a proof-mass of 1 kg, a resonant frequency of 0.3 Hz, and a (structural-damping) quality factor of 100.

The noise projections are compared with the self-noise floors of the GS-13 (using its conventional coil magnet readout) and a Trillium T240, both as measured at LIGO. We find that between 0.01 and 2 Hz the suspension thermal noise of the GS-13 would limit the resolution of a future optically readout inertial sensor, based on GS-13 mechanics. To fully exploit the sensitivity of the interferometer presented in this thesis, mechanics with a lower suspension thermal noise would

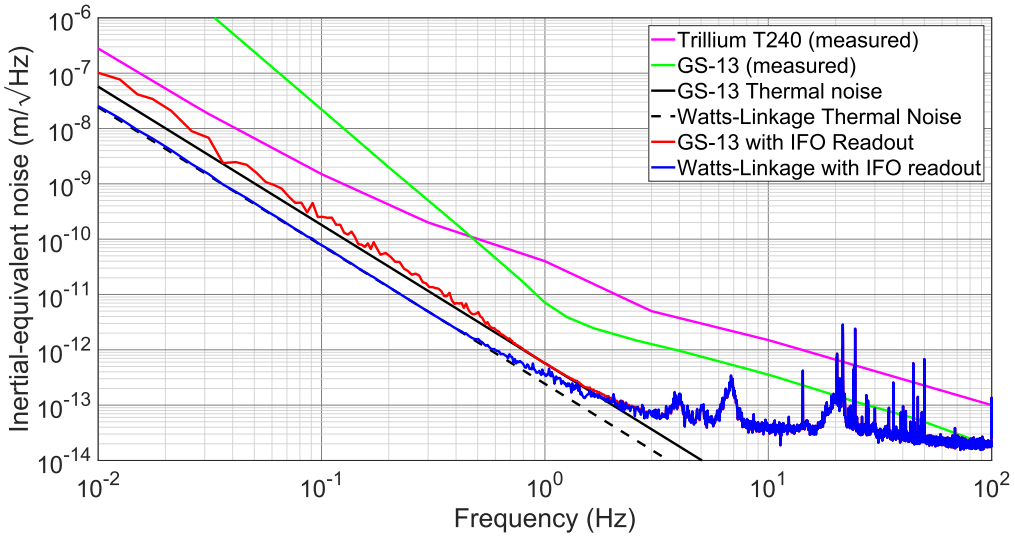

Figure 2.12: The sensitivity of HoQI projected onto a GS-13 (red) and a Watt's linkage (blue) is compared with a GS-13 using conventional readout (green), and a Trillium T-240 force-feedback seismometer (magenta). The (calculated) suspension thermal noise of the GS-13 (black) [9] and Watt's linkage (dashed black) are also shown.

have to be evaluated. Increasing the structural Q of the spring reduces this thermal noise, and an improvement in the resolution between the optically readout GS-13 and Watts linkage can be seen when the interferometer is used. Despite the thermal noise limitation, using HoQI to interrogate a GS-13 could increase the sensitivity by a factor of 100 at 100 mHz and would improve it at all frequencies up to 100 Hz. The estimated noise curves of the GS-13 and the Watts linkage converge at frequencies above 2 Hz as they are limited by the interferometer readout noise.

2.5 Further Development and Improvements

For integration onto inertial sensors and suspension chains in LIGO, the interferometer must be further reduced in size, while keeping the same basic layout and at least maintaining the current sensor resolution. FIG 2.13 shows a photo of a smaller version of HoQI, measuring around 8 by 6 cm and 3 cm tall. The size reduction was achieved by designing new mounts for a number of different components, most noticeably the waveplates. The waveplates are held in the mounts

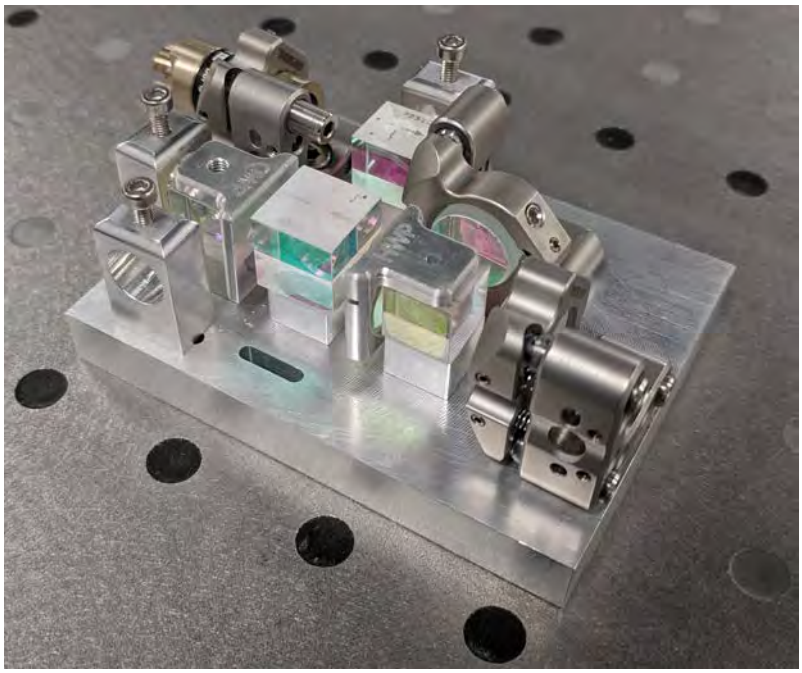

Figure 2.13: A compact prototype of the HoQI sensor used to evaluate the noise performance of the sensor when compacted further. The device pictured measures 8 cm by 6 cm,

and during initial alignment can be rotated to markings on the face of the mount to optimise the linearity of the device. These mounts are bonded to the tops of the beamsplitters and are designed to not touch the coatings on the beamsplitters. For ease of assembly and alignment slots for placing beamsplitters have been milled into the interferometer baseplate to ensure that the beamsplitters are placed in the correct location. The beamsplitters and their spacers are glued together, such that their sides are flush and are then glued to the baseplate. These changes mean that the alignment of the test arm of the interferometer can be achieved by adjusting the input fibre coupler alone.

To ensure that the interferometer is linear in its response, the beam input to the HWP needs to be in a single, well defined polarisation state. The cube PBSs that have been used up to this point have poor PERs. The cubes used in HoQI transmit p-polarised light and reflect s-polarised light. Typically, good PBSs have extinction ratios on the order of 1000:1, meaning if we look at the transmitted port of a PBS, we should see one part in 1000 in the s-polarised state. On

transmission, the extinction ratios were found to be satisfactory, having extinction ratios between 1400:1 and 2000:1. However, on reflection, the extinction ratios were measured to be between 10:1 and 80:1, these low extinction ratios will result in polarisation leakage. This matter is made worse as the ratio of these states will depend on the initial polarisation state, having the effect of reducing the fringe visibility, as a fraction of the beam will interfere, forming spurious interferometers. To reduce the error in the interferometer sensitivity caused by the poor PERs of the beamsplitters, a 45° mirror is used after the input fibre coupler to increase the purity of the polarisation state entering the interferometer. This increase in the purity of the input polarisation ensures that the beam incident on the HWP gets split into only two seperate polarisation states of equal power. This is important if corner cubes are used in place of mirrors in the arms.

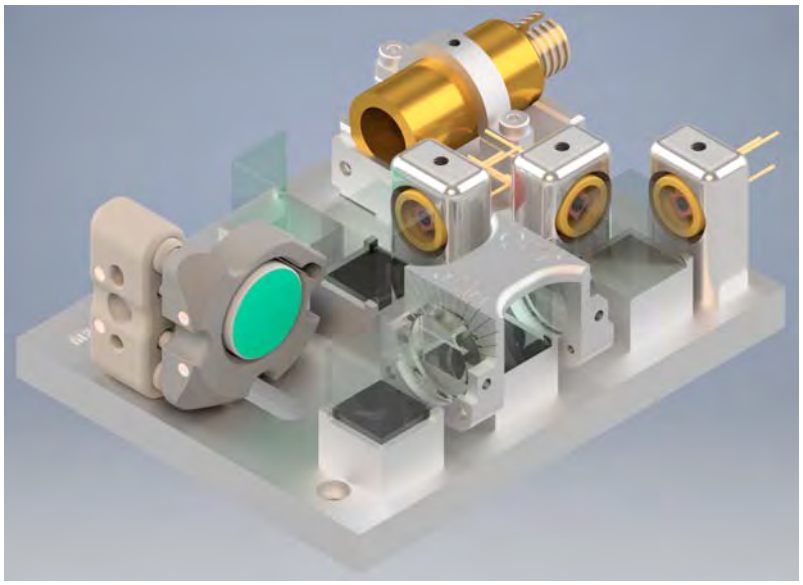

Figure 2.14: A render of an updated design of HoQI featuring new waveplate mounts, a new fibre mount and beamsplitter holders as well as a 45 degree mirror to account for the poor reflection PERs of the PBS used. Discussion about the specific design of the baseplate and associated new mounts can be found in chapter 4

Currently, mirrors are used on the ends of the reference and test arms, and as such a half inch kinematic mount is used to provide the alignment for the reference arm to ensure overlap with the beam from the test arm. Corner cubes are being investigated for use, as these would all but

eliminate the angle sensitivity of the device, though this would require a small redesign of the base-plate. These changes have greatly simplified the construction and alignment of the interferometer, as such, one can be built, aligned in less than an hour. A Computer Aided Design (CAD) render of these design changes is shown in FIG 2.14; and are discussed in detail in chapter 4.

2.6 Summary

Compact homodyne phasemeters provide an excellent method of accurately measuring the motion of external targets over large working ranges. In the course of this chapter the HoQI sensor has been developed, with noise sources such as polarisation, frequency and intensity noise have each been reduced. Over the course of this development the design of the interferometer has been refined to improve the ease of assembly and alignment of the sensor, reducing the total time for construction and alignment from several hours to less than an hour.

These improvements result in the self noise of the HoQI sensor to have a much higher resolution than other displacement sensors used in LIGO despite being tested in air without any seismic isolation present. When compared to the CPS and BOSEM sensors, HoQI has a 1000 times higher resolution than both these sensors at 10 Hz. When compared with other sensors of their type, such as the EUCLID, HoQI is 10 times more sensitive at 10 Hz.

When used as readout mechanisms for inertial sensors, compact interferometers should be able to reduce the self noise of an inertial sensor, such as the GS-13 by approximately by a factor of 1000 at 10 mHz. Such a device is predicted to be close to limited by the thermal noise of the suspended masses spring below 1 Hz.

Chapter 3

Non-Linearities in Homodyne Phasemeters

Nothing in life is to be feared; it is only to be understood. Now is the time to understand more so that we may fear less.

Marie Curie

This chapter describes the creation of a polarisation model used to propagate the effects of optic misalignments and imperfections that create non-linearities in homodyne phasemeters. Text and figures in this chapter are sourced from a mature draft of the paper 'Quantifying cyclic non-linearities of interferometers in the frequency domain', of which I was the principal author. My role in the paper was in the creation of the model and simulating the effects of non-linearities for a couple of use cases. At of the time of writing, this is unpublished though a mature draft of the paper can be found on the LIGO DCC [82]; sections of the paper that are not of my work, detailing the analytical derivation of non-linearities are excluded from this thesis. Text and figures have been added in addition to those sourced from the paper to provide a more detailed explanation of the topic.

3.1 Introduction to Non-Linearities

Over the past 40 years, increasingly small optical readouts, based on homodyne interferometry, have been developed, such as those described by Downs [83] and Speake [54]. The resolution reached by the most recent compact versions is on the order of 10^{-13} m/$\sqrt{\text{Hz}}$ at 1 Hz [5, 73]. A review of linearity in both heterodyne and homodyne phasemeters can be found in [50]. This high resolution is possible due to careful reduction of noise sources, such as frequency, intensity and birefringence noise. However, there is still a mismatch between the measured phase and the real phase in the readout. These distortions are referred to as non-linearities, named after the type of coupling they have on the readout. Typically, these non-linearities are induced by imperfections or misalignments in the optics. The linearity is defined as the relation between the real and measured optical phase and any discrepancy between these is referred to as a non-linearity.

Previous research into non-linearities has quantified the results in terms of RMS errors in the time domain [84, 85, 86], focusing on reducing the RMS error between the measured and real phase. Research on developing simple models has been developed by de Groot [87] on heterodyne interferometry, while Stone and Howard [88] detail such a model on homodyne interferometers. Data processing techniques, such as ellipse fitting, have previously been used to measure the ellipse parameters and correct for them in post processing, first shown by Heydemann [89] and used on optical inertial sensors by Zumberge [71]. While the RMS error caused by non-linearities is important, the frequencies affected by these phase errors are important for the response of isolation systems or inertial sensors with high Q factors, e.g. non-linearities present in seismic signals in the kHz region would have little impact in applications such as Advanced LIGO. This chapter aims to understand how these non-linearities affect the signal and how they contribute to the signal in the frequency domain. Such high Q dynamic systems increase the motion seen by the device substantially and are presented later in this chapter.

Next generation inertial sensors with significantly higher resolution use optical readouts instead of coil magnet based readouts. These sensors are needed to improve control systems present on

large physics experiments [90]. Sufficiently quiet isolation platforms are mandatory for the proper functioning of these instruments and need to be improved in order to expand the capability of these detectors [47]. Sensors, such as interferometric inertial and displacement sensors, are required to achieve a high level of isolation and to sense the residual motion of the equipment.

Fringe counting interferometers, such as those described in chapter 2, are one such candidate for improving the controllability of the detector and for reducing sensor noise coupling. They rely on measuring two quadratures of the optical phase, by splitting the optical signal in two and phase shifting one with respect to the other by 90 degrees. This method allows for the precise measurement of the optical phase of one part in 10^8 of a wavelength, while being able to track the optical phase over multiple wavelengths. The main culprits for inducing non-linearities in interferometers, such asHoQI, are the half and quarter waveplates, and the beamsplitters. The quarter waveplate, responsible for generating the two quadrature signals, is the largest single source of non-linearities in the interferometer. As described in chapter 2 the optical phase is extracted by taking the arctangent of two quadrature signals and in the absence of offsets, is given by,

$$\text{PD1} = \frac{P_{in}}{8}(a\sin(\phi_{opt})), \tag{3.1}$$

$$\text{PD2} = \frac{P_{in}}{8}(a\sin(\phi_{opt}+\frac{\pi}{2}+\delta)), \tag{3.2}$$

where $\frac{\pi}{2}$ is the expected phase shift between the two quadratures and δ is a small error in the phase shift. When δ is zero, the optical phase can be extracted easily using,

$$\phi_{opt} = \arctan\left(\frac{\text{PD1}}{\text{PD2}}\right) \tag{3.3}$$

However, when δ is non zero, the two quadratures given by PD1 and PD2 are no longer orthogonal, meaning that the method of extracting the optical phase given by equation 3.3 is no longer valid. This results in a discrepancy between the extracted optical phase and the real phase as measured by the interferometer. Errors caused by the half waveplate and deficiencies in the beamsplitters

cause unequal powers in each polarisation state. This results in the offsets and powers as seen by each quadrature not canceling; invalidating the assumptions made in equation 3.3.

3.2 Correction of Non-Linearities

It is possible to not only reduce but in some cases eliminate the effects of non-linearities on the readout by performing a technique known as ellipse fitting shown by Heydemann [89]. The ellipses in the Lissajous plot, shown in FIG 2.3 in chapter 2 can be fitted using a least squares method, described by Rosin [91],is summarised in the section below, and derived fully in Appendix A.1

Ellipse Fitting Routine

The ellipse fitting routine takes in x and y data from the Lissajous figure. In this case $x =$ SPD$-$CPD and $y =$ SPD$-$MCPD, where CPD, SPD, MCPD are the cos, sin and $-$cos photodiode outputs. The script then creates a matrix of parameters, containing values of x^2, xy, y^2, x and y, which in the following notation are contained in vector X. In least squares fitting we take a function, $R^2 = G^2(x_i, y_i)$, where G is the standard quadratic form of a rotated ellipse, given by equation 3.6, and minimise it with respect to each of the parameters, denoted by the vector H, where H, X and G are defined as,

$$H = \begin{bmatrix} A & B & C & D & E \end{bmatrix}, \tag{3.4}$$

$$X_i = \begin{bmatrix} x_i^2 & x_i y_i & y_i^2 & x_i & y_i \end{bmatrix}, \tag{3.5}$$

$$G = \begin{bmatrix} Ax_i^2 & Bx_i y_i & Cy_i^2 & Dx_i & Ey_i & F_i \end{bmatrix}, \tag{3.6}$$

where H is the coefficients of the ellipse parameters. We then define the function that we wish to

minimise, R^2 as,

$$R^2 = \sum_{i=1}^{n} \left(Ax_i^2 + Bx_iy_i + Cy_i^2 + Dx_i + Ey_i + F\right)^2, \tag{3.7}$$

$$= \sum_{i=1}^{n} (HX_i + F_i)^T (HX_i + F_i), \tag{3.8}$$

$$= \sum_{i=1}^{n} H^T X_i^T X_i H + F_i(X_i H) + F_i(X_i^T H^T) + F_i^2. \tag{3.9}$$

We then differentiate this function with respect to H, to find the expression for our minimised parameters, given by,

$$\frac{\partial R}{\partial H} = H^T X^T X + \sum_{i=1}^{n} F X_i = 0, \tag{3.10}$$

$$H^T X^T X = \sum_{i=1}^{n} -F x_i, \tag{3.11}$$

$$H = -\sum_{i=1}^{n} \frac{F x_i}{X^T X}. \tag{3.12}$$

When running the interferometer, the ellipse parameters given by equation 3.12 can be determined by measuring motion on the order of a wavelength. Now that the ellipse parameters have been determined, they can be converted into the gains, offsets and rotations of the ellipse; the full derivation of this is shown in Appendix A.1 and is summarised below. FIG 3.1 shows the significance of each of the ellipse parameters.

The equation of a simple non rotated ellipse, with semi-major and minor axes, a and b is given by,

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1. \tag{3.13}$$

In principle, due to imperfections or mis-alignment of the quarter waveplate, equation 3.13 must be rotated by angle α, giving the rotated x and y variables that we shall fit, this is given by,

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{pmatrix} x' - x_0 \\ y' - y_0 \end{pmatrix}. \tag{3.14}$$

Coefficient Label	Value
A	$a^2\sin^2(\alpha) + b^2\cos^2(\alpha)$
B	$(b^2 - a^2)\sin(2\alpha)$
C	$a^2\cos(\alpha) + b^2\sin(\alpha)$
D	$-2x_0A - y_0B$
E	$-2y_0C - x_0B$
F	$x_0y_0B + x_0^2A + y_0^2C - (ab)^2$

Table 3.1: Table showing the links between the physical ellipse parameters and the least squares coefficients, a, b are the gains, x_0, y_0 are the offsets and α is the rotation of the Lissajous from circularity.

Substituting in our values for x' and y', which represent the optical phase, and adding an offset, we recover the form of our rotated ellipses before fitting, these are,

$$x = (a\cos(\phi) - x_0)\cos(\alpha) - (b\sin(\phi) - y_0)\sin(\alpha), \tag{3.15}$$

$$y = (a\cos(\phi) - x_0)\sin(\alpha) + (b\sin(\phi) - y_0)\cos(\alpha). \tag{3.16}$$

This can be substituted into equation 3.13 which links our least squares parameters, to physical offsets, gains and ellipse rotations, and are summarised in table 3.2.

These coefficients can then be linked back to physical parameters, such as the gains, offsets and rotations between the two quadratures, this derived in full in Appendix A.1 and is summarised in table 3.1.

Reshaping the Ellipse into a Unitary Circle

With the physical ellipse parameters extracted, the ellipse can be re-shaped into a circle. This process limits the amount of up-conversion of noise in the output ASD. The ellipse parameters

Coefficient Label	Value
α	$\frac{1}{2}\arctan(\frac{B}{A-C})$
a	$\sqrt{\frac{(A+C)-\sqrt{(A-C)^2+B^2}}{2}}$
b	$\sqrt{\frac{(A+C)+\sqrt{(A-C)^2+B^2}}{2}}$
x_0	$\frac{2DC-EB}{B^2-4AC}$
y_0	$\frac{2AE-BD}{B^2-4AC}$

Table 3.2: Table linking the coefficients from the least squares fitting to physical photodiode gains, offsets and the error on the two quadrature, a, b are the gains, x_0, y_0 are the offsets and α is the rotation of the Lissajous from circularity.

referred to in the rest of this section are defined by FIG 3.1.

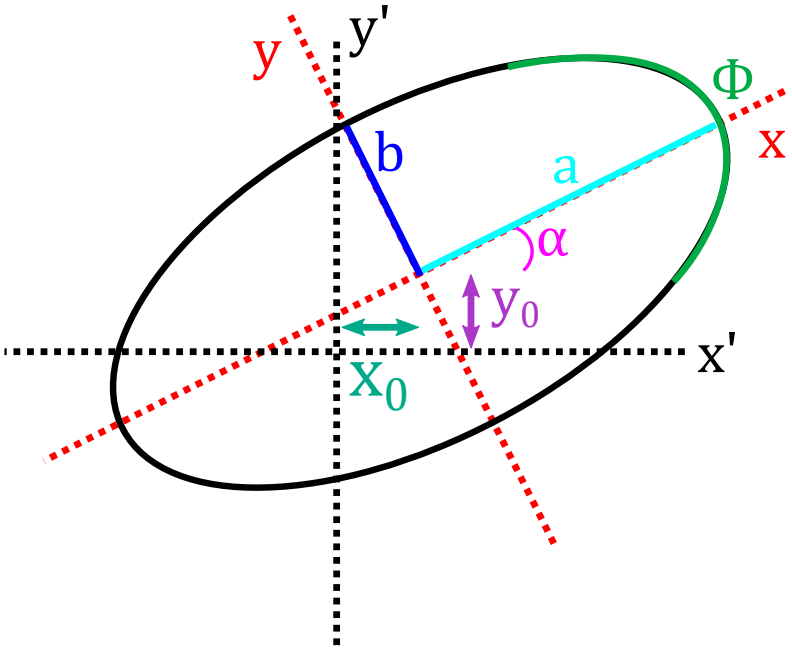

Figure 3.1: Figure showing the definitions of each of the ellipse parameters used to linearise the output of the interferometer. The following values are outputted from the ellipse fitting algorithm and are used to correct the ellipse: the gains, a, b, the offsets x_0 y_0 and the rotation α.

By taking equations 3.15 and 3.16,we apply a rotation matrix with the argument $-\alpha$, giving the

un-rotated coordinates $x' - x_0$ and $y' - y_0$, giving the equation 3.17.

$$\begin{pmatrix} x' - x_0 \\ y' - y_0 \end{pmatrix} = \begin{bmatrix} \cos(\alpha) & \sin(\alpha) \\ -\sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{pmatrix} (x' - x_0)\cos(\alpha) - (y' - y_0)\sin(\alpha) \\ (x' - x_0)\sin(\alpha) + (y' - y_0)\cos(\alpha) \end{pmatrix} \tag{3.17}$$

The offsets x_0 y_0 can be trivially removed. The final step is to scale the ellipse such that its maximum absolute value is unity. The least squares fitting technique sets the constant term of the general form of an ellipse to equal 1. We can then use this to scale the semi-major and semi-minor axes accordingly. Remembering that the fitted ellipse is of the form,

$$F = A'x_0^2 + B'x_0y_0 + C'y_0^2 - (a'b')^2. \tag{3.18}$$

The scaling factor, K is defined as $a' = \frac{a}{K}$ and $b' = \frac{b}{K}$, where a and b are the parameters returned by the least squares fitting process and is equal to,

$$K^4 = \frac{A'x_0^2 + B'x_0y_0 + C'y_0^2 - 1}{(ab)^2}. \tag{3.19}$$

Therefore, with some re-arranging, we can divide by the scaled coefficients a' and b', allowing the phase to be extracted by taking the arctangent, given by equation,

$$\phi = \arctan\left(\frac{ay'}{bx'}\right). \tag{3.20}$$

FIG 3.2 shows raw interferometric seismometer data that has been processed using the ellipse fitting technique to produce a unitary fitted Lissajous. This highlights that the fitting technique can be run successfully when measuring motion smaller than an optical wavelength.

3.3 Examples of Non-Linearities

FIG 3.3 shows an example of what non-linearities look like in the frequency domain. Here an optical inertial sensor is driven by a platform at 5 Hz with an amplitude of a few micrometers.

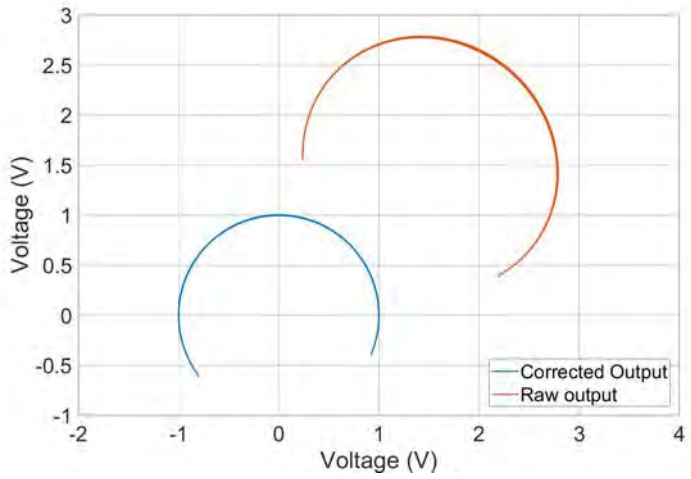

Figure 3.2: Figure showing a typical output of the ellipse fitting technique on real interferometric seismometer data. The raw trace (red) undergoes ellipse fitting and is corrected into a unitary circle (blue).

The signal is processed using the ellipse fitting routine detailed above and the non-linearities are corrected. To illustrate the problems that non-linearities can have on the readout, 10% errors were induced into each of the ellipse parameters, these being the gain, offsets and quadrature errors. The ASDs of these signals are then taken to illustrate the effect of non-linearities on the readout and are shown in FIG 3.3.

These sources look like low frequency noise which has been up-converted to higher frequencies, taking note of the higher order harmonics of the 5 Hz driving frequency, which are clearly visible at 25, 30, 35 Hz. The effect of this noise is spread out into higher frequencies and is intensified by the large driving signal. This forms a shelf-like feature in the ASD, giving rise to their name of 'non-linearity shelves'.

To investigate the effect that non-linearities will have in the presence of seismic input motion, a geophone is used to measure ground motion with two different readout mechanisms, the conventional coil-magnet readout and an non-linear interferometer which was rigidly attached to the side of the geophone can. The signal measured by the coil readout is assumed to be linear for the purposes of this test. The second readout is provided by a non-linear interferometer which is attached rigidly

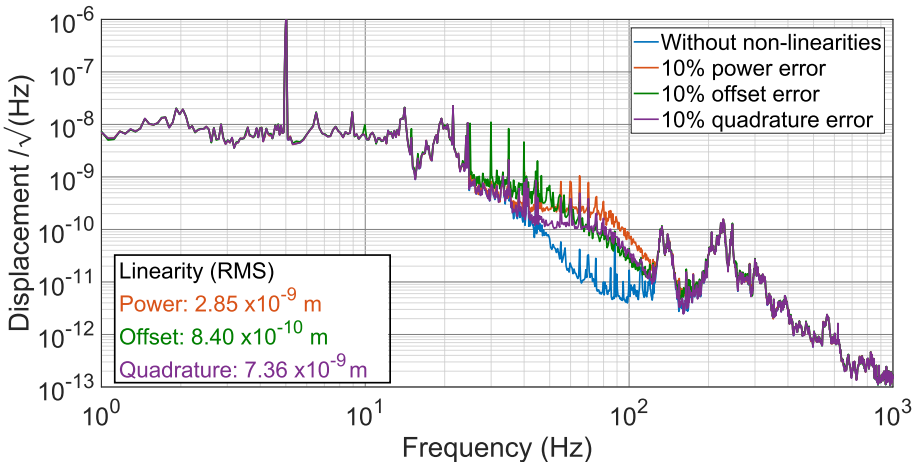

Figure 3.3: Figure showing the effects of non-linearities on the readout of the interferometer. Here, a 10% error has been injected into the ellipse fitting routine when correcting the ellipse for the three different cases: power, offsets and the angle between the two quadratures. The RMS values for each of these cases are displayed on the figure to highlight the importance of evaluating the linearity in the frequency domain.

to the outside of the geophone can, this interrogates a test mirror attached to the proof mass of the seismometer. Using the coil readout as a linear reference the measured motion between the two sensing methods can be compared. Once the signals have been calibrated into displacement, any differences that arise between the two readout mechanisms will either be caused by non-linear effects of the interferometer or opto-mechanical coupling differences between the different readout methods. Given the close, rigid alignment of the interferometer to the outer geophone case, the effect of the mechanical coupling differences is likely to be small compared to the non-linear effects that have been deliberately introduced into the interferometer.

FIG 3.4 shows a comparison between two simultaneous measurements of an optical geophone, with one of the readouts being provided by an interferometer and the other by a coil-magnet system. This was placed in a quiet lab environment and left to measure ground motion for two minutes. From this, we see that the motion measured is 1×10^{-7}m; this equates to a tenth of an optical fringe, the majority of the motion occurring at approximately 0.1 Hz. Despite the size of the input motion

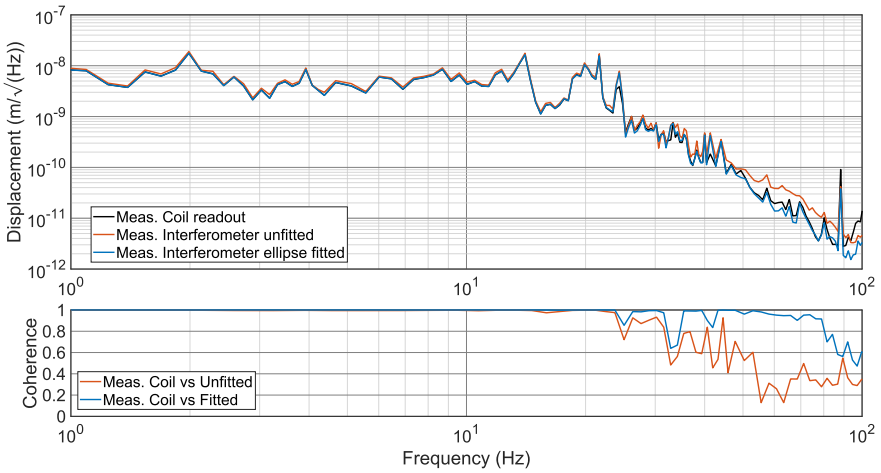

Figure 3.4: Top: Figure showing the effects of non-linearities on the readout of the interferometer in the presence of small input motion, i.e. motion that is much less than an optical fringe. The ellipse fitted curve (blue), is identical to the unfitted curve (red) but undergoes ellipse fitting before the ASD of the signal is taken and plotted. The ellipse fitted curve more closely resembles the coil readout trace, which is the input signal and is considered linear.
Bottom: The coherence i.e. the similarity between the input and the fitted (blue) and unfitted (red) traces show that the fitted signal more closely matches the input signal as it is not subject to nonlinear effects.

being small, there is a noticeable discrepancy between the measured optical and coil readouts and is predominantly caused by non-linear effects in the interferometer. This is confirmed, as when the measured interferometer data undergoes ellipse fitting (shown in yellow), the measured motion decreases to the same level as the linear coil sensor between 50 and 80 Hz.

The coherence C_{xy} between two signals, x and y is the correlation between these signals at a given frequency, f, it is defined by the equation,

$$C_{xy}(f) = \sqrt{\left(\frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)}\right)}, \tag{3.21}$$

where G_{xy} is the cross spectral density of the two signals and G_{xx} and G_{yy} are the respective

Power Spectral Density (PSD)s of the two signals.,

For perfectly coherent signals, the coherence will have a value 1 across all frequencies, whereas for incoherent signals the coherence will be equal to 0. A more rigorous example notation of the coherence can be found in the article by Gardner [92]. In the case of FIG 3.4, the coherence should be close to 1, since both the coil and the interferometer should be measuring the same motion.

In the same frequency region, when the ellipse fitting has been applied to the data, the coherence between the coil and the measured interferometer increases in the same region, indicating that the cause of this lack of correlation between the coil and interferometer was due to non-linear effects. At high frequencies above 60 Hz, the coherence drops between the two readout methods of the same proof mass on the fitted data, indicating that there is some other cause of discrepancy between the two readout methods. Moreover, by comparing both FIG 3.3 and FIG 3.4, there is a clear link between the level of non-linear error in the readout and the size of the input ground motion.

3.4 Modeling of Optical Non-Linearities

To investigate other potential sources of non-linearities, a MATLAB simulation has been developed, which incorporates many sources that can induce non-linearities in the readout, by propagating the electric field through a series of Jones matrices. Stone [88] used this technique on a simple homodyne phasemeter while de Groot [87] used this to examine how non-linearities manifest in heterodyne interferometers. The Jones matrix model includes components to account for effects such as: finite extinction ratios on the input fiber and the polarising beam splitters, the rotations of both the half and quarter waveplates as well as the retardance errors of the waveplates. The input motion for this model is customisable and can be shaped by arbitrary **ZPK** systems before the optics simulation takes place.

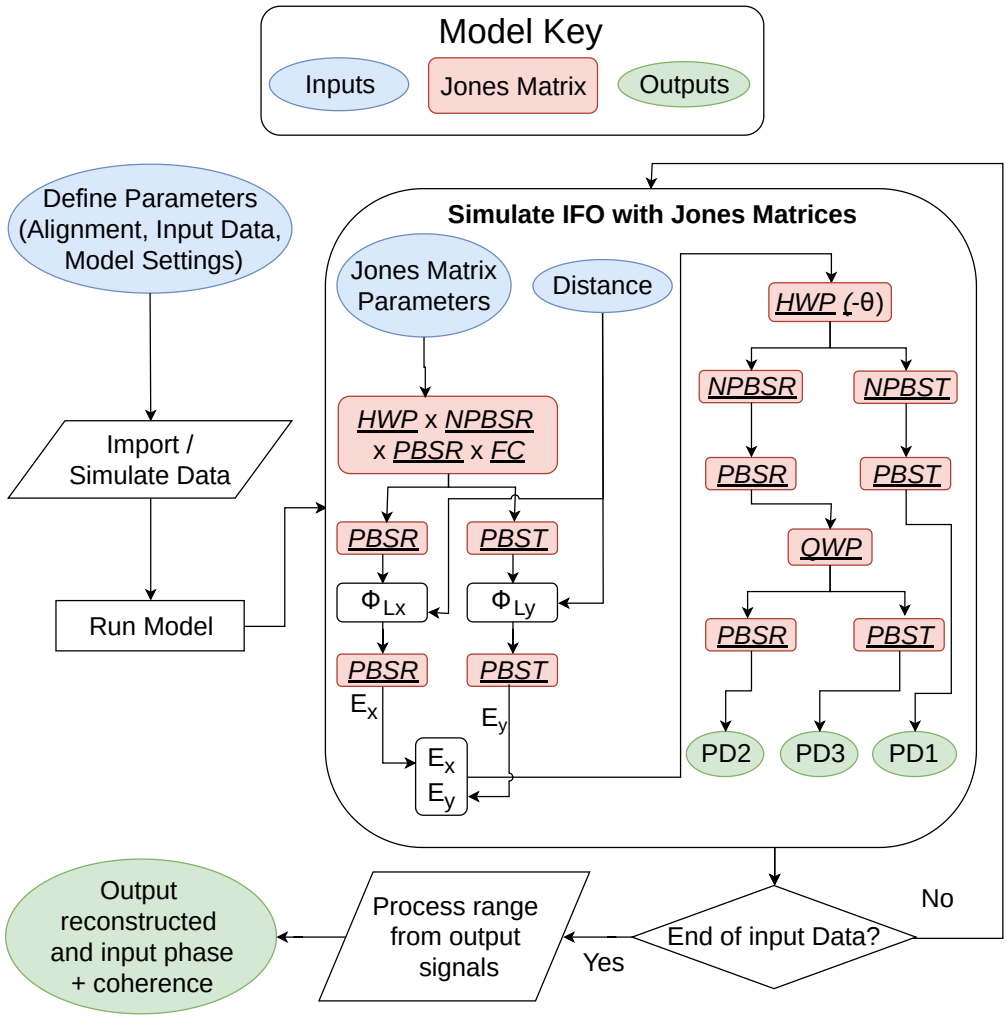

Figure 3.5: Figure describing how signals are propagated throughout the interferometer model. Initial parameters of the interferometer (shown in blue), such as waveplate rotation and beamsplitter PERs, are specified as well as vector describing the lengths of each arm of the interferometer. An input electric field is propagated through a series Jones Matrices (red) producing the outputs (green).

FIG 3.5 illustrates how signals and options are passed through the model to estimate non-linear effects on the readout. In this model an initial electric field of mixed polarisation states and a displacement vector, describing the arm lengths of the interferometer, are specified. The input electric field, whcich contains two polarisation states, is propagated through Jones Matrices, the matrices and values of which are described in Table 3.3, while the order of propagation is specified

Table 3.3: Table showing matrices used in the polarisation propagation model to simulate the effects of various optical components. The variable δ represents the retardance error in the waveplates and θ represents the angle of the fast axis of the waveplate relative to the horizontal axis.

Optic Label	Model Equation	Model Parameter
Fiber	$\sqrt{PER_{Tp}}$	$\sqrt{0.99}$
PBS T	$\begin{bmatrix} \sqrt{PER_{Tp}} & 0 \\ 0 & \sqrt{PER_{Ts}} \end{bmatrix}$	$\begin{pmatrix} \sqrt{1} & 0 \\ 0 & \sqrt{0.001} \end{pmatrix}$
PBS R	$\begin{bmatrix} \sqrt{PER_{Rp}} & 0 \\ 0 & \sqrt{PER_{Rs}} \end{bmatrix}$	$\begin{pmatrix} \sqrt{0.01} & 0 \\ 0 & \sqrt{1} \end{pmatrix}$
NPBS R	$\begin{bmatrix} r_p & 0 \\ 0 & r_s \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
NPBS T	$\begin{bmatrix} t_p & 0 \\ 0 & t_s \end{bmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
HWP	$\overbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}}^{\text{Rotation}(-\theta)} \overbrace{\begin{bmatrix} e^{\frac{i\pi+\delta}{2}} & 0 \\ 0 & e^{\frac{-i\pi+\delta}{2}} \end{bmatrix}}^{\text{Phase Delay HWP}} \overbrace{\begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}}^{\text{Rotation}(\theta)}$	$\begin{pmatrix} \theta = 45 + 2^\circ \\ \delta = 3.6^\circ \end{pmatrix}$
QWP	$\overbrace{\begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}}^{\text{Rotation}(-\theta)} \overbrace{\begin{bmatrix} e^{\frac{i\pi+\delta}{4}} & 0 \\ 0 & e^{\frac{-i\pi+\delta}{4}} \end{bmatrix}}^{\text{Phase Delay QWP}} \overbrace{\begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}}^{\text{Rotation}(\theta)}$	$\begin{pmatrix} \theta = 45 + 2^\circ \\ \delta = 3.6^\circ \end{pmatrix}$

by FIG 3.6, taken from Cooper [5]. The distance input vector is used to generate the phase ϕ_{Lx}, ϕ_{Ly}, experienced by the beams traveling along the arms of the interferometer. The beam is combined at PBS1 and PBS3 to produce the three photodiode signals, which can be processed in the same way as the measured signals. To avoid aliasing issues, the model is run at a sample frequency of 1 kHz and run for hundreds of seconds to generate enough averages at 10 mHz. Using a sample frequency of 1 kHz and max time of 500 seconds takes a dual core laptop approximately a minute to run. A drawback of this model is that it cannot be used to calculate the effect of spurious interferometers, caused by reflections of beams from AR coated surfaces, nor does it take into account of any opto-mechanical couplings in devices.

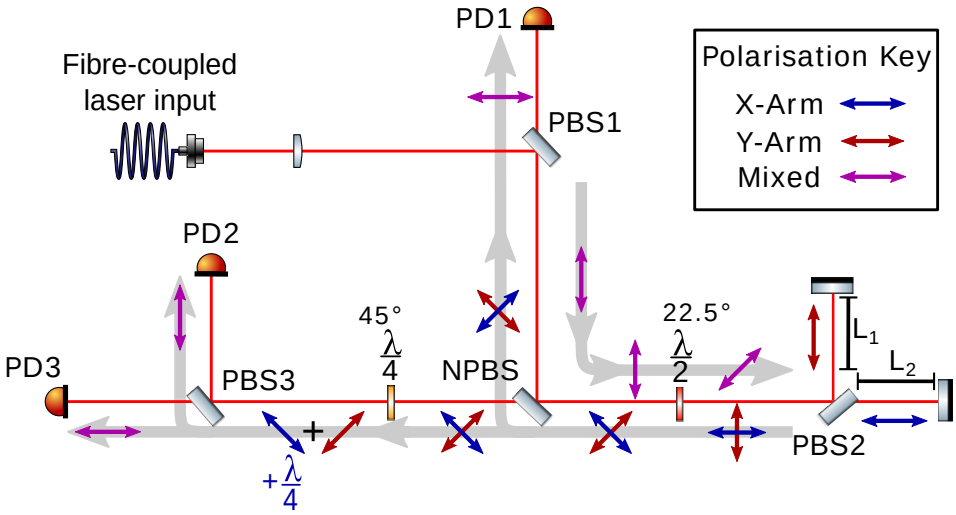

Figure 3.6: Figure showing the interferometer layout used in the MATLAB model, taken from Cooper [5].

FIG 3.7 shows a comparison between a measured inertial sensor, using both coil and interferometer readouts as well as a prediction of the motion given by the MATLAB model. The model's predicted motion is generated by using the linear coil signal as an input to the MATLAB model, which is then calibrated using the same method as the real interferometer. The values used to generate the simulation result are adjusted to re-create the measured time and frequency domain output. These are shown in Table 3.4. The RMS error in the displacement of the fitted and unfitted

interferometer, when compared with the coil readout, is 5×10^{-9}m and 1.2×10^{-8}m respectively, corresponding to a non-linearity of around 1%. When compared with other interferometers, shown by Table 3 in a review by Watchi [50], the linearity of the presented devices is comparable to others of its type. In this test the half and quarter waveplates were deliberately misaligned to inject some non-linear effects in the signal. By comparing the simulated interferometer readout with the measured interferometer readout before fitting, we see that the model is able to simulate the expected non-linearity across the frequency range of interest. As such, the model is able to confidently quantify the level of non-linearity that will be present in the measured phasemeter; the model can now be used to evaluate the effects of non-linearities on other use cases. Moreover the model can be used to identify the most sensitive components present in HoQI. In this specific case, the largest single source of non-linearity is caused by the quarter waveplate.

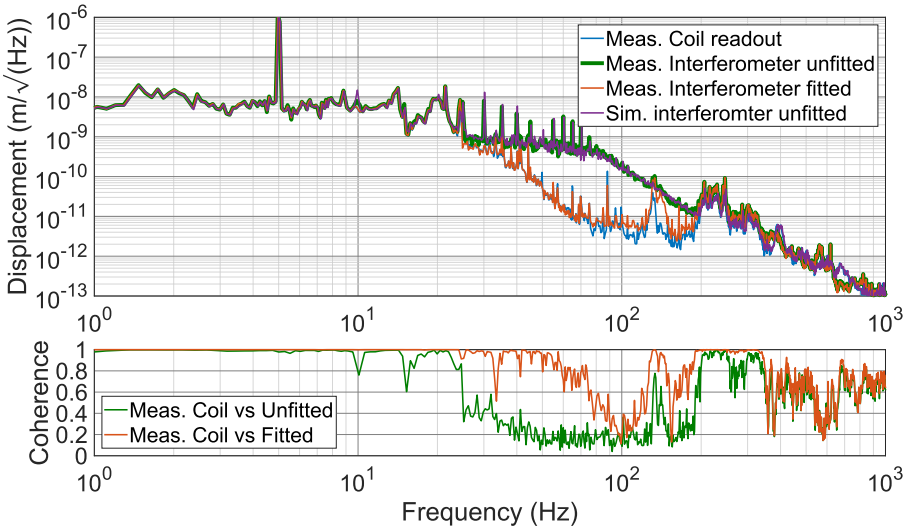

Figure 3.7: Figure showing the measured coil (blue), with the measured interferometer before and after ellipse fitting (green and red respectively) compared to the simulated interferometer that has not undergone fitting (purple), the parameters used to generate the simulated motion are shown in table 3.4. The RMS error in the displacement of the fitted and unfitted interferometer is 5×10^{-9}m and 1.2×10^{-8}m respectively.

Table 3.4: Table showing the parameters used in the polarisation propagation model to re-create the motion measured by an optical geophone.

Component	Misalignment
Input Fiber PER	1:200
Quarter Wave Plate	9.5 Degrees
Half Wave Plate	6.5 Degrees
Wave Plate Retardance Error	3.6°
PBS1 Transmission PER	1:1000
PBS 1 Reflection PER	1:50
PBS2 Transmission PER	1:1000
PBS2 Reflection PER	1:50
PBS3 Transmission PER	1:1000
PBS3 Reflection PER	1:50
NPBS Reflection (S:P)	50:50
NPBS Transmission (S:P)	50:50

3.5 Effects of Non-Linearities on Applications of HoQI

Now that the MATLAB model has been verified against real data, we can now examine the effect of non-linearity on two applications of the interferometers in question: as a readout of a high Q seismometer, and as a readout mechanism for high Q suspension systems, such as those in use in LIGO.

3.5.1 Interferometric Readout of High Q Seismometer

Due to their superior noise performance, interferometers can be used as a readout mechanism for geophones [5] and seismometers [72, 73, 93], resulting in a factor 100 increase in resolution at 10 mHz when compared to a geophone with a coil-magnet readout. Such inertial sensors have a Q of around 1, in the case of the L-4C. In the absence of large errors in the interferometer fitting parameters, the low Q of the geophone spring limits the impact of non-linear effects spoiling the readout at high frequency at the expense of limiting the resolution of the sensor at low frequency, due to the thermal noise contribution. Inertial sensors, such as those proposed by Cooper [5], are limited by suspension thermal noise between 10 mHz and 2 Hz; this can be reduced by increasing the Q of the inertial sensor. For this application, we have applied a filter modeling the effects of measuring ground motion using an interferometer connected to inertial mass with a resonance frequency of 1 Hz and a Q of 1000. This was achieved by filtering the input data by a ZPK system that describes the transfer function between the ground and the suspended mass.

The parameters used in the simulations presented in this section are given by Table 3.5.1. The numbers chosen for this simulation represent realistic values for polarisation extinction ratios and typical errors in alignment when constructing the devices. For the high Q inertial sensor, the model was run at a sampling frequency of 2000 Hz for a total of 1000 seconds. The result of this is shown in FIG 3.8.

FIG 3.8 shows the simulated motion as measured by an interferometer when placed on a high Q geophone, using optics parameters given by Table 3.5.1. Below its resonance frequency of 1 Hz, the

Table 3.5: Table showing the parameters used in the polarisation propagation model.

Component	Misalignment
Input Fiber PER	1:200
Quarter Wave Plate	2 Degrees
Half Wave Plate	2 Degrees
Wave Plate Retardance Error	3.6°
PBS1 Transmission PER	1:1000
PBS 1 Reflection PER	1:50
PBS2 Transmission PER	1:1000
PBS2 Reflection PER	1:50
PBS3 Transmission PER	1:1000
PBS3 Reflection PER	1:50
NPBS Reflection (S:P)	50:50
NPBS Transmission (S:P)	50:50

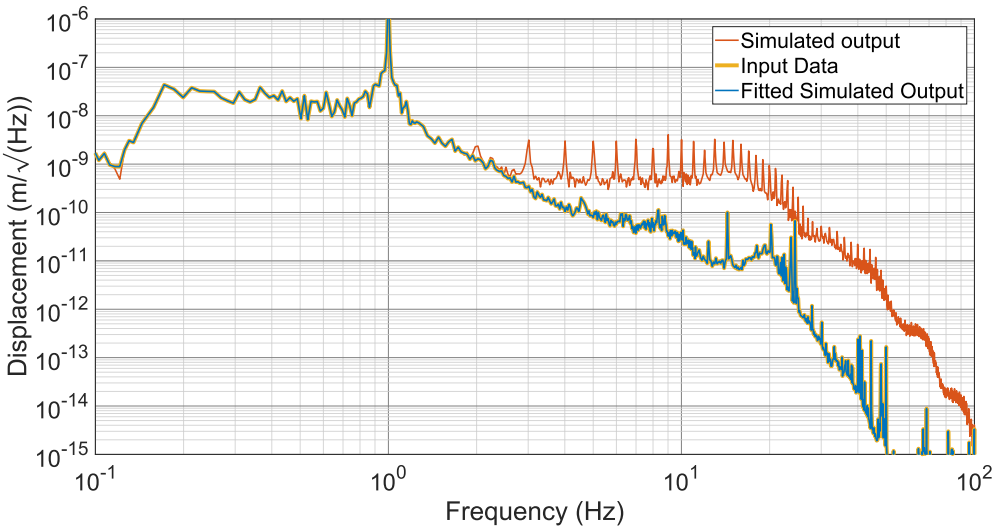

Figure 3.8: Figure showing the predicted motion measured by an interferometer when placed on a high Q geophone, with a resonance frequency of 1 Hz and a Q of 1000, compared against the true input motion (red). The blue trace represents the simulated motion as measured by the interferometer before any ellipse fitting is applied. The true input motion is obscured by the simulated and fitted interferometer measured motion (yellow), showing that as long as fitting is applied, the true motion can be extracted from the interferometer without non-linear effects spoiling the resolution.

predicted interferometer motion is identical to that of the input motion. Above 2 Hz the simulated motion is dominated by the non linear noise couplings forming a 'shelf' that spreads low frequency motion into higher frequencies. Higher order resonances of the 1 Hz peak can clearly be seen at multiples of the resonance frequency that gradually decay in amplitude. Encouragingly, these non-linear effects can be corrected for by running ellipse fitting to the point where the impact of non-linear effects no longer spoil the expected motion measured by the sensor. At 30 Hz higher order non-linearity shelves can be seen, caused by higher order terms in the Fourier series that describe the non-linear effects, further spoiling high frequency resolution without additional fitting.

3.5.2 Interferometers on High Q Suspensions

Interferometers, such as those described previously, can be used as position sensors in LIGO's quadruple suspension system. To do this, the quadruple suspension model, as detailed in [94], is used to generate transfer functions between the suspension point (ST2) length and the UIM, where the interferometers are proposed to be located. An estimate of the ST2 suspension point motion is modified by the suspension transfer function and used as an input into the polarisation model to simulate the effects of non-linearities in this use case. Taking into account the locations of the test and reference reflectors on the suspension chain, the differential motion between the two mirrors will be given by $\Delta_X = X_{\mathrm{UIM}} - X_{\mathrm{ST2}}$. In order to damp the resonances from the UIM, the sensors, and therefore the non-linear contributions to the noise performance, need to be below $10^{-13}\mathrm{m}/\sqrt{\mathrm{Hz}}$ at 10 Hz to be below the expected stage 2 suspension point length.

FIG 3.9 shows the projected stage 2 motion through the quad model as input to the interferometer model, described in section 3.4 shown in purple. The signal, as seen by the interferometer readout is shown in red, while the ellipse corrected motion is shown in blue, which is hidden behind the input motion trace. Without the ellipse fitting correction, the sensor's measured signal is dominated by non-linearities in the 0.4-20 Hz region, caused by up-conversion of the suspension peaks. The lower subplot shows the coherence of the signal before and after the ellipse correction. With this we see a significant loss in coherence between the real and measured motion. After fitting, we see the coherence increases to 1 across the whole frequency range.

During observation time, the suspension resonances are damped to lower the RMS motion to allow for easier control of the numerous cavities. FIG 3.10 shows the same simulation as FIG 3.9 but with the suspensions damped using the damping provided by the current suspension sensors and current damping filters. With this level of damping, the impact of non-linearities can only be seen at around 1.5 Hz, owing to the greatly reduced amplitudes of the suspension resonances. When the ellipse fitting correction is activated, the effects of non-linearity on the readout are all but eliminated.

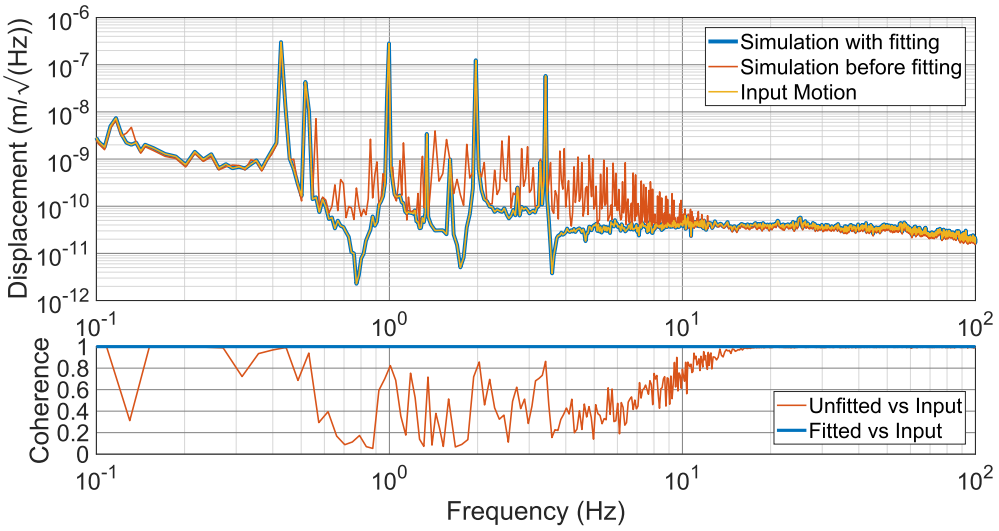

Figure 3.9: Figure showing the effects of non-linearities with damping of the suspension resonances turned off. Top: Shows the amplitude spectral density of the input motion (yellow), with the motion measured by the interferometer before and after ellipse fitting correction, shown in red and blue respectively. Bottom: Shows the coherence between the input signal and the interferometer signal before and after ellipse fitting, shown in red and blue respectively.

In both cases, this represents the best case scenario for the coherence measured by these devices. Here we have not taken into account sensor noise, loss of coherence due to mechanical noise sources, or spurious interferometers present in the readout, in order to simplify the simulation. In reality, the corrected signals would lose coherence where the signal becomes sensor noise limited.

3.6 Summary

Non-linearities are frequency and amplitude dependent errors that can be present in long range interferometers such as, but not limited to, HoQI. These are caused by imperfections in the optics used in the interferometers, which alters the polarisation that is propagating through the interferometer, leading to errors in the phase unwrapping of the three quadrature signals. The result of this is that, in the presence of large motion, the signal at low frequency is up-converted to higher

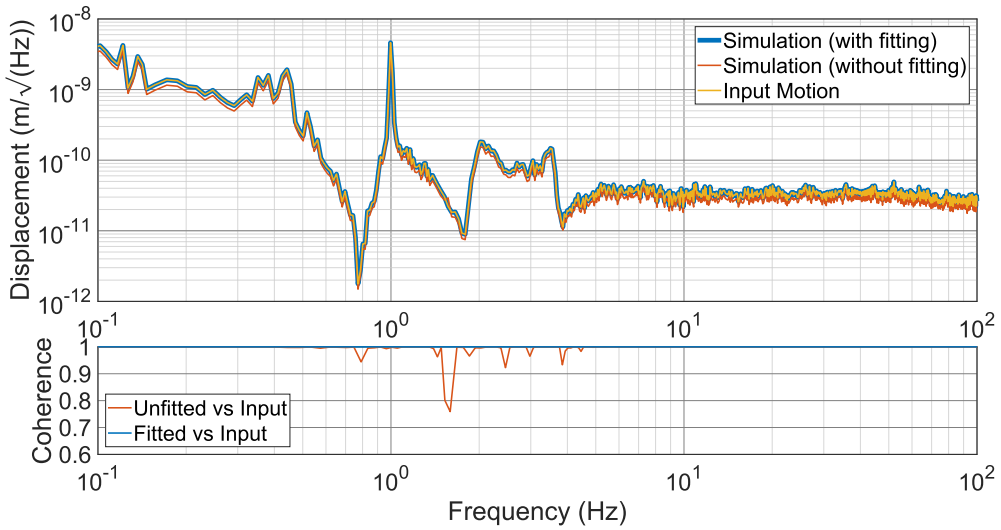

Figure 3.10: Figure showing the effects of non-linearities with damping of the suspension resonances turned on by looking at the simulated signals that would be measured by interferometers. Top: Plot of the ASD of the input motion (yellow), with the motion measured by the interferometer before and after ellipse fitting, shown in red and blue respectively. Bottom: A plot of the coherence between the input signal and the interferometer signal before and after ellipse fitting, shown in red and blue respectively.

frequencies spoiling the resolution of the proposed sensor.

A method of reducing the effect of these non-linearities has been evaluated and applied to both simulated and real data, demonstrating that these non-linear effects can be quantified and corrected for. This correction, known as ellipse fitting, can be applied in post processing to the measured data and eliminates the 'fake' motion caused by the non-linear effects. This increases the coherence between motion measured by a linear reference and the intentionally non-linear interferometer.

To investigate this, a MATLAB model that propagates an electric field through a series of Jones Matrices describing the interferometer has been constructed. This model allows for the level of non-linearities to be accurately quantified given some arbitrary input motion. This not only allows for the effect of non-linearities to be quantified when evaluating new components, but also allows applications of interferometers like HoQI to be evaluated to see whether non-linearities will

negatively effect the readout. In the two cases studied, a high Q seismometer and as a readout for high Q suspensions, non-linear effects are present but can be adequately corrected by ellipse fitting, restoring the linearity and the resolution of HoQI.

Chapter 4

Interferometric Inertial Sensors

Why do we fall? So that we can learn to pick ourselves up.

Thomas Wayne

This chapter describes the construction of a compact optical inertial sensor using the interferometer described in chapter 2. Parts of the chapter reference a paper, with a preliminary title of 'Development of a near suspension thermal noise limited geophone', which is in pre-publication at the time of writing. I am the lead author and have designed and constructed the devices described here. Text and figures from the manuscript have been copied verbatim and expanded on for increased clarity and further discussion.

To isolate the test masses from ground motion, LIGO uses a complex system of passive [39, 78] and active [69] isolation. The active isolation is largely provided by the ISI, and reduces the inertial motion at frequencies from 0.1 to 10 Hz. The ISIs are suspended on springs from the ground, providing passive isolation above their resonance frequencies at approximately 1 Hz. Despite the success of these systems, the lowest frequencies in the LIGO detection band, 10-20 Hz, are still limited by technical noises that are driven by residual motion at even lower frequencies [48, 49]. The residual motion must be reduced in order to detect gravitational waves below 10 Hz [47].

The ISIs are sensed by a mixture of displacement sensors, geophones, and force feedback seismome-

ters. The sensors are 'blended' or 'fused' together to use the 'best' information in each frequency band. Geophones, such as the Sercel L-4C and Geotech GS-13 are used throughout the detector to measure the motion of both the ground and the isolated platforms. While the performance of force feedback seismometers is inherently superior to that of geophones at frequencies below 0.5 Hz, they are considerably more expensive. As such, only the most critical isolation platforms, the BSCs, which house the primary test-masses, and the beamsplitter use these broadband seismometers. Other chambers that house the auxiliary optics must rely on geophones to provide inertial platform measurements.

Coil-magnet geophones are limited by their intrinsic readout noise, caused by the Johnson noise in the geophone's coil. The noise budget of an L-4C is shown in FIG 4 of Kirchofff et al. [95]. The other main noise source, suspension thermal noise, is around a factor of 200 lower at 10 mHz. With a sufficiently quiet readout method, the resolution of these sensors can be substantially improved at low frequencies. In the past interferometers have been used to decrease the readout noise contribution in broadband seismometers [72, 73] and thus increase the resolution of these devices. In an earlier paper [5], we demonstrated sufficient readout performance to reach the suspension thermal noise from 10 mHz to 2 Hz of an L-4C geophone, should no additional noise couplings arise.

Improvements to the resolution of geophones will in turn allow for a re-design of the ISI's control loops, reducing sensor noise injection, as shown by chapter 5 of this thesis. In particular, they will allow for much improved inertial isolation between 0.1 and 0.3 Hz while maintaining the same RMS velocity of the platform.

4.1 Principle of Inertial Sensors

Inertial sensors are sensors that rely on inertial mass, often called a proof mass, to make measurements of displacement, velocity or the acceleration relative to this mass. The simplest example of this is a mass on a spring, depicted in FIG 4.1. An inertial sensor will measure different variables depending on whether the oscillation is below or above the resonance frequency. Below the reso-

nance frequency, the suspended mass will follow the ground motion and thus the inertial sensor will only be sensitive to the acceleration between the ground and the suspended mass. Above the resonant frequency, in the presence of no external forces, the sensor will directly measure displacement relative to the proof mass.

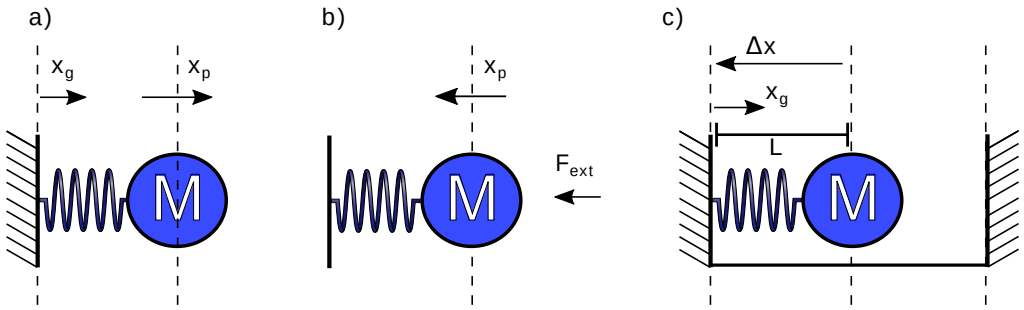

Figure 4.1: Diagrams of the inertial sensor for each of the transfer functions described. a) shows ground to platform coupling, b) shows how external forces affect the proof mass and c) shows how a measurement made by the inertial sensor is converted into ground motion.

To determine how ground motion couples into motion of the platform or mass motion we need to define a number of transfer functions. A transfer function is a frequency domain function that maps an input of a system to an output. Working in the frequency domain allows for easy switching between acceleration, velocity and displacement by multiplying by a factor of $i\omega$ for every differentiation. Every time an integration is performed the term is multiplied by a factor of $\frac{1}{i\omega}$. There are three transfer functions that are important in inertial sensors, these are:

1. $x_g \rightarrow x_p$, this describes how the ground motion couples into platform motion. At frequencies below the resonance x_g maps directly to x_p. Above the resonant frequency, x_p is isolated to a factor of $\frac{\omega_0}{\omega^2}$, and at frequencies above Qf_0, the response scales with $\frac{1}{\omega}$. Here ω_0 is the angular resonant frequency, k is the spring constant, m is the mass and Q is the quality factor of the spring.

$$\frac{x_p}{x_g} = \frac{\frac{i\omega\omega_0}{Q} + \omega_0^2}{\left(-\omega^2 + \frac{i\omega\omega_0}{Q} + \omega_0^2\right)} \tag{4.1}$$

2. $F_{ext} \rightarrow x_p$, describes how any external forces couple into the platform motion. This is useful for mapping, for example, how forces applied by an actuator will affect platform motion, or seeing how ground motion is suppressed in passive isolation. This has a non zero DC value of $\frac{1}{k}$, where k is the spring constant.

$$\frac{x_p}{F_{ext}} = \frac{1}{m(-\omega^2 + \frac{i\omega_0\omega}{Q} + \omega_0^2)} \tag{4.2}$$

3. $\Delta X \rightarrow x_g$, describes how a measurement of ΔX can be converted into inertial equivalent displacement. The inverse of this is known as plant inversion. This will be used to describe how ΔX; measured by different readout mechanisms, translates into ground motion. There is a constant offset, caused by the length of the spring, this is denoted by the length L and does not feature in the equation 4.3. Graphs of these transfer functions are shown in FIG 4.2.

   Different types of inertial sensor include accelerometers, seismometers and geophones. The rest of this thesis will focus on the latter two of these devices, as they are used in the detectors which measure ground and platform motion respectively, on both the Hydraulic External Pre-Isolator (HEPI) and ISI stages. As shown by chapter 5 the ISIs are limited by inertial sensor noise, either directly, as is the case below 100 mHz or indirectly as is the case with the case above 1 Hz. In this case sensor blending must occur at higher than ideal frequencies to prevent excess inertial sensor noise coupling in at low frequency.

$$\frac{x_g}{\Delta x} = \frac{\left(-\omega^2 + \frac{i\omega\omega_0}{Q} + \omega_0^2\right)}{\omega^2} \tag{4.3}$$

Equation 4.3 reveals the problem behind creating a sensitive inertial sensor, at frequencies below the resonance, the **apparent motion measured by the** inertial sensor becomes unbounded in

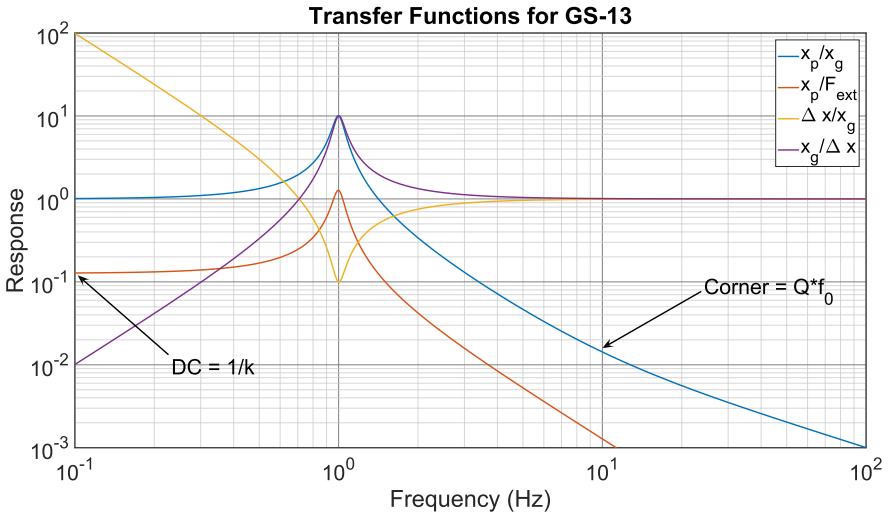

Figure 4.2: Here are the different transfer functions showing the frequency response of different systems, relating ground to platform motion ($\frac{x_p}{x_g}$), external forces to platform motion ($\frac{x_p}{F_{ext}}$), ground motion to distance measured by the sensor ($\frac{\Delta x}{x_g}$), sensor responses to ground motion ($\frac{x_g}{\Delta x}$)

displacement. This can be seen by equation 4.2 and FIG 4.2.

4.2 Noise Sources

The main noise sources for the interferometric L-4C are suspension thermal noise and readout noise. The suspension thermal noise is calculated assuming a resonant frequency of 1 Hz, a quality factor of 1 and a mass of 1 kg using equations given by Saulson [96]. The readout noise used in this chapter is the measured interferometer noise taken from FIG 2.9 from chapter 2. This contains the electronic noise of the readout electronics and the frequency noise as measured by the interferometer, though the path length difference has been minimised to reduce its impact on the total readout noise curve. The readout noise trace was chosen as it presents a 'worst-case' measurement for the intrinsic noise of the interferometer, as this was not tested in vacuum with seismic isolation. The lack of isolation may have resulted in the interferometer described in chapter 2 measuring noise sources that might not be present in future tests, e.g. where seismic isolation is used, or when the device is placed in a vacuum chamber. The readout noise has several

contributions, but the dominant two are electronic noise (or dark noise) and frequency noise, caused by laser fluctuations. The frequency noise is calculated by the equation,

$$\delta L = \frac{L\delta f}{f_0}, \tag{4.4}$$

where L is the length mismatch between the two arms, f_0 is the (optical) carrier frequency of the laser and δf is the laser frequency noise. Since frequency noise is a form of readout noise, it also passes through the plant inversion filter in equation 4.3. To reduce the input laser frequency fluctuations we have used a 1064 nm solid-state Innolight Mephisto 500NE laser which has frequency noise of $10^5 \frac{Hz}{\sqrt{Hz}}$ at 0.1 Hz. Our readout system had an arm length mismatch of 6 mm in the initial prototype.

The electronic noise is taken from FIG 3 in [5] and was taken by injecting a constant current that was generated by putting a low-noise resistor in series with a very stable voltage reference. The readout noise trace is based on the sensitivity curve presented in [5], which was measured with a much smaller path length difference of 0.7 mm. The frequency noise contributions in the readout noise measured in [5] are assumed to be small, owing to the significantly reduced arm length mismatch. A noise budget highlighting the dominant noise sources of an interferometric L-4C is shown in FIG 4.3. This shows that at 10 mHz the optical L-4C should be a factor of 100 more sensitive than the coil readout L-4C.

Below 60 mHz the device will be limited by frequency noise coupling, as this has a f^{-3} dependency below the resonant frequency of the L-4C. In the future this can be reduced by either minimising the length mismatch between the two arms, or stabilizing the laser frequency with a second interferometer with significantly different arm lengths. At 10 mHz frequency noise is only a factor of two higher than the suspension thermal noise, so only a small change to the frequency noise coupling will be required. Suspension thermal noise limits the resolution of the device between 60 mHz and 2 Hz. At frequencies above 2 Hz frequency noise and readout noise are the only significant sources of noise. In reality, it will be challenging to measure the self noise of the sensor in

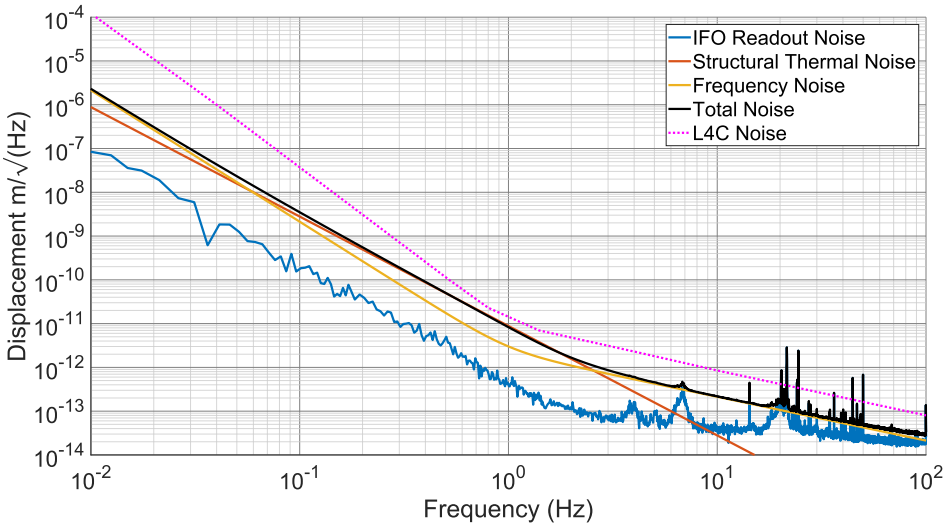

Figure 4.3: A noise budget of the optical inertial sensor. The interferometer readout noise (blue), structural thermal noise (red) and frequency noise (yellow) are summed in quadrature to produce the expected resolution of the optical L-4C (black). For comparison the resolution of an L-4C with coil readout is shown in magenta.

this frequency range due to large ground motion. The effect of this ground motion can be reduced through vibration isolation and 'huddle' testing [95].

4.3 Development of an Interferometric Inertial Sensor

The next section details the development of the first three prototypes of optical inertial sensors using HoQI as a readout mechanism for an L-4C.

4.3.1 Initial Prototype

Previous iterations of HoQI relied on having angle adjustment, for the laser beam alignment in two degrees of freedom as well as position adjustment in the other, to ensure the beams were optimally overlapped and to reduce frequency noise respectively. As a result, the first design incorporated both of these alignment tools into the design, which was broadly based off the first compact design

presented in section 2.4.1 in chapter 2.

The optical baseplate used in this design was slightly larger than the first measuring 14 cm × 14 cm, though it retained the use of mounted half inch optics for ease of tuning. This would be of particular interest to be able to easily change the rotation of each of the wave plates to examine the effects of non-linearities on the readout. One improvement made from this was to mill into the baseplate by 3 mm to designate the location of each of the beam-splitters and their corresponding mounting cubes, this ensured that the beamsplitters could be glued in the exact position and rotation angle for optimal alignment. The large operating range that is required for seismic measurements will amplify any translational offsets caused by angular misalignments of the optics.

To try and maximize the coherence between the measured motion using the coil readout and the interferometer readout, both the optical baseplate and the geophone can were mounted onto a single larger baseplate. To interface the interferometer with the geophone can, the L-4C was carefully disassembled by drilling out the retaining pins near the top of the geophone can. With these pins removed, a specially designed jig was used to separate the geophone can from the internals, including the coil and proof mass. This process is shown in FIG 4.4. With the can removed a hole could safely be milled into the side of the can, large enough to fit a custom half inch mirror mount to the side of the proof mass, without damaging the internals of the geophone itself. This had to be a sufficient distance away from the side of the can to overlap with a 45 degree steering mirror which aligned the beam onto the horizontally glued mirror and corrected for any angular misalignments when attaching the mount to the proof mass. An image of the mirror mount, as well as the overall assembly can be found in FIG 4.5.

The L-4C has a peak to peak travel of 6.25 mm, a resonance frequency of 1 Hz and a moving mass of 1 Kg [97]. To give the sensor the largest range, the resting point of the proof mass before adding any additional mass is assumed to be half of the peak to peak travel. This equates to a spring extension of 3.125 mm and a spring constant of $39.5\frac{\mathrm{N}}{\mathrm{m}}$. By adding the mirror assembly to the moving mass, the spring will extend further. By using the same logic, the maximum permitted extension for the spring will be a quarter of the total peak to peak motion (1.6 mm). Using this

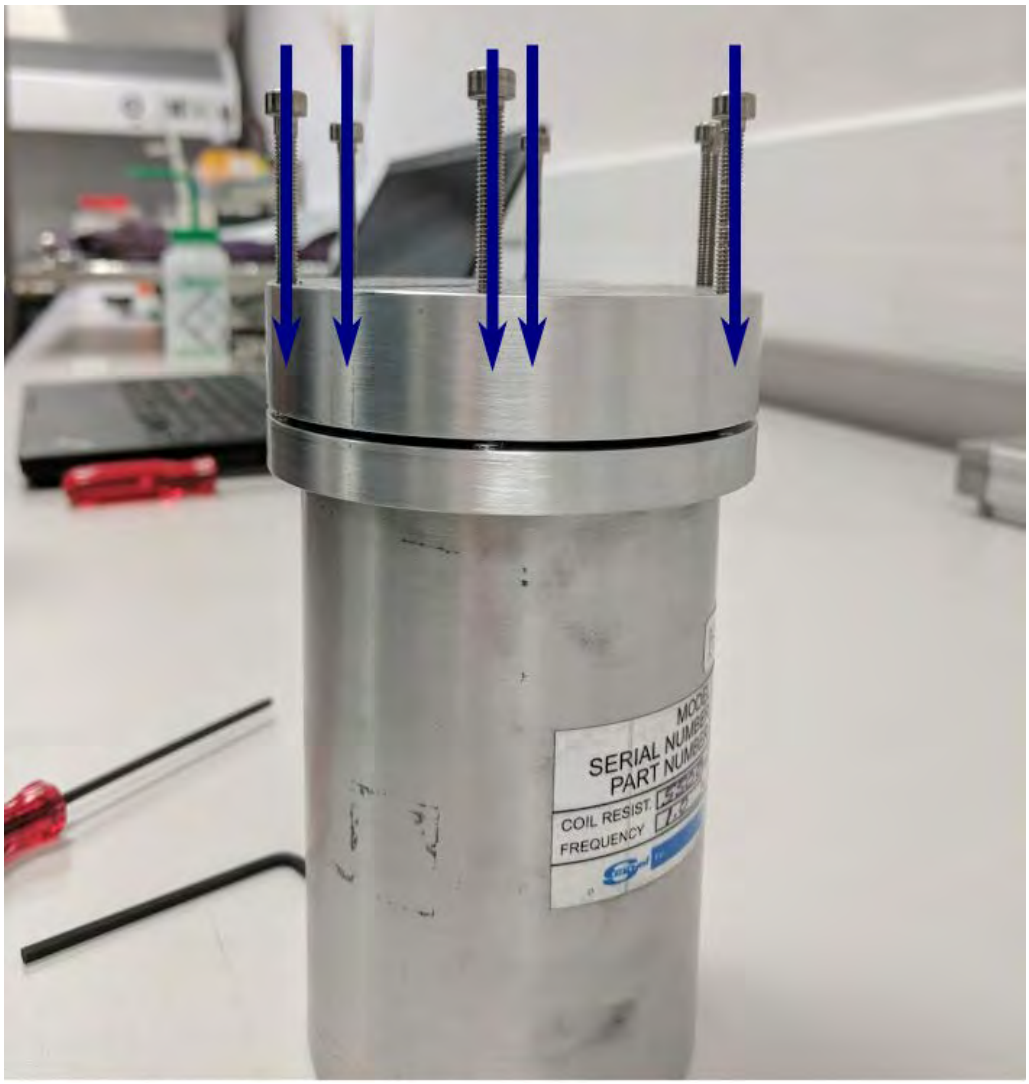

Figure 4.4: Image of the process of removing the outer can from a Sercel L-4C. To remove the can force is applied, shown by the blue arrows, by turning screws that press on the outer can of the geophone. Each screw is turned gradually by half a turn before moving onto the next screw and repeating, ensuring the can is removed evenly around the inner proof mass.

information, the new maximum mass of the mirror assembly can be calculated by using equation 4.6.

$$F = kx \tag{4.5}$$

$$\omega_0 = \sqrt{\frac{k}{m}} \tag{4.6}$$

Substituting numbers into this equation, gives a maximum mass for the whole assembly of 6.28 g. The mirror used in these tests is a Thorlabs BB05-E03 and has a mass of 2 g, meaning the rest of the mount on the mirror assembly needs to be less than 4.28 g. To achieve this mass requirement the mount was light-weighted by drilling holes, chamfering edges and reducing the thickness of the mount, the final mount weighs just 2.5 g, well within the mass requirements.

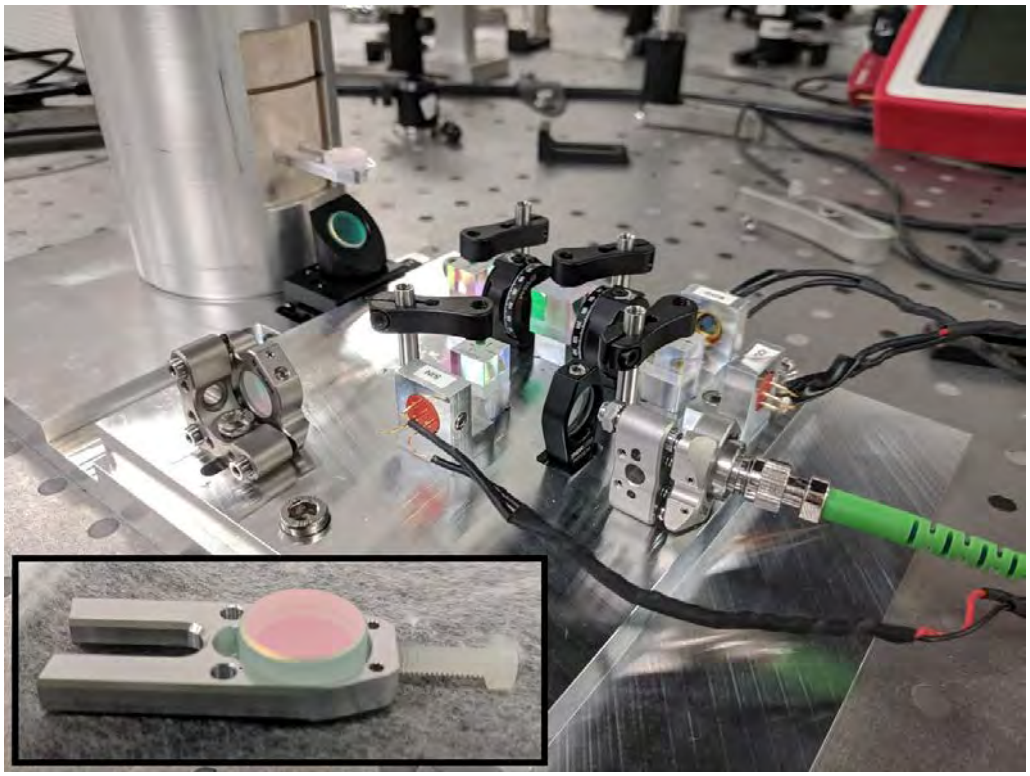

Figure 4.5: Image showing the initial first prototype of the optical inertial sensor.

To evaluate the performance of the inertial sensor, one was aligned with 5 mW of 1064 nm light which was coupled through a Polarisation Maintaining (PM) fibre; which has an extinction ratio of 200:1. The wave plates were aligned to maximise the circularity of the ellipse. The alignment is achieved by adjusting the fibre coupler such that the laser beam is reflected from the center of the test arm's mirror. The accurate placement of the beamsplitters and photodiodes results in the test arm beam being centered on each of the three photodiodes. The reference arm is then aligned to ensure optimum overlap of the two beams to maximise the usage of the ADC. Two other geophones were placed as physically close to the optical L-4C as possible to maximise the coherence between

the two. A photo of the setup is shown in FIG 4.6.

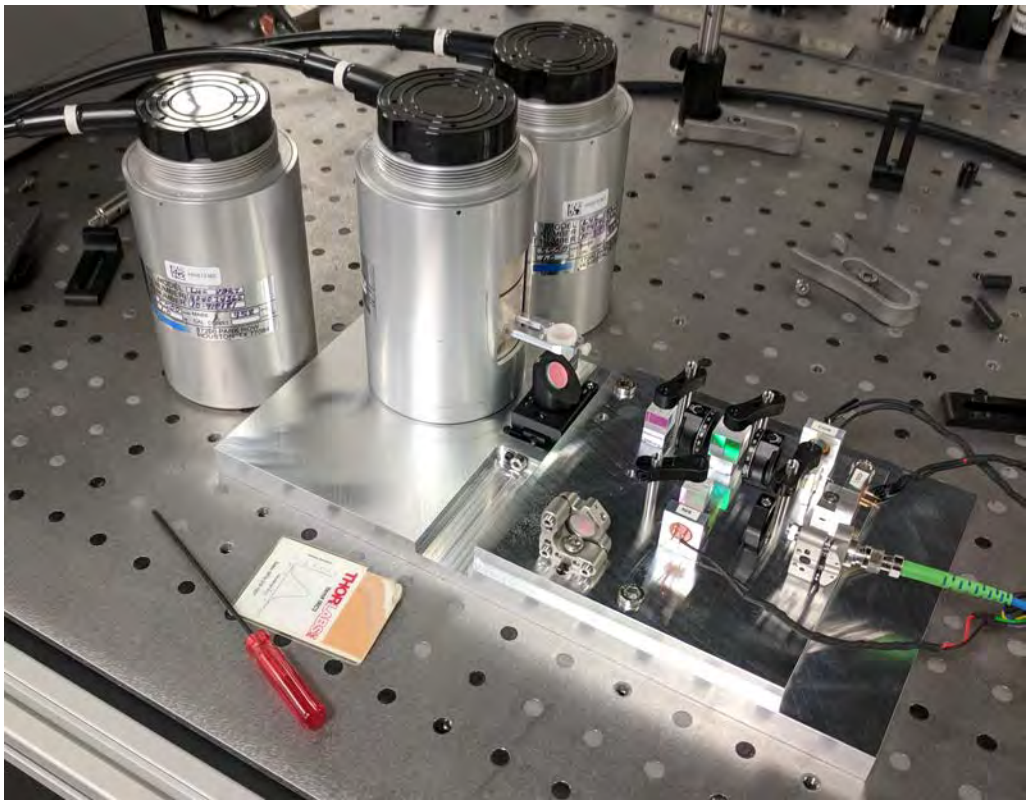

Figure 4.6: Image showing a photo of the initial huddle test. All three L-4C coil outputs are measured simultaneously along with the optical measurement of one of the L-4Cs.

The results of this huddle test are shown in FIG 4.7, in this test the inertial sensors are placed on the optical bench inside a box, with small rubber feet acting as the only form of seismic isolation. While the box will act to shield the sensor from some air current and temperature fluctuations, it will not be perfect. This lack of isolation will leave the sensors subject to air current and temperature fluctuations at frequencies below 0.1 Hz, while at high frequency the sensors will measure the residual ground motion through the rubber pads and any vibration from residual acoustic noise. Improvements to the isolation of the box will also reduce acoustic noise couplings above 10 Hz.

Coherent information between the optical and coil readout L-4Cs is optimally removed using the method described in [98], as previously applied in [95], leaving only the incoherent signal, shown by

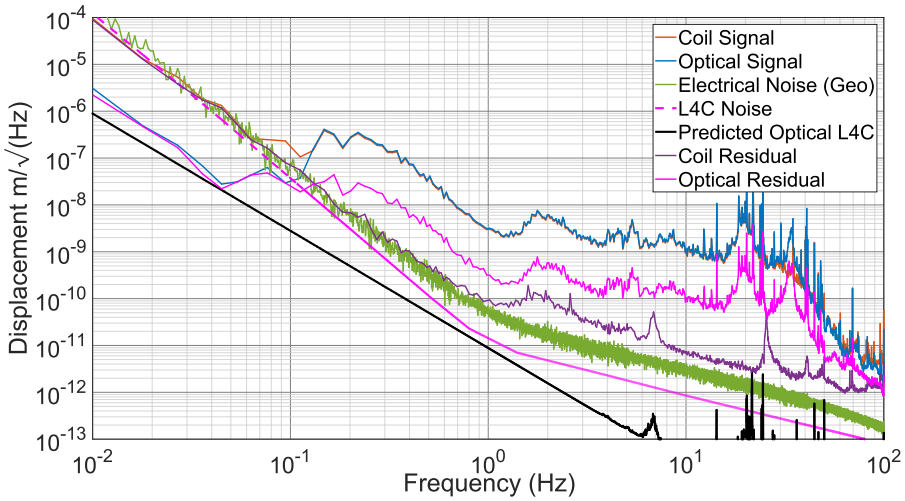

Figure 4.7: Figure showing the results from the first huddle test measurement, here the measured signals from the optical inertial sensor (blue) and the signal measured by the coil based geophone (red). These are plotted against the measured electronic noise (green), the expected optical L-4C resolution (black) and the coil L-4C resolution (dashed magenta). The residuals for the interferometer (magenta) and the coil (purple), show the residual motion measured by the two readout motion. All signals have been plant inverted to account for the dynamics of the L-4C.

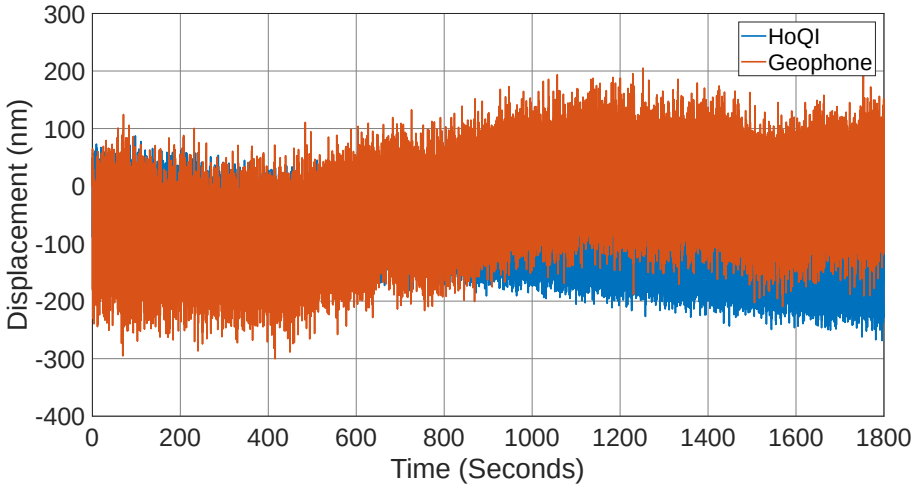

Figure 4.8: Figure showing the long term stability of the optical inertial sensor when compared to the geophone using the data collected in the initial huddle test shown in FIG 4.7.

the magenta trace in FIG 4.7. The same method can be repeated for the coil geophones, this time only using two geophones in the coherence-removal process shown by the purple trace in FIG 4.7. The electronic noise and the coil L-4C sensitivity curves are shown by the green and dashed magenta curves, which represent the best subtraction possible using the coherent subtraction method as they are the fundamental limits of the L-4C when tested. The electronics used for this measurement were modified EUCLID [54] electronic boxes. For interest, the calibrated signals, as measured by the optical and coil L-4C are also displayed in blue and red respectively. In the ideal case, this residual will be the sum of the fundamental noises in the system: the readout noise, thermal noise and frequency noise. In reality the device will likely still measure residual ground motion that is incoherent between the devices. A time series plot of the huddle test is shown in FIG 4.8 and does not include the response of the sensor as described in equation 4.3.

This test shows that at 10 mHz the measured resolution of the optical inertial sensor is only a factor of 3 above the expected optical L-4C noise, which in part is limited by the thermal noise of the spring of the geophone, and hits the suspension thermal noise of the proof mass at 40 mHz. The coherent subtraction, or the residual motion, between coil geophones sits on a combination of the coil geophones and the electronic noise between 100 mHz and 1 Hz, whereas the optical residual diverges and is a factor of 5 higher than the coil residual. This indicates that the coherence between the geophones in this frequency ranges is insufficient to subtract the motion. This lack of coherence is also present at higher frequencies, hinting that there is some additional source of motion or noise coupling in to the optical readout, as non-linearities in the readout have been eliminated by the ellipse fitting method described in chapter 3. The coil residual is within a factor of three of the combined electronic and L-4C self noise with two devices subtracting from the main measurement devices. As shown by Kirchhoff [95] the subtraction increases with the inclusion of more devices. Moreover, with higher resolution electronics and improved seismic isolation, the noise floor of the coil and optical L-4C noise should be exposed.

The result of the initial test is that the optical L-4C behaves as expected at frequencies below 100 mHz, and is likely limited by suspension thermal and frequency noise. To improve the resolution

of the sensor above this frequency, the coherence needs to be improved between the two inertial sensor readout methods to subtract the ground motion.

4.3.2 MK2 Optical Inertial Sensor

A clear problem with the first prototype was the lack of coherence between the optical and coil read out mechanisms, so a number of changes were made to the construction of the interferometer baseplate and integration with the L-4C to address this issue. The first modification that was made was reducing the size of the sensor substantially and bolting the interferometer baseplate to the geophone can directly. To accommodate this, the baseplate was reduced in size from $14 \times 14\,\text{cm}$ to $6 \times 7\,\text{cm}$, a picture of this is shown in FIG 4.9

The fibre collimator was changed from a Thorlabs F230APC-1064 to a Schäfter and Kirchhoff 60FC A4 coupler to avoid the need to use an additional lens at the start of the setup, while the small total path length of the interferometer ensured that this length was within the Rayleigh range of the collimating lens. Instead of being fixed to the baseplate via retention arms, the beam cubes and their spacers are now glued in place into the designated milled areas on the baseplate to ease alignment. In order to reduce size of the interferometer baseplate the waveplates are now attached through custom mounts that are glued to the tops of the beamsplitters. The spaces between the components have also been reduced, further reducing the size.

The lack of any kinematic mount in one of the arms presented challenges when considering the alignment of the interferometer as a whole. A greater focus had to be placed on the alignment of the initial fibre coupler and on the placement of the mirror on the inertial mass to ensure the beam successfully overlaps with the reference beam on the recombination beamsplitters. For this iteration, a Thorlabs K05 mirror mount was adapted to fit the new fibre coupler, though this resulted in the fibre coupler rotating and being challenging to align, often resulting in large drifts.

Shortly after construction, it was discovered that the coil readout and the coil springs were damaged during or shortly after assembly. This resulted in the resonant frequency and the quality factor

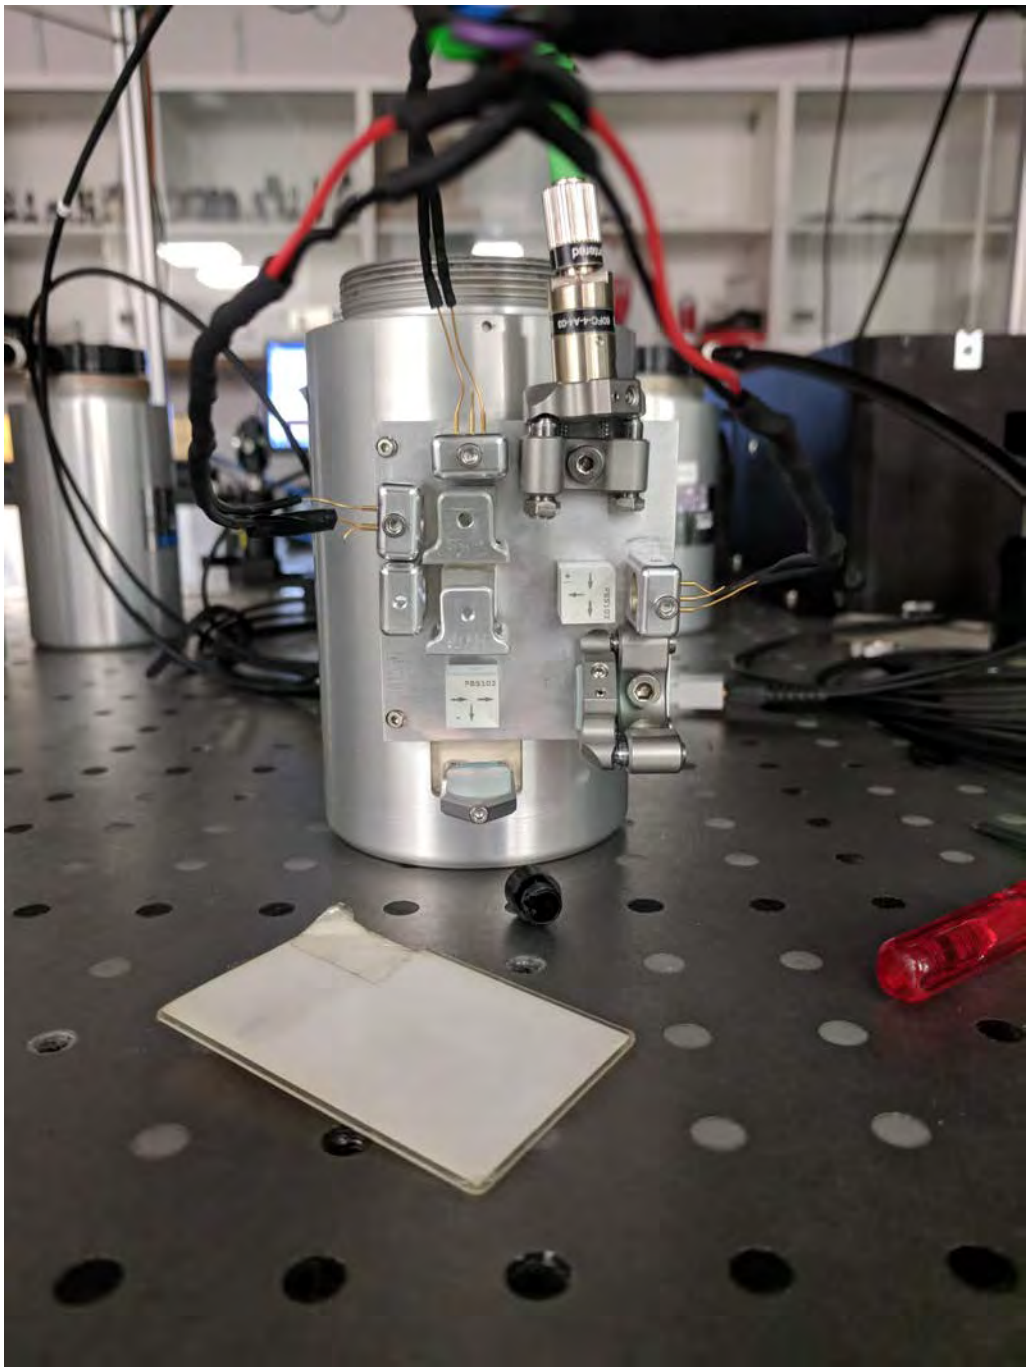

Figure 4.9: Image showing the updated inertial sensor featuring new waveplate holders, glued optics and a baseplate bolted directly to the L-4C.

of the geophone changing significantly and the calibration changing between data runs. As such the data acquired by the optical L-4C was unreliable. However the usability of the L-4C could be evaluated to determine which areas needed improving in future designs and which building

practices could be adapted to prevent the breaking of future devices. From this point forward the optical geophones were assembled upside down, such that the proof mass is resting on its stops at all times limiting the damage that could occur to the springs. Through this limited testing, there was still a number of shortcomings that could be addressed for the next iteration, namely, the fibre coupler mount, beamsplitter placement and waveplate mounts.

4.3.3 MK3 Optical Inertial Sensor

During testing of corner cubes for use in the A+ upgrade, it was discovered the polarising beamsplitters had a poor PER in reflection; as discussed briefly in chapter 2. This meant that instead of having a clean polarisation state incident on the half wave plate, the polarisation state was mixed. This resulted in multiple spurious interferometers being present in the readout spoiling the resolution of the sensor. As such a 45 degree fixed mirror was added such that the first PBS would operate in transmission where the extinction ratio is much larger reducing the influence of stray light from the incorrect polarisation.

During testing of previous devices, the fibre coupler mount, previously a modified Thorlabs K05 Kinematic Mirror Mount was found to be contributing to the decrease in coherence between the interferometer signal and the coil signal, as well as causing significant stability issues that impeded the operation of the inertial sensor. As such, a new fibre mount was designed specifically for the fibre couplers - Schäfter + Kirchhoff 60FC, was designed. A render is shown in FIG 4.11.

The design of the fibre coupler mount was driven by the need to keep the coupler stable not only during operation but also when constructing the device, where the whole inertial sensor would be rotated. The fibre mount sits on three small spacers to ensure the mount doesn't rock while in place. It is fixed to the baseplate via a single M2 screw at the back. Small slots at the front of the coupler allow some angular alignment of the input beam. These can be fixed in place by a recessed, captive M2 nut and bolt. During initial alignment, the input fibre coupler can be used to ensure optimum reflection off the test mirror, such that the beam is incident on each of the

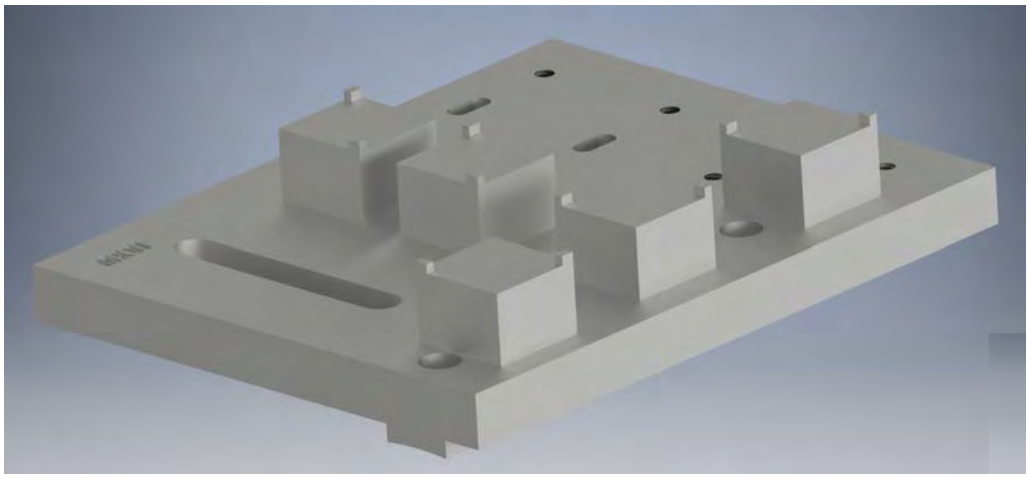

Figure 4.10: Image showing a render of the latest baseplate design. In this version the beam splitter spacers are milled out of a single piece of aluminum with the rest of the baseplate. Each of the spacers has two tabs to further improve placement accuracy of the beam splitters and slots have been milled to allow for translational movement of the reference mirror to reduce frequency noise coupling. Mounting holes for photodiodes, the fibre coupler and the baseplate itself have been positioned not to interfere with one and other and are accessible when fully built.

fixed photodiodes. The accurate placement of the beamsplitters allows the reference mirror to be aligned to match the alignment of the test mirror spot on the photodiodes with the tuning of the reference mirror. The fibre coupler itself is clamped to the main body of the mount, and can be held in place with a single screw for extra stability.

While meeting the size requirements of the previous iteration the waveplates themselves were unable to be rotated after the mounts were glued to the tops of the beam splitters. As a result of this, the waveplates were not at their ideal operating angle and were misaligned by a few degrees. This would lead to significant levels of non-linearity, as shown by chapter 3. To account for this a new waveplate mount was designed to hold both waveplates in a single mount and could be held in place with M2 screws while the interferometer is running, fine markings were etched every 5 degrees to allow for easier controlled tuning of the waveplate rotation. Once at the optimal alignment the waveplates can be glued in place and the holding screws removed, reducing birefringence due to stress.

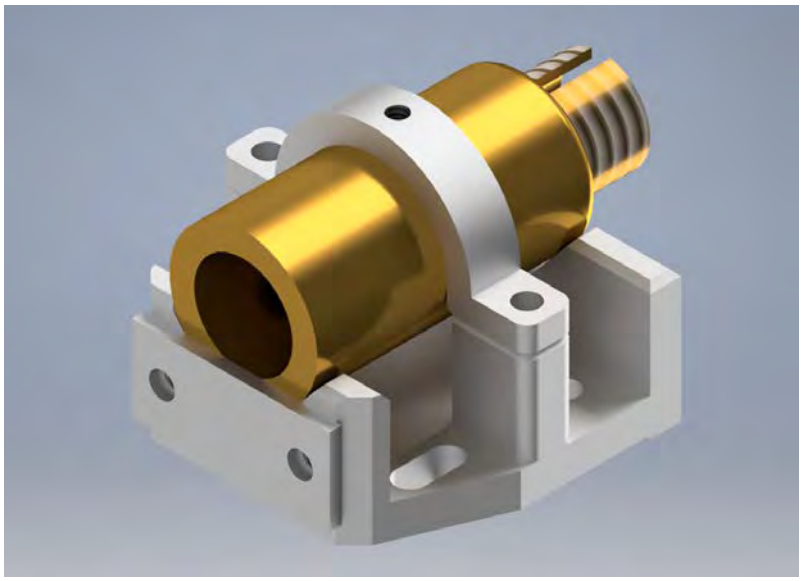

Figure 4.11: Image showing an updated design for the fibre coupling mount. The coupler is held in place in three places to avoid over constraining the coupler. Fixed place holes and a slot are used to provide angular adjustment for the whole interferometer.

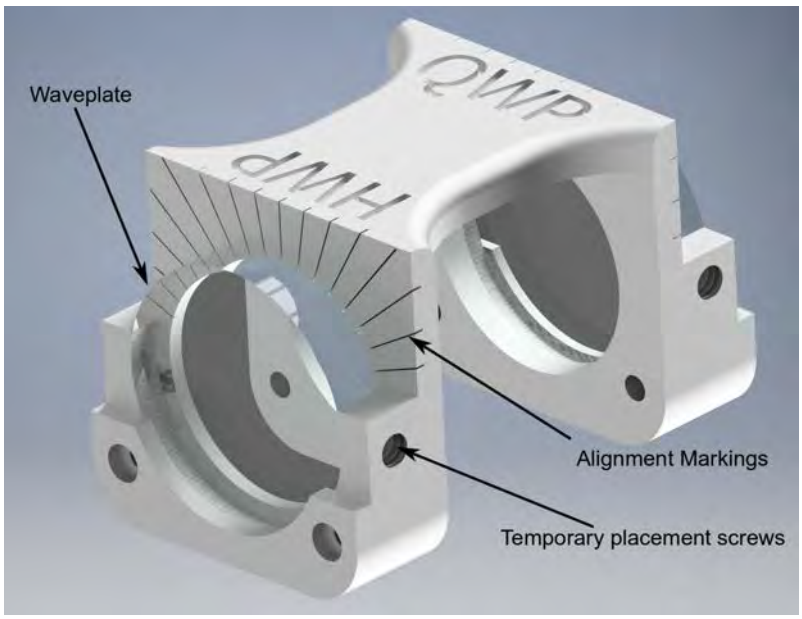

Figure 4.12: Rendered image showing the design of the waveplate mount used for the current prototypes of optical inertial sensors. Here the waveplates can be rotated during alignment to ensure they are at the optimum operating angle to reduce non-linearities in the device.

FIG 4.13 shows a render of the optical inertial sensor, with an assembled version is shown in FIG 4.14. The overall design is compact enough to fit on the side of the geophone, which it is mounted using three M2 screws to avoid the baseplate tilting with respect to the can of the inertial sensor, reducing tilt coupling in the device. A mirror is mounted using a light weighted mount and glued to the side of the moving mass. This is done upside down so the milled side of the geophone can be used as a flat horizontal reference, reducing the angular misalignment of the test mirror.

MK3 Initial Results and Future Work

FIG 4.15 shows a sample of data measured using the MK3 optical inertial sensor, in this test an L-4C is measured by an interferometer and a coil simultaneously, allowing some coherent information to be subtracted. Shown in red is the signal measured by the interferometer, the coherent subtraction is shown in blue and is compared against the best previous subtraction with the MK2 interferometer (magenta); taken with different ground motion, along with the coil L-4C noise (dashed magenta). The theoretical performance of the optical inertial sensor is shown in black. The figure highlights the improvements that have been made through the various design changes detailed above that result in much greater coherence between the two readout mechanisms. This increase in coherence means that as much as an order of magnitude more motion can be subtracted at 1 Hz, despite only using a single sensor compared to the magenta trace from the first prototype, that used three sensors. As shown by Kirchhoff [95], the coherent subtraction increases substantially when subtracting motion using more than one geophone, as the tilt motion of the table can be subtracted.

The arrows on the plot indicate the likely steps required to measure down to the predicted noise floor of the device. Below 100 mHz the device is likely limited by thermal and air current noise, which can be reduced by placing the devices in an insulated box. This is because the measured signal from previous tests was only a factor of 2 above the suspension thermal noise and likely a similar level of insulation would result in the noise floor of the sensor being reached.

Between 100 mHz and 1 Hz the subtraction is reaching the noise floor of the coil L-4C, which is

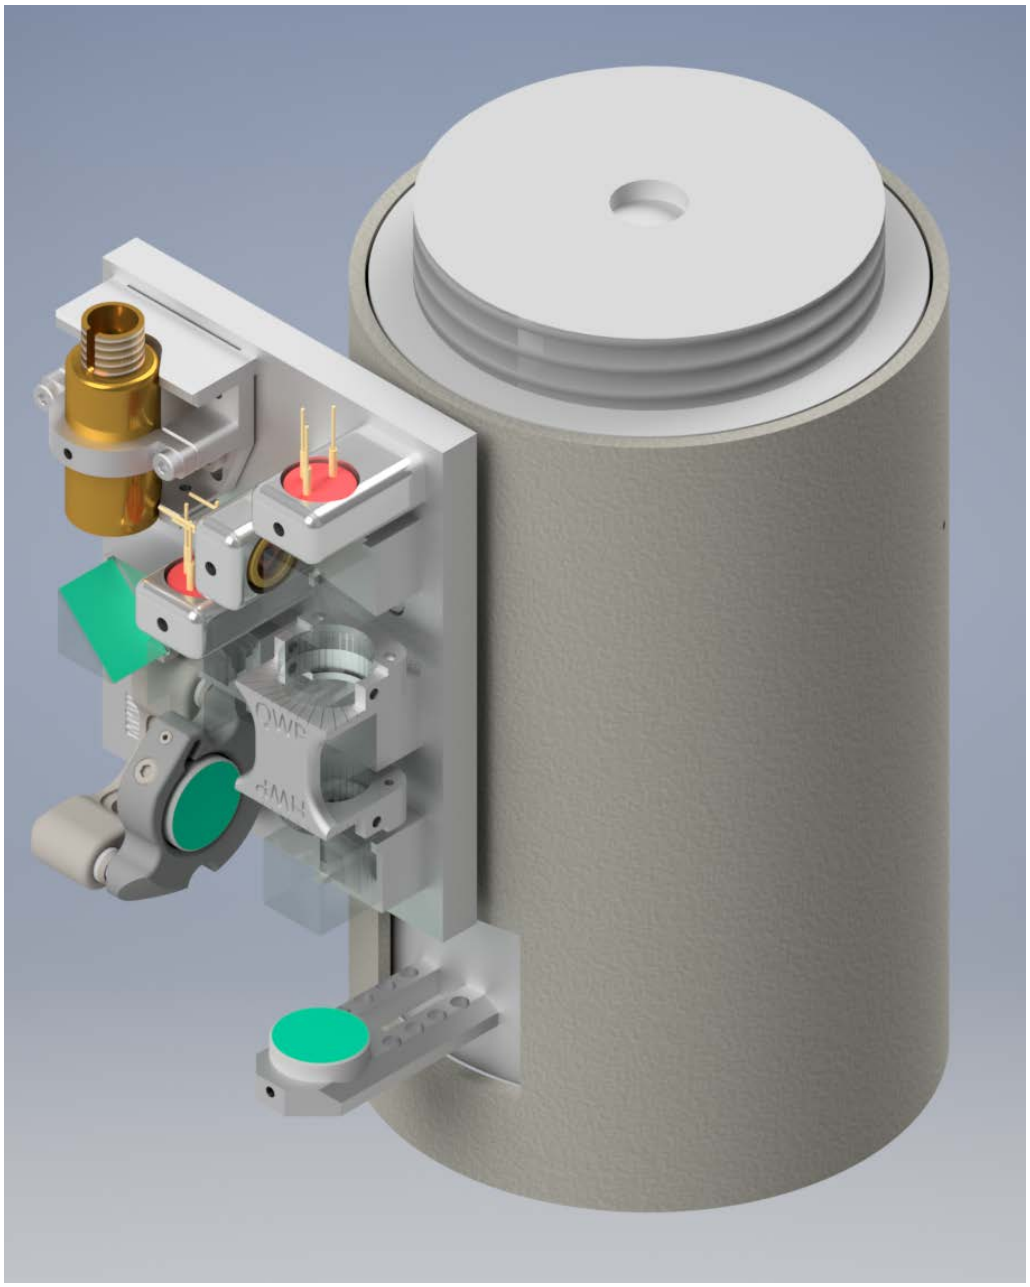

Figure 4.13: Image showing a rendered version of the optical inertial sensor.

only a factor of 5 noisier than the predicted resolution of the sensor, which at this frequency is limited by suspension thermal noise. This motion can be reduced by measuring multiple optical inertial sensors, which should have a noise floor lower than a coil L-4C in this frequency range, allowing more coherent motion to be subtracted and the noise floor to be reached.

Between 1 and 10 Hz, despite the coherence dropping with frequency, adding more geophones, both

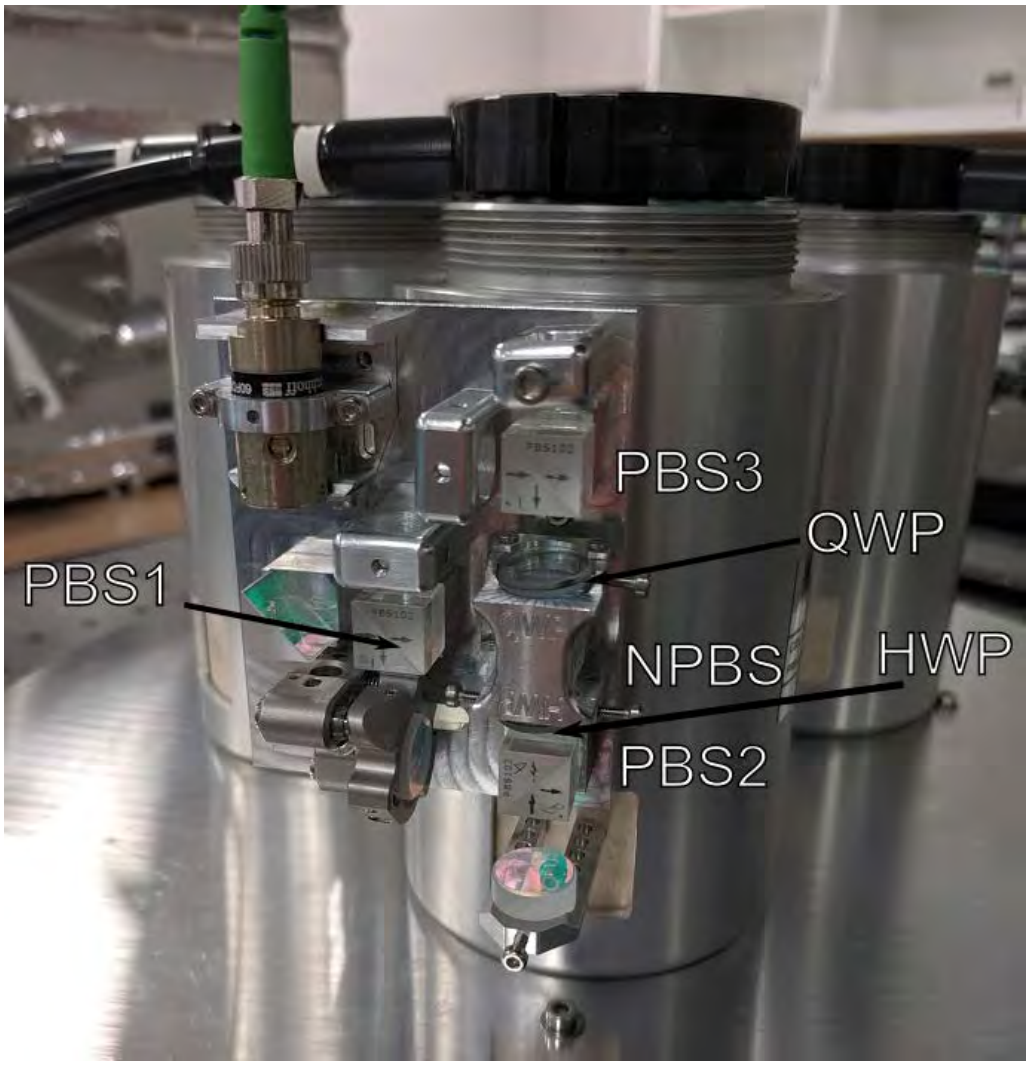

Figure 4.14: Image showing an assembled MK3 inertial sensor being tested on a platform with two other coil readout L-4Cs.

readout by coils and interferometers, should reduce the motion seen by the main optical geophone again as shown by Kirchhoff [95]. Active isolation will be required in this region to further reduce the motion seen by all interferometers.

Above 10 Hz, the coherence between devices drops sharply, thus active and passive isolation will be required in order to measure the noise floor of the geophone in this frequency range. A single stage of isolation, with a 10 Hz resonance frequency and a $\frac{1}{f^2}$ roll-off would be sufficient to achieve this, as only a factor of 100 isolation is required at 100 Hz. Taking data during the night, with

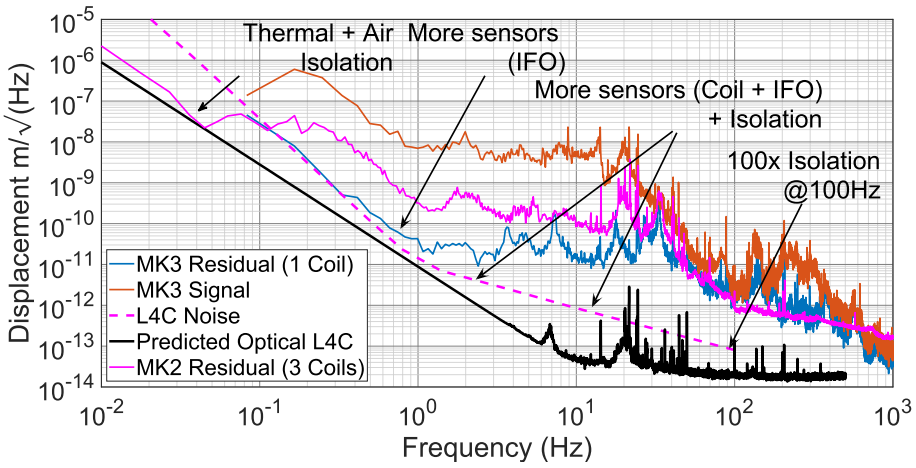

Figure 4.15: Figure showing the initial results from the MK3 optical inertial sensor. Here the predicted motion (black) is compared with the signal measured by the MK3 (red) and its coherent residual (blue) by subtracting coherent motion measured by a single coil geophone. Coherent subtraction results from the MK2 inertial sensor, obtained by subtracting coherent motion measured by three coil geophones are shown in magenta, with the traditional L4C noise shown in dashed magenta.

reduced human driven seismic activity would also improve this measurement.

Performing these adjustments to the testing procedure should result in a huddle test where the sensor noise of the L-4C can be measured successfully. At low frequencies, this would result in the suspension thermal noise of the L-4C being measured as the dominant noise source below 2 Hz. To measure the resolution of the sensor above this frequency, seismic isolation will have to be used to reduce the ground motion to a sufficiently small level.

4.4 Summary

L-4C geophones are inertial sensors that measure ground motion using a coil and a magnet. At frequencies lower than approximately 0.06 Hz, the sensors begin to measure their own self noise, caused by the Johnson noise of the readout coil [95]. Interferometers, such as those presented in

chapter 2 are an alternate readout mechanism, and in the past have been used to improve the resolution of other inertial sensors [72].

In this chapter, a compact interferometric inertial sensor, that uses HoQI as a readout has been developed using the mechanics of an L-4C. Such a sensor, while limited in testing due to manufacturing issues discussed in detail in section 4.3, has been shown to have a higher resolution when compared with the coil-magnet readout geophone. These two readout methods were tested simultaneously for half an hour. The interferometric readout has a resolution of a factor of 60 higher than the coil-magnet readout. Moreover, at approximately 50 mHz the interferometric sensor reaches the suspension thermal noise of the springs in the L-4C.

While the low frequency performance of the interferometric inertial sensor could be immediately realised, measuring the sensor noise at frequencies higher than 60 mHz proved challenging. Multiple sensors were tested in unison allowing for subtraction of coherent motion from the interferometric sensor. This allowed the self noise of the coil L-4C to be reached between 60 mHz and 1 Hz. Above these frequencies, additional sensors, both coil and interferometer based, can be used for improved coherent subtraction, as well as active and passive seismic isolation will be required to measure the self noise of the interferometric inertial sensor.

Chapter 5

Improvements to Seismic Isolation in Gravitational Wave Detectors

Control, control, you must learn control!

Master Yoda

This chapter describes my work as a LIGO fellow in the development of a MATLAB model of the HAM ISIs in use at both LIGO observatories and the improvements to suspension systems by using HoQI. The objective of the HAM model was to predict the measured in-loop platform motion using only out-of-loop sensors. This requires the model to be modular, allowing for new filters and sensors to be evaluated, serving as a prediction of the impact to the isolation performance new sensors would provide. In addition, the predictive nature of the model allows each of the control loop paths to be calculated separately. Permitting the model to function as a noise budgeting tool; useful for commissioning or diagnosing problems with the ISI. Text and figures in this section have been copied from the HAM ISI model tech note [99] of which I was the principle author and was the main contributor to the code used in the model. It is available at `https://dcc.ligo.org/LIGO-T1800092`. The text and figures in the subsection on controlling Signal Recycling Cavity Length (SRCL) motion is copied from [100], of which I was the principle

author. It is available at `https://dcc.ligo.org/LIGO-T1900107`.

5.1 Introduction

To reduce the coupling of ground motion into the gravitational wave detection band, LIGO uses both ISIs and multiple stages of pendulums to isolate the optics from ground motion; the number of stages of isolation depends on the payload that is being suspended. Methods of improving performance of the ISIs will be the main focus of this chapter. Specifically the single stage HAM ISIs, shown in FIG 5.1, which isolate the auxiliary optics, such as the signal recycling cavity, from ground motion. The two stage BSC ISIs that are responsible for part of the isolation of the test masses and beamsplitter will not be discussed, due to its more complex nature. As it necessitates the modeling of numerous cross couplings, which is beyond the scope of this chapter. Both the HAM and BSC ISIs use a combination of active and passive isolation, used at low (sub 30 Hz) and high (above approximately 1 Hz) frequencies, respectively.

5.1.1 Introduction to Control Loops

To reduce the motion of the ISIs, the motion measured by sensors fitted to the platform is fed back to actuators through a control loop. These actuators then drive the platform to counteract the input ground motion. A simple loop diagram of this process is shown in FIG 5.2. When the loop open, i.e. the feedback is turned off, the output motion is simply,

$$\text{Out} \quad = \quad \text{P} \times \text{Input}, \tag{5.1}$$

where P is the plant. When the control loop is closed the transfer function becomes slightly more complicated as the output depends on the previous cycle of the control loop. In the steady state this can be analysed in the frequency domain. By working backwards through the loop, it can be shown that the closed loop response is given by,

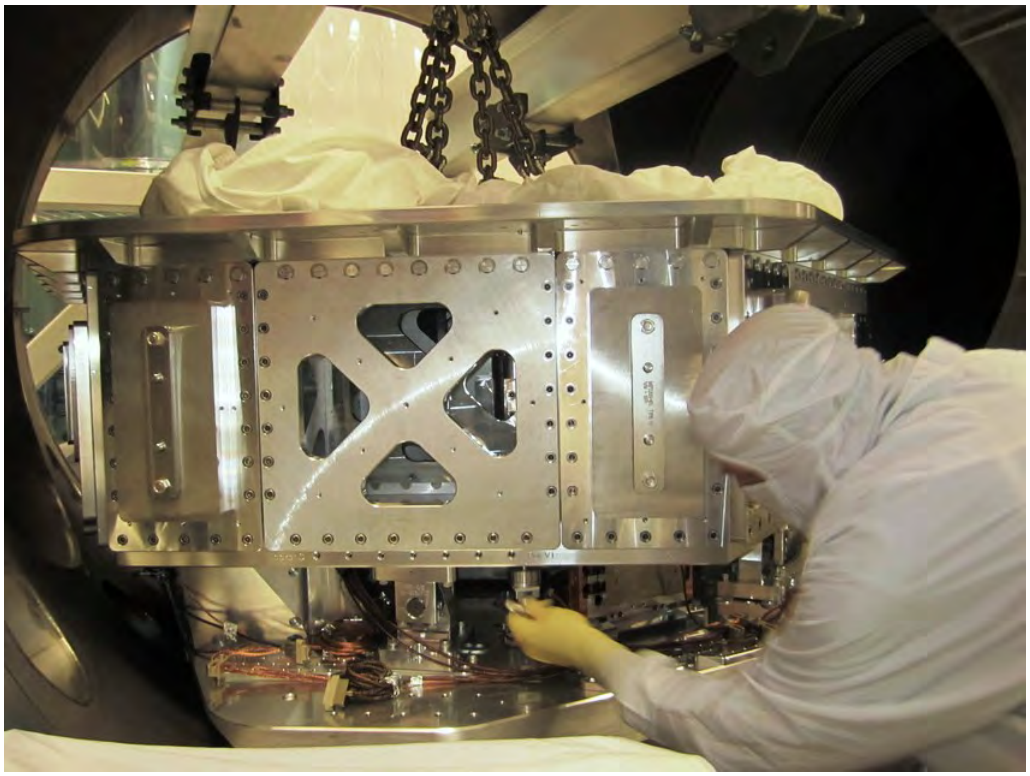

Figure 5.1: An image of a fully assembled Advanced LIGO ISI. Credit: LIGO Laboratory [10]

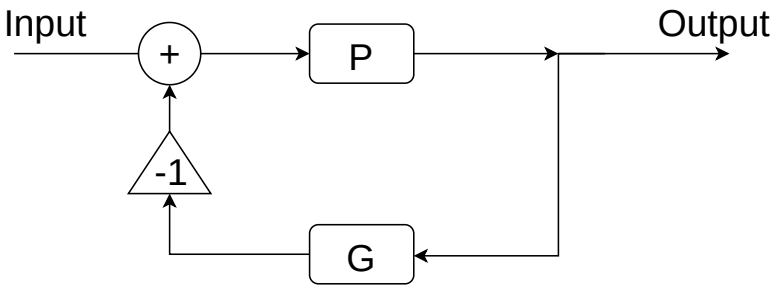

Figure 5.2: A schematic of a simple control loop providing feedback, P defines the plant, of the platform and G is the gain applied to the feedback loop.

$$\begin{aligned} \mathrm{Out}(\omega) &= -G(\omega)\mathrm{Out}(\omega) + \mathrm{P}(\omega)\mathrm{In}(\omega), \\ \mathrm{Out}(\omega)\,(1+\mathrm{G}(\omega)) &= P(\omega)\mathrm{In}(\omega), \\ \mathrm{Out}(\omega) &= \frac{P(\omega)\mathrm{In}(\omega)}{1+G(\omega)}. \end{aligned} \tag{5.2}$$

Here the plant multiplied by the input motion is suppressed by a factor of $\frac{1}{1+G(\omega)}$, where $G(\omega)$ is

the gain of the system. A more detailed description of control loops can be found in the book by Goodwin [101] or the tutorials by Shapiro [102].

5.1.2 Seismic Isolation in LIGO

The ISI is sensed by a combination of displacement and inertial sensors. These take the form of 6 CPS and a further 6 GS13 geophones. The working principle of the latter is explained in chapter 4. A schematic of sensors on the HAM ISI is shown in FIG 5.3. In addition to these sensors, L-4C geophones are used to measure local ground motion inside each chamber, with the exception of HAM1 which does not feature an ISI at the time of writing. Only HAM4 and HAM5 have L-4Cs on stage 0 of the ISI. All HAM chambers have have eight L-4Cs, split evenly between horizontal and vertical sensors and are located on the HEPI. In addition to these sensors, there are STS-2 seismometers located near HAM2, HAM5 and ITMY in the case of LLO, and a single STS at LHO, which are used for measuring low frequency ground motion and for sensor correction.

The control loop of the HAM ISI is shown in FIG 5.4. Using the X degree of freedom as an example, blocks denoted with F^x represent a filter, these filters describe the amplitude and phase of a sensor used in feedback or feedforward when summed into the rest of the loop. Blocks denoted with a K_x represent controllers, which specify the gain and phase of a controller at a given frequency, these take form of the isolation controller, K_x^{I} and the damping controller K_x^{D}. These are responsible for the active isolation control and the damping control respectively. Contributions to the overall motion caused by the self-noise of each of the sensors on the ISI are denoted by n_{SN}^x, where SN is the name of the sensor. Input motion paths, as measured by sensors on site, are denoted as x_{ST0}.

In the model this motion is determined by stitching together multiple sensors to minimise sensor noise injection, as is described in detail in section 5.2.1. In reality, sensor noises are combined with the measured input motion. The control loop presented in FIG 5.4, and therefore the model, does not take into account higher order effects, such as back reaction from the suspensions and couplings from HEPI. Mechanical cross couplings apart from the tilt to translation coupling between RX and

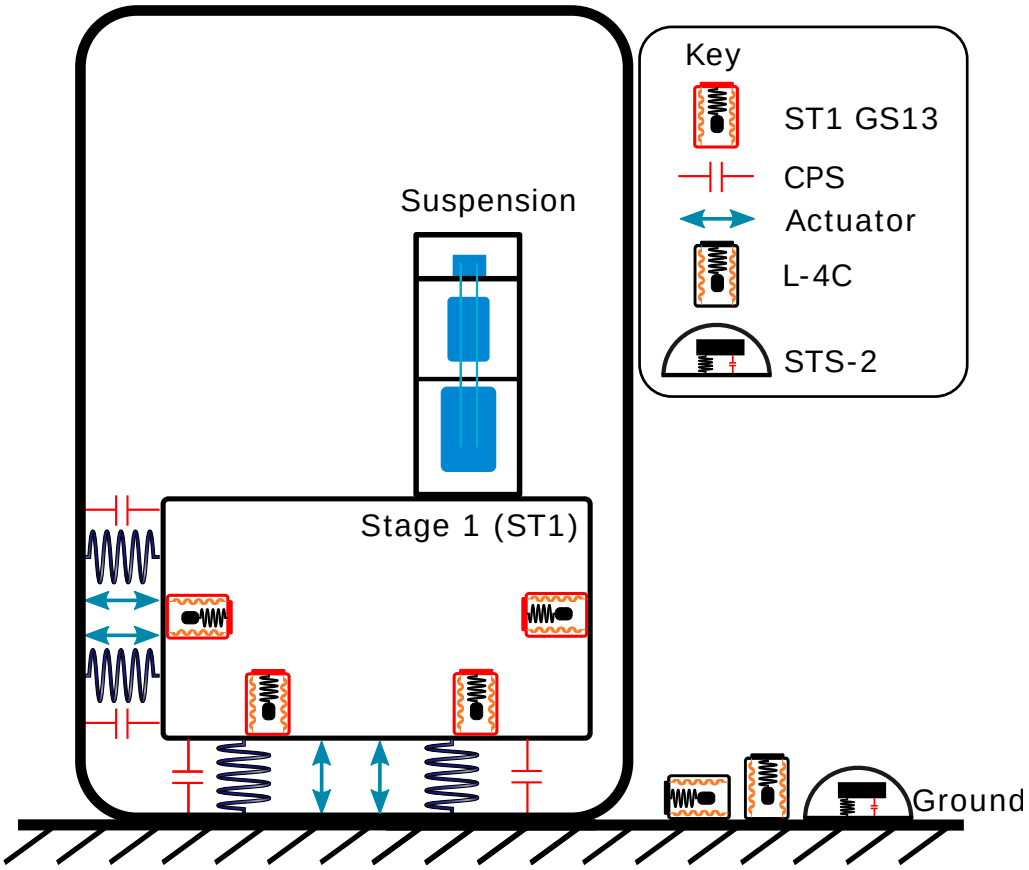

Figure 5.3: A schematic of a HAM ISI showing the sensors and actuators used to control the isolated platform. HEPI is not shown here for simplicity.

Y and RY and X are also not accounted for. These have been omitted to reduce the complexity of the model, as their integration point and coupling factor in the model are not known at the time of writing. If the accuracy of the model needs to be improved further then these additional cross couplings and suspension back reaction would need to be included.

The control of the seismic isolation platforms is complicated by tilt to translation coupling, the process where low frequency tilt is indistinguishable from translation. FIG 5.5 show how this effect occurs in a geophone, however this effect also occurs in the seismometers and isolation tables, as discussed in this thesis. This occurs because below their resonance frequency seismometers are only sensitive to acceleration or the local gravitational field when the device is tilted. This looks identical to translational motion and therefore the tilt motion is indistinguishable from real translational

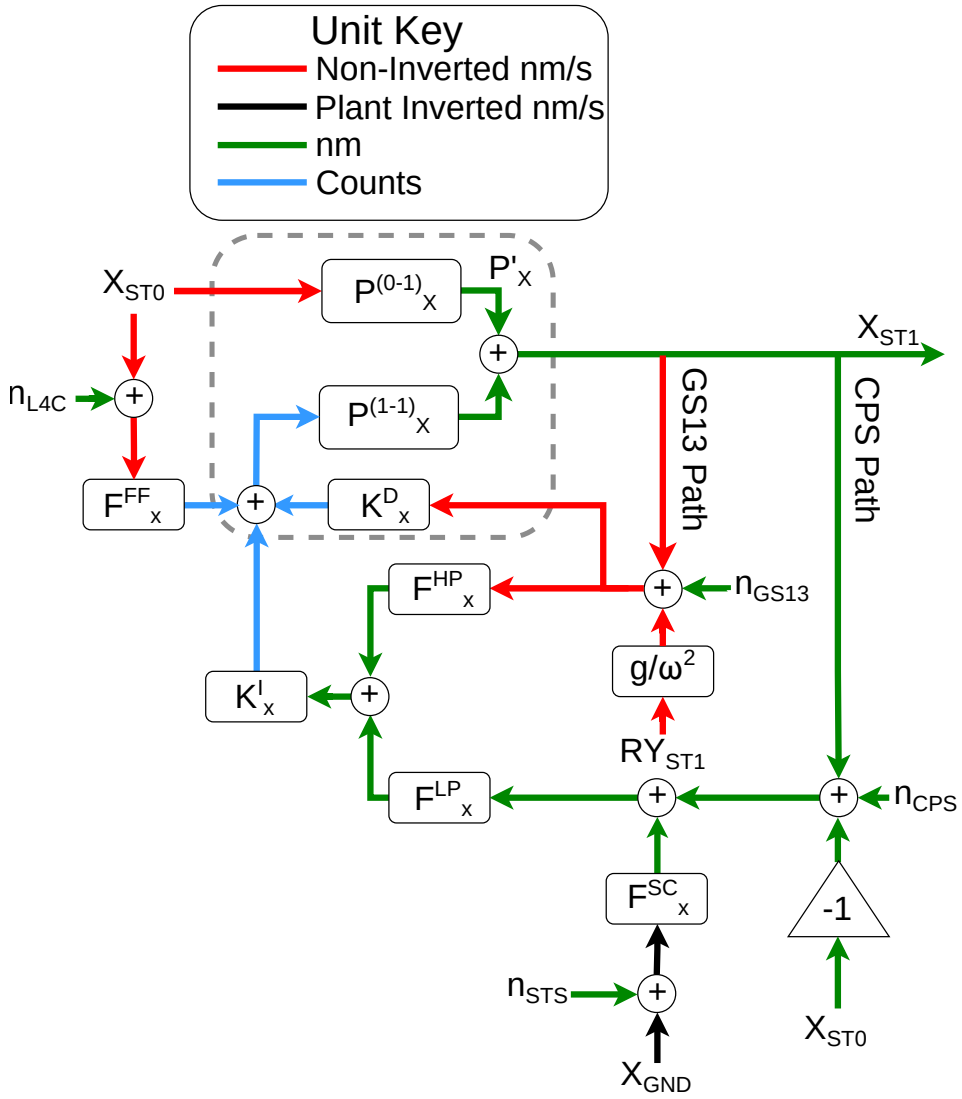

Figure 5.4: A control loop diagram of a HAM ISI, showing filter modules e.g. the high pass filter $F_{\mathrm{x}}^{\mathrm{HP}}$, plants e.g. $P_{\mathrm{x}}^{(0-1)}$ describing the ground to platform transfer function, $P_{\mathrm{x}}^{(1-1)}$ describing the actuator to platform transfer function, and controllers $K_{\mathrm{x}}^{\mathrm{I}}$ e.g the isolation filter. Sensor noises are described with the notation n_{SN} while ground motion inputs are denoted as x_{ST0}, adapted from [11].

motion. The acceleration seen by the mass when tilted by a small angle θ, is proportional to $g\theta$. This can be converted to displacement trivially by Fourier transforming and integrating twice, giving the familiar tilt to translation coupling factor of,

$$\delta x = -\frac{g}{\omega^2}. \tag{5.3}$$

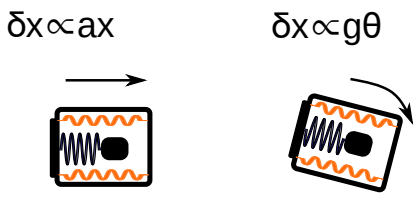

Figure 5.5: Figure showing how tilt couples into translation in horizontal geophones. Using the small angle approximation, tilt is indistinguishable from any translational motion below the resonance frequency of the seismometer and couples to translational displacement by a scaling factor of $\frac{g}{\omega^2}$. Here g is the acceleration due to gravity and ω is the angular frequency.

This couples RY platform motion to X, and RX motion into Y with a coupling factor of $-\frac{g}{\omega^2}$ and is denoted by the block RY_{ST1} in FIG 5.4. The passive response of the platform is taken into account by the block $P_x^{(0-1)}$ and the response of the actuators denoted by $P_x^{(1-1)}$. X_{GND} represents the STS-2 path used for sensor correction in X, Y and Z.

Due to the unbounded noise of the GS13s below their resonance frequency, sensor blending is employed through the use of high and low pass filters, denoted as F_X^{HP} and F_X^{LP}, respectively. These are used to attenuate unwanted noise injection from the GS13s at low frequency and CPSs at high frequency. Due to their superior noise performance at high frequencies and their inertial nature, GS13s are used as a reference for platform motion. The CPS measures the difference between platform motion and ground motion. In addition to the blending filters, feedforward paths, such as the sensor correction, F_x^{SC} are used to add ground motion to the CPS to reduce the effect of the microseismic peak at 0.1 to 0.3 Hz on platform motion. The design of these sensor correction filters is covered in chapter 6. Another feed forward path is added to subtract coherent motion at around 10 Hz using the local L-4Cs, denoted by F_x^{FF}.

The colours in FIG 5.4 denote the units of each filter and input motion as measured in Control and Data System (CDS). These must be converted to a common unit in order to correctly calculate

the final platform motion.

During normal operation, only certain data channels are saved at full rate, examples of which are the ground seismometers (L-4Cs and STS-2s), the displacement sensors (CPS) before and after sensor correction and the on platform seismometers. Other important terms that are present in the control loop, such as sensor noises, are not saved. However, their contributions to the overall stage 1 motion of the ISI are important as they are the limiting the isolation performance of the ISI at some frequencies. An example of this is the on platform GS13s, which limit the isolation performance in X below 0.05 Hz, while the CPS noise is limiting between around 0.6 and 10 Hz in RX, as shown in FIG 5.12 and FIG 5.14 respectively.

Moreover, the filters discussed above are digital filters and as such can be changed relatively easily using LIGOs CDS. However there is limited time to evaluate the performance of these new filters, as experimental filters are only tested when the detector is not in observation mode. Instabilities or unwanted effects of these filters can cause the detector to lose lock and prevent observation. Even outside of observation times, the performance of the filter takes several hours to evaluate and has to be conducted around other upgrades to the detector. Local noise sources, such as people walking around the chambers and earthquakes, can cause issues with these tests. Evaluating the improvements to the ISIs through the use of new high resolution sensors, such as those described in chapters 2 and 4, is an even more complex task. Installing new sensors is a time intensive process, which is why it often takes multiple years to qualify sensors for use inside the vacuum enclosure. In order to quantify improvements to the ISI, both in terms of optimising filters and replacing current sensors with their state of the art counterparts, a predictive and modular model of the ISI has been created and is described below.

5.2 Description of the HAM Model

5.2.1 Inputs

An overview of the main inputs to the HAM model are shown in FIG 5.6. Here the user selects mandatory options such as the start time, duration, chamber, degree of freedom and interferometer to run the model on. Other options, such as which local L-4Cs to use and the selection of which feed-forward paths to activate can be selected. The code checks to see if data matching these options is already cached to file, to reduce load on the Network Data System (NDS) servers and to speed up the model. A similar process is repeated for both the filter files and for information about which filter was active at the time.

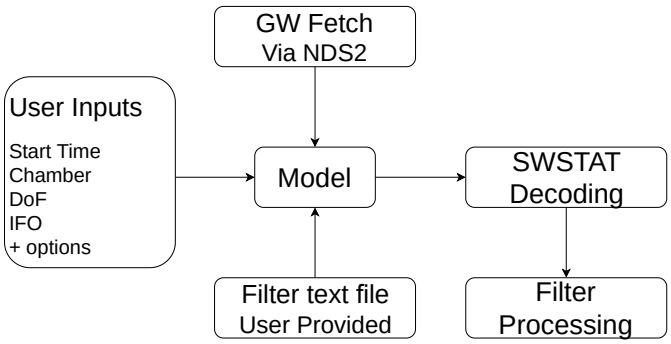

Figure 5.6: An overview of the input components used by the HAM ISI model to calculate the expected platform motion.

Measurements of the local ground motion is provided by the on site seismometers, instead of using an approximation and fitting to a typical ground motion curve for use in all cases. The model uses real out-of-loop sensors to allow for the fact that, unlike sensor noise, the input ground motion is not static in amplitude, will change over time, and can appear in the measured stage 1 GS13 signal. If fits to ground motion were used, these local noise sources would not be taken into account; causing discrepancies between predicted and measured motion.

To avoid the measured ground motion being saturated by sensor noise, multiple sensors are combined to provide a stitched input ground motion spectra and are different for translation and tilt

degrees of freedom. The code converts the downloaded time series data into the frequency domain by taking an ASD over a configurable number of averages. This is done to improve the speed and simplicity of the model. A downside to this approach is that phase information from the sensors is lost, and as such feedback paths must be weighted by the coherence, to prevent the over subtraction of motion. This coherence is measured between the sensor providing the feedforward signal and the on platform GS13 during the O2 winter break when the platforms were in damped only mode. During this time all isolation loops and feedforward paths are turned off. Weighting the feedforward paths by their respective coherences will limit the amount of subtraction possible due to feedforward.

The input motion paths are summed with sensor noises in quadrature to avoid over predicting subtraction from feedforward loops. The sensor noises are taken from datasheets and previous measurements [12] and are interpolated across the model's working frequency vector.

Ground Motion Stitching

Currently, for both tilt and translational motion, there exists no single sensor placed on the ground on site that is capable of measuring ground motion over the required frequency range of 10 mHz to 100 Hz without measuring cross-coupling or sensor noise. As such we must use combinations of sensors to model the input ground motion; tilt and translational degrees of freedom require different combinations of sensors and as such are discussed separately.

For the translational degrees of freedom (X,Y,Z), a combination of local L-4Cs and the corner station STS-2 broadband seismometer are used to construct the input ground motion. Rather than using a blending filter, which may be a source of confusion, the sensors are stitched together at 0.8 Hz to account for the different characteristics and physical locations of each sensor. This stitching process cuts the vectors at the stitching frequency and concatenates them. Any error in the stitching process would be obvious, as there would be a step in the measured ASD. The stitching frequency could then be changed to eliminate this step. Like the GS13s, the L-4Cs present on each ISI are inertial sensors. As such, their measured signal decreases as $\frac{1}{f^2}$ below their

resonance frequency of 1 Hz. Therefore, these sensors are not ideal for measuring low frequency motion below 0.1 Hz in translation, and 0.5 Hz in tilt.

Due to their lower resonant frequency and principle of operation, the broadband STS-2s encounter this problem at a much lower frequency of 8 mHz. During times of high wind velocity the STS-2s begin to couple comparable ground tilt to the measured translational motion, and as such are not a perfect sensor. This is covered in more detail in chapter 6. Nevertheless they are the best sensor available to measure low frequency seismic motion on site and thus are used in the model to measure ground motion at low frequency.

FIG 5.7 shows the stitching of these two sensors, 0.8 Hz is chosen as a suitable stitching frequency due to the overlapping amplitudes of the two signals in the majority of cases. This method requires signals from both seismometers to be converted into displacement units, with the response of the mechanics in each sensor to be accounted for. The STS-2 is integrated, transforming the signal into displacement. The L-4Cs have a more complicated response and must be corrected by equation 4.3 in chapter 4.

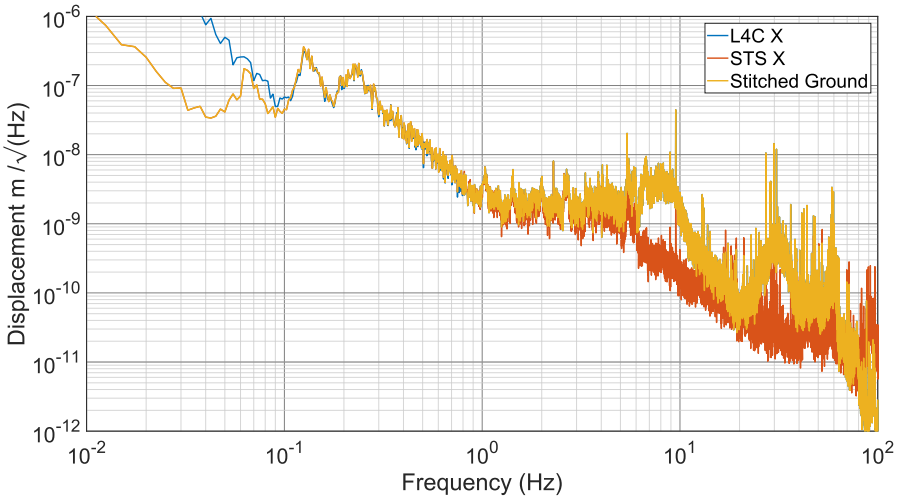

Figure 5.7: Figure showing the input signals used to make a ground motion estimate for the HAM model. Here the L-4C signal (blue) is used at frequencies above 0.8 Hz and is stitched with the STS-2 (red) which is used at frequencies below 0.8 Hz.

The sensor stitching for the rotational degrees of freedom is site specific due to the different sensors available for use at the two sites. In the case of LHO, there is no ground rotation sensor present in the corner station. The L-4Cs on each chamber are used as a reference above 0.3 Hz for ground rotation. Below this frequency, the STS-2's in Z are fitted to approximately match the amplitude of the Beam Rotation Sensor (BRS) at ETMX. The result of this is shown in FIG 5.8. This approximation is not perfect, but should be sufficient to get within a factor of a few of the real motion.

The ground rotation is sufficiently different across the 4 km sites that the BRS at ETMX cannot be used to reliably measure the ground rotation at the corner station, however it can be used to estimate the ground rotation to the correct order of magnitude. The exact scaling of this is then fitted to the sub 0.1 Hz region of the measured motion trace of a HAM ISI. The Z degree of freedom of the STS-2 is used as this should be the most resilient to tilt coupling at low frequency. The L-4Cs begin to show evidence of being noise limited in the tilt degrees of freedom from as high as 0.5 Hz. However, 0.3 Hz was chosen as a compromise between limiting sensor noise injection and avoiding large steps in the ASD. This method of tilt estimation could be improved further by better shaping of the response of the STS-2 in Z to more closely match the ASD of the BRS, or by choosing a different stitching frequency. The agreement with the GS13 on stage 1 is sufficient for the purposes of the model, as the GS13s used to measure the platform tilt are limited by sensor noise at a similar frequency.

During preparations for O3, BRSs were installed at the corner and end stations of LLO, and as such these can be used as a measure of low frequency ground rotation in the corner station. They can be used for the ground rotation input, provided the start time for the model is after October 2018, any time before this date the LHO ground estimation technique using LLO data is used. Both the internal sensor noise and the location of the BRS prevents it from being used above 2 Hz; the measured ground rotation is not common among chambers above this frequency and therefore cannot be measured by a single BRS. Above 2 Hz the chamber's local L-4Cs are used to provide rotational information, as these are able to measure dynamics of HEPI. An example of

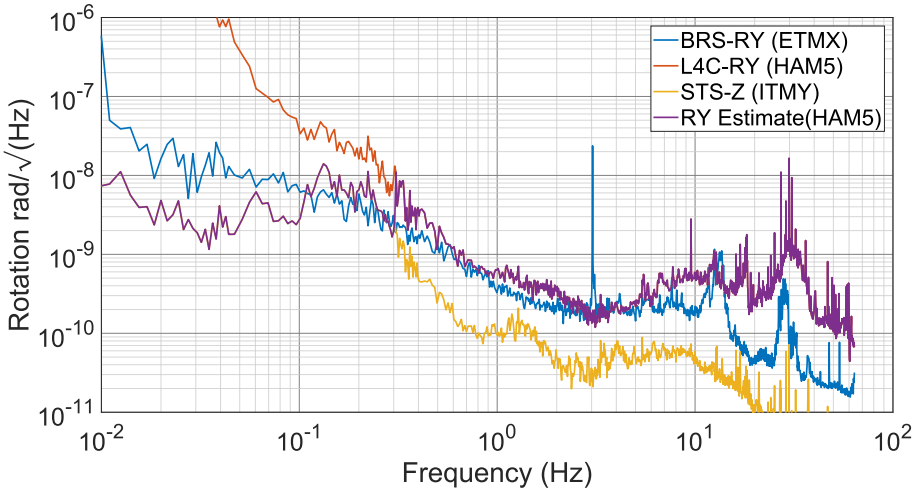

Figure 5.8: Figure showing how ground rotation is estimated at LHO. The BRS (blue) is located at one of the end stations and is used as a guide for stitching the local STS2-Z (yellow) and L-4C-RY (red) to make a ground rotation estimate (purple).

this stitching is shown in FIG 5.9. Like the translation case, the response of the L-4C must be accounted for to convert it into displacement units.

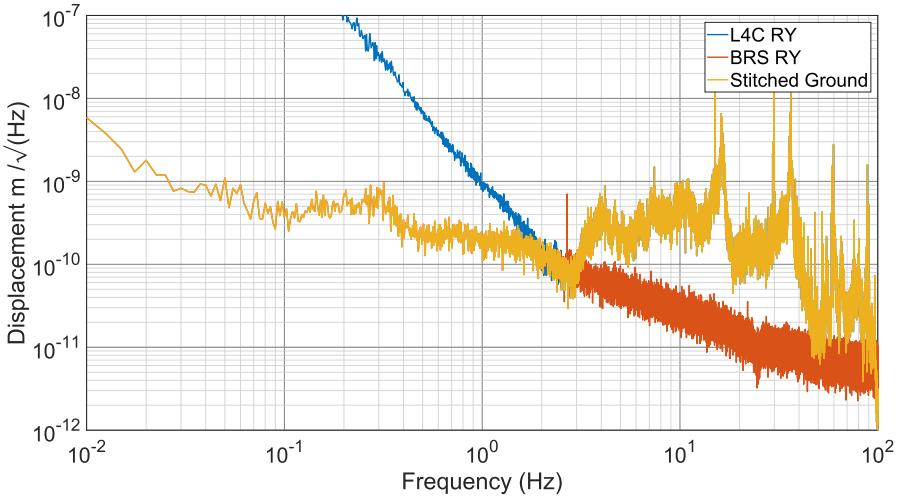

Figure 5.9: Figure showing the stitching of input signals at LLO using the BRS (red) at frequencies below 2.5 Hz and L4Cs (blue) above 2.5 Hz, producing a estimated ground rotation spectrum (yellow).

5.2.2 Platform and Actuator Responses

The response of the passive isolation of the platform and the actuators must be taken into account, which are referred to as the ground and force plants respectively. The ground plant is the transfer function between Stage 0 (ST0) (the ground L-4C's) and the stage 1 GS13s, or the HEPI L-4Cs and the stage 1 GS13s, which is referred to as the ST0 to Stage 1 (ST1) plant or $P^{(0-1)}$ in FIG 5.4. The transfer function defining the response of the table was measured during the O2 winter break where the isolation platforms were set into damped only mode, in which all feedforward paths and the isolation loop were disabled. The force plant is the transfer function between ST1s actuators and the ST1 GS13s, and is referred to as the ST1 to ST1 plant and is denoted as $P^{(1-1)}$ in FIG 5.4. This was determined by measuring the transfer function between the actuators and the GS13 on stage 1. ZPK systems were produced by fitting a system to the measured transfer function, these are then used to create the isolation and damping loop gains used in the model. An advantage to using ZPK systems is that they can be easily extrapolated to any arbitrary frequency vector by taking the frequency response of the system. A downside to this method is that when the responses become complicated at high frequencies, information about their dynamics will be lost.

Having a good fit to the plant is essential for the isolation and damping open loop gains. Improper fits will cause issues, such as an incorrect open loop gain. For LLO, ground plants from Hanford are used due to problems when fitting the ground transfer functions. The plants used should be sufficient due to the similarities between the ISIs, though this could be a source of inaccuracies with feed forward projection when running the model on LLO data. This can be fixed by interpolating the measured transfer function to the size of the model's frequency vector and performing a matrix multiplication on the two vectors.

5.2.3 Filters

The HAM model requires 6 control filters: the Sensor Correction (SC) filter, High Pass (HP), Low Pass (LP), L-4C Feed Forward (FF), isolation filters (ISO) and damping filters (DMP). Each of these filters is recorded in text files sorted by filter name and filter module. These filters can either be downloaded directly from the site, ensuring that the active filters at the selected start time are used, or specified as a text file. Information about which filter is active on the ISI is stored in a bit filtered number between 0 and 65,532, representing 16 different switches for a particular filter bank. The encoding of these is specified in Table 5.1. These bit filtered numbers are stored in SWSTAT channels resulting in a single number which is then decoded into 16 switches. For example, if the value of this channel is 11, this can be represented as 1101 in reverse binary, using the conversion table this indicates that filter modules (FM) 1,2 and 4 were active at the time.

Once the numbers from the SWSTAT channels have been pulled from NDS and decoded to represent the active filter modules, the function `readfilterzpk.m` can be used to read the filter. These can be specified by the user or downloaded automatically from the Subversion (SVN) archive.

5.2.4 Coherence

Without a time domain model the sensor correction and feedforward paths may overestimate the amount of subtraction that can be obtained when estimating platform motion. To correct for this, the sensor correction and feed forward paths are weighted by the coherence between their measurement sensors and the stage 1 GS13 when the isolation table is in a damped only state. The measured coherence becomes more accurate the more averages it is measured over, as such 35 averages are taken over an hour long stretch of data. This strikes a balance between having a large number of averages and measuring across the frequency region of interest.

This feature is only enabled on LHO data due to a lack of suitable damped only time at LLO. As such, examples of improvements to the ISIs shall be discussed using only Hanford data. This coherence data is then interpolated to span the working frequency vector used in the input ground

Table 5.1: Table showing the encoding of the bit filtered SWSTAT channels used by the model to pull filters in use.

Name	Value	Bit
FM1	1	0
FM2	2	1
FM3	4	2
FM4	8	3
FM5	16	4
FM6	32	5
FM7	64	6
FM8	128	7
FM9	256	8
FM10	512	9
Input On/Off	1024	10
Offset In	2048	11
Output On/Off	4096	12
Limit	8192	13
Unknown	16384	14
Decimation	32768	15
Hold	65536	16

motion ASD to transform it into an arbitrary length vector. An example of the weighting used for the SC and FF paths is shown in FIG 5.10.

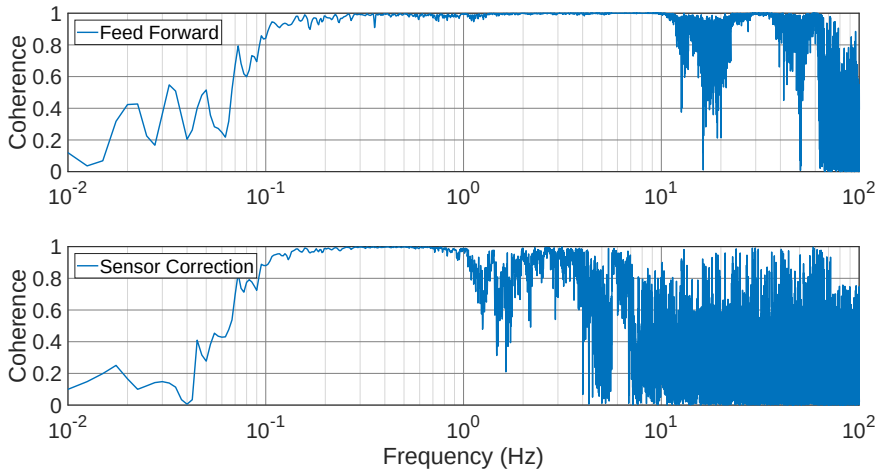

Figure 5.10: Coherence between the ST0 L-4C and GS13 (top) and STS-2 and GS13 (bottom) used in feed forward and sensor correction.

5.2.5 Calculation of Platform Motion

Depending on the signal source, some inputs should be summed in quadrature and others **summed** coherently. All noise propagation terms, collectively known as SN, from the STS-2, CPS, GS13 and L-4C, are summed in quadrature along with the tilt component (only on X and Y). In reality sensor noises have random phases, preventing multiple sources of noise from canceling. In this model, sensor noises do not have any phase information and they are instead summed in quadrature.. The sensor correction path (SC) is used in X, Y and Z only. The ground motion coupling through the CPS path (CPS) and the feed forward path (FF) is used in all degrees of freedom apart from RZ. These ground motion terms are summed coherently together and then added in quadrature with the rest of the terms. The basic equations that describe the estimated stage 1 motion and signal are described in equations 5.4 and 5.5. An overview of how the model calculates the final motion is shown in FIG 5.11.

$$\text{Motion}_{\text{ST1}} = \sqrt{(FF + SC + CPS + GND)^2 + SN^2 + TILT^2} \tag{5.4}$$

$$\text{Signal}_{\text{ST1}} = \sqrt{(FF + SC + CPS + GND)^2 + SN^2 + TILT^2 + GS13n^2} \tag{5.5}$$

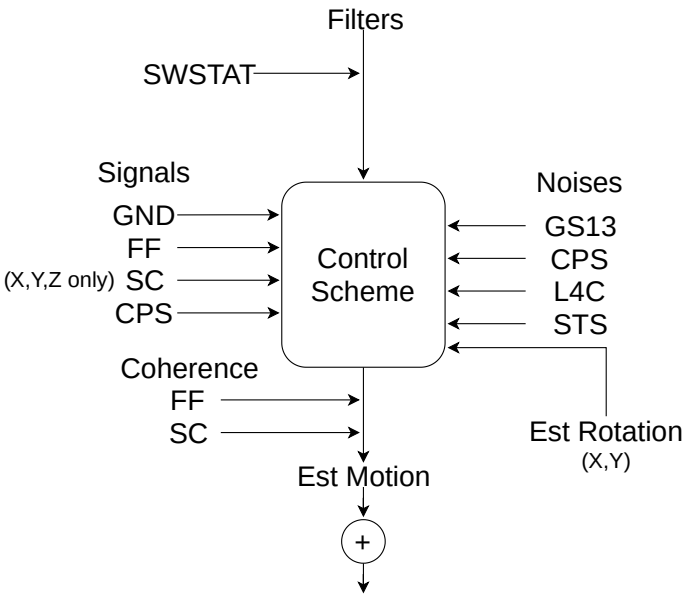

Figure 5.11: How filters, sensor noises and signals are fed into the model's control scheme.

Using the X degree of freedom as an example, as it is the most complicated degree of freedom, we can calculate the expected stage 1 motion in X after it has passed through the control loop, as shown by Kissel in [11]. The final motion as seen by sensors on the platform in the X degree of freedom is given by equation 5.6, the full derivation of which is given in Appendix A.3

$$
\begin{aligned}
x_{\mathrm{ST1}} &= \underbrace{\frac{P_x^{'(0-1)}}{1-\epsilon G_x^{'I}} x_{ST0}}_{Residual\ ground} + \underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} F_x^{FF} x_{ST0}}_{Feed\ forward} \\
&- \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} x_{ST0}}_{CPS\ ground\ injection} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} F_{gnd}^{SC} x_{GND}}_{Sensor\ correction} \\
&+ \left(\underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} K_x^D}_{Damping\ tilt\ coupling} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F^{HP}}_{Isolation\ tilt\ coupling} \right) \frac{g}{\omega^2} RY_{ST1} \\
&+ \left(\underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} K_x^D}_{Damping\ GS13\ noise} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{HP}}_{Isolation\ GS13\ noise} \right) n_{GS13} \\
&+ \underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} F_{ST0}^{FF} n_{L4C}}_{L4C\ noise\ injection} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} n_{CPS}}_{CPS\ noise\ injection} \\
&+ \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} F_{gnd}^{SC} n_{STS}}_{STS\ noise\ injection}
\end{aligned}
\tag{5.6}
$$

This process can be repeated for the other degrees of freedom to calculate the estimated stage 1 motion. X and Y are by far the most complicated of the two degrees of freedom, as the tilt to translation coupling term must be accounted for. To avoid injecting large amounts of GS13 noise into the control loop below around 0.3 Hz we must calculate, for the X and Y, the respective tilt degree of freedom first. The estimated motion, rather than the signal, can be used as the tilt motion term instead of using the measured signal, which at low frequency is dominated by self sensor noise. In reality, the stage 1 sensor noise from the tilt degree of freedom does not couple into the translational degrees of freedom, meaning this method is valid. The tilt degrees of freedom (RX, RY, RZ) have no tilt coupling terms, or sensor correction terms and therefore no STS-2 noise injection.

In the Z degree of freedom there are no tilt coupling terms present. The ground motion terms, i.e. ones denoted by x_{ST0} or x_{GND}, are summed coherently and then summed in quadrature with the rest of the terms, producing the estimated stage 1 motion.

To evaluate the accuracy of the model the estimated signal, rather than the estimated motion, is compared with the measured signal of the stage 1 GS13s. To produce this the estimated motion is summed in quadrature with the sensor noise of the witness sensor, in this case the GS13. This is used instead of the CPS as it is a direct measure of the true platform motion.

5.3 Outputs of the HAM Model

FIG 5.12 shows a typical output of the model when run on HAM5 in the X degree of freedom at LHO. The input data was taken during a long lock stretch during O2 on the 6th August 2017 at 8AM UTC. The full control loop was simulated with all parts of the isolation loop activated. Here the estimated motion (blue) and estimated signal (red) are plotted against the stitched input ground motion (black), the measured GS13 signal (green) and the noise of the GS13 witness sensor (dashed magenta). Overall the model provides a good fit to the measured platform motion with a few notable exceptions at 0.3 Hz, 1 Hz and 30 Hz. The ISI signal is clearly limited by in-loop GS13 noise at frequencies below 0.1 Hz and above 70 Hz. To investigate discrepancies between the model and the measured values, the other main output of the model, the individual noise contribution plot, can be analysed.

FIG 5.13 shows the contribution from each of the individual paths in the control loop shown in FIG 5.4. The ground transmission and feed forward paths have been coherently summed together to form the residual ground trace. In addition, the ST0 input on the CPS path and the sensor correction loop have been summed together to form the CPS ground injection trace. The tilt-to-translation coupling is shown in magenta, taken from the previously calculated ST1 RY motion. Sensor noises are shown in dotted lines to distinguish them from real seismic input motion.

Addressing the discrepancies of the models in order of ascending frequency, the error in the predicted motion at 0.3 Hz and 1 Hz are caused by excessive tilt coupling. Looking at the RY contribution plot, shown in FIG 5.14, we find that the former of these is likely caused by over-predicting the RY input ground motion, as referenced in section 5.2.1. At this frequency the model is likely

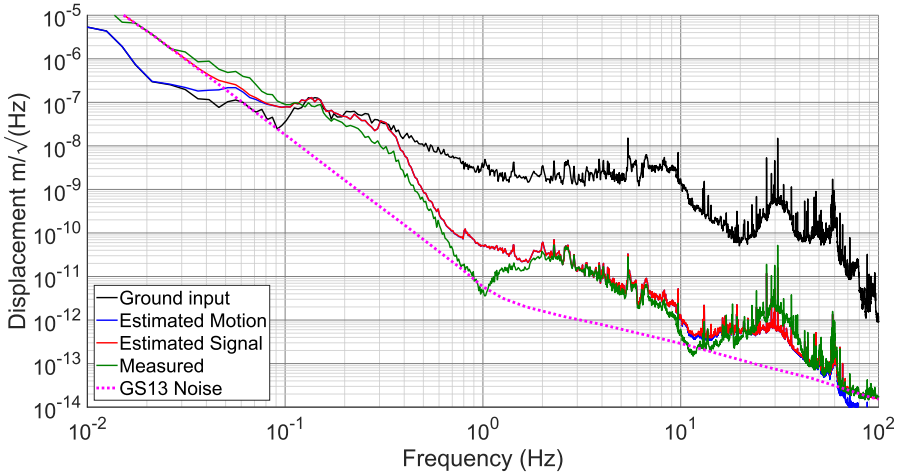

Figure 5.12: Figure showing a typical output of the HAM ISI model when run in the X degree of freedom. This shows the estimated motion (blue) and signal (red) compared against the stitched input ground spectrum (black), the GS13 noise (dashed magenta) and the measured ST1 motion (green).

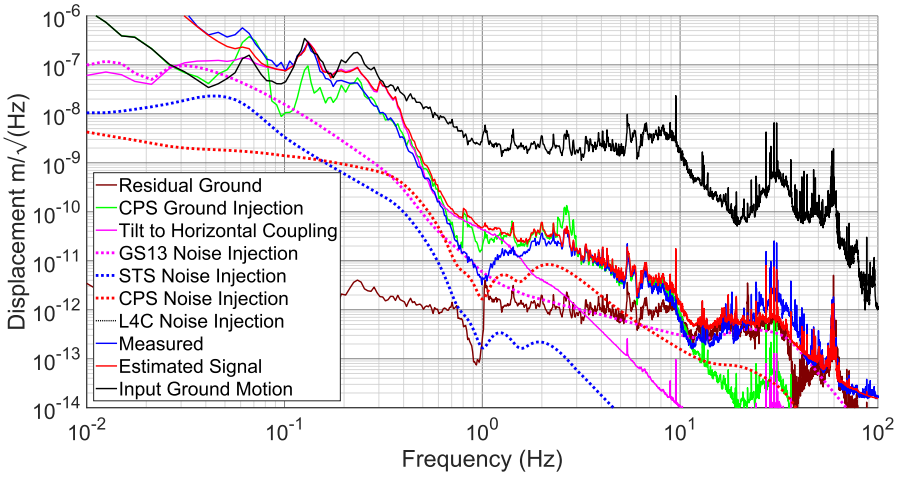

Figure 5.13: Figure showing the individual contributions in the X degree of freedom that sum together to produce the final estimated signal shown in FIG 5.12 highlighting the limiting noise terms at certain frequencies.

injecting L-4C sensor noise incorrectly into the control loop. This can be adjusted by altering the stitching frequency of the ground rotation estimation. The discrepancy at 1 Hz is more complicated, as the estimated motion of the platform agrees well with the measured ST1 motion in RY.

At this frequency the RY motion is limited by the CPS noise. This hints that the in-loop GS13 is suppressing the motion of the platform below the L-4C sensor noise. Such a discrepancy will need to be studied in the future in more detail.

At 30 Hz, the model under-predicts the motion as measured by the stage 1 GS13s. This is likely due to an imperfect fit on either the ground transmission or actuator transfer function, denoted by $P^{(0-1)}$ and $P^{(1-1)}$, respectively, in FIG 5.4. This could be resolved by interpolating over the raw transfer function vectors instead of fitting the data. This will be included in a future version of the model.

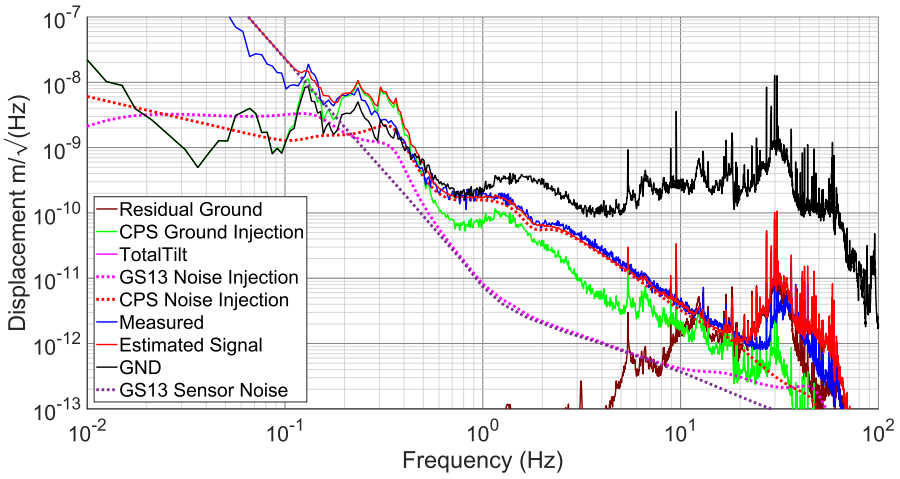

Figure 5.14: Figure showing the individual contributions from each path in the HAM ISI control loop in the RY degree of freedom. Note that CPS sensor noise injection, shown in dashed red is a limiting noise source from 0.6 to 10 Hz.

5.3.1 Calculating Suspension Point Motion

The model can be used to calculate the motion as seen by the suspensions on the ISI. This involves calculating each degree of freedom in Seismic (SEI) co-ordinates (X, Y, Z, RX, RY, RZ) into suspension co-ordinates: length, transverse, vertical, roll, pitch and yaw (L, T, V, R, P, Y). The matrix needed to do this is downloaded from site for each suspension on the selected chamber. An example of this is shown in Table 5.2 and an example of the projected motion is shown in FIG

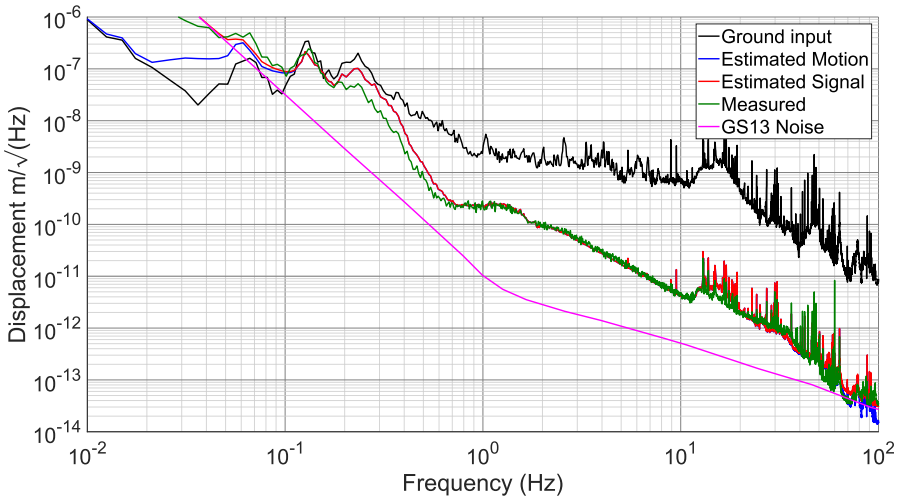

Figure 5.15: Figure showing the suspension point motion projection of the HAM ISI model comparing the measured motion (green) against the model projected signal (red).

5.15.

5.4 Improvements due to new Filters

5.4.1 Reducing SRCL motion

During O1 and O2, the SRCL was close to limiting the performance of DARM [103]. SRCL itself is limited above 10 Hz by shot noise, the controller for the SRCL cavity has a UGF of around 40 Hz. The bandwidth of this controller is limited by the optic motion in the SRC. This system couples shot noise into the SRCL optics above 10 Hz, injecting noise into the real cavity length signal. To eliminate this shot noise coupling into DARM, the bandwidth of the controller must be reduced while maintaining the same RMS performance.

To reduce the bandwidth of the SRCL controller, the suspension point motion of the SRCL cavity needs to be reduced. This in turn is limited by the ISI motion of HAM4 and HAM5 between 0.7 and 4 Hz shown by FIG 5 of [103]. The ISI motion is limited by a combination of CPS ground injection from 1.5 to 4 Hz and tilt to translation coupling from 0.7 to 1.5 Hz. Using the HAM

Table 5.2: Table showing the SEI to Suspension (SUS) co-ordinate transformation matrices needed to project the estimated ISI motion to suspension point motion.

	X	Y	RZ	Z	RX	RY
L	-0.014	0.999	-0.1691	0	-1.0958	-0.0153
T	-0.9999	-0.014	0.4578	0	0.0153	-1.0958
V	0	0	0	1	0.4554	0.1755
R	0	0	0	0	-0.014	0.9999
P	0	0	0	0	-0.9999	-0.014
Y	0	0	1	0	0	0

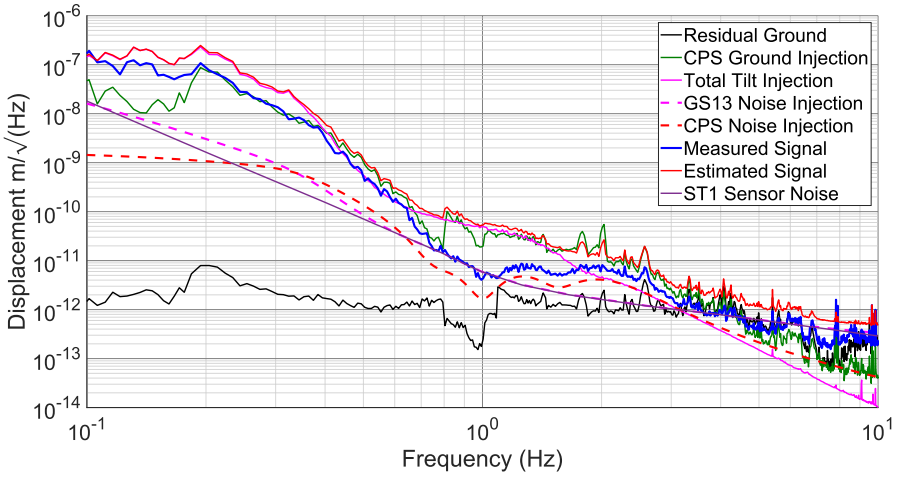

Figure 5.16: Figure showing the individual contributions that sum together to form the estimated motion of the HAM4 ISI in Y, highlighting the limiting noise sources.

model described in [99], the RX and RY motion is found to be limited by CPS noise injection and is shown in FIG 5.17.

FIG 5.16 shows the different contributions from each of the individual paths in the HAM ISI control loop, shown in FIG 5.4. This shows that platform tilt is one of the dominant noise sources at 1 Hz. FIG 5.17 shows contributions for the RX platform motion, in the frequency region of 0.7 to 4 Hz the CPS noise, that is injected through the low pass filter is shown to be a dominant noise source.

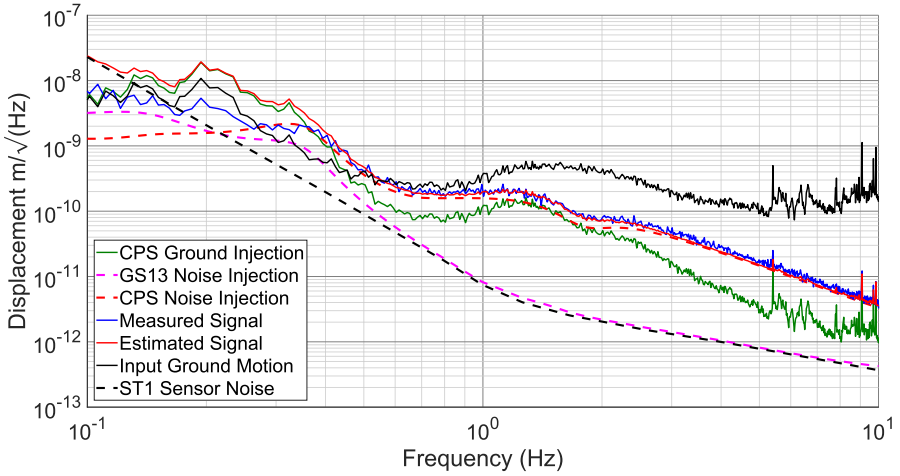

Figure 5.17: Figure showing the individual contributions that sum together to form the estimated motion of the HAM4 ISI in RX, highlighting the limiting noise sources.

Reducing ISI Motion

To reduce the injected CPS noise, we can either change the high and low pass blends in RX to attempt to suppress this noise between 0.7 and 4 Hz, and potentially sacrifice performance at lower frequency, or replace capacitive sensors on the ISI with higher resolution sensors, described in [5]. The reduction in platform motion using both of these options has been evaluated using the HAM ISI model by investigating the suspension point motion of the signal recycling cavity located on HAM4 and HAM5. This investigation required changing blend filters in the RX degree of freedom and propagating this change forward to the Y degree of freedom.

Changing RX Blend Filters

FIG 5.18 shows the current high and low pass blending filters in use at LHO compared with a new set of blending filters designed to reduce the CPS injection in the RX and RY degrees of freedom.

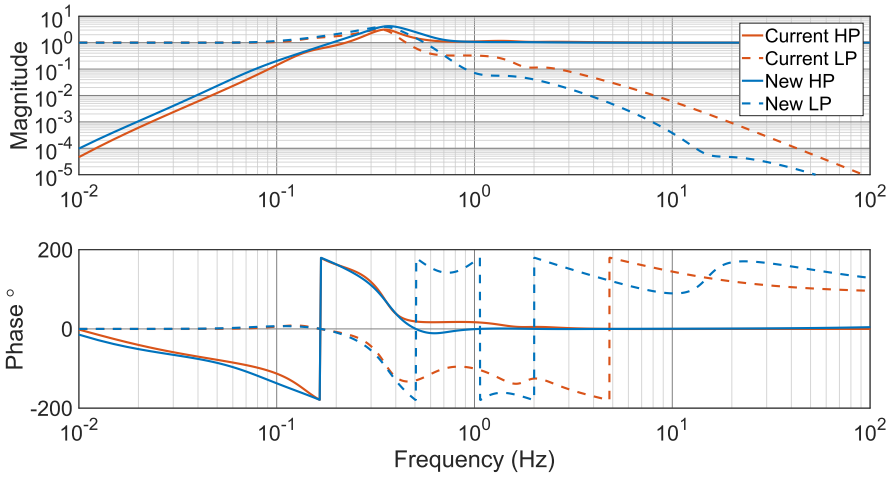

Figure 5.18: Figure showing the comparison of the new (blue) and old (red) low and high pass blends shown by the dashed and un-dashed lines respectively.

Replacing Sensors

Using the model the ADE 0.25 mm, a type of capacitive position sensor, and the HoQI interferometric sensors [5] are evaluated with both the old and new blending filters. These new displacement sensors will function identically to the CPS sensors they replace, as they still measure the difference between platform and ground motion. Hence, we can just directly replace the sensor noise curves with new sensors and project these through the HAM model. FIG 5.19 shows a comparison between the sensors used in this evaluation.

SRCL Noise Projections

The HAM model was run to calculate each permutation of the SR3 suspension point motion on HAM5, using different blending filters and sensors to replace the CPS, shown in FIG 5.18 and 5.19. The model was configured using the parameters shown in table 5.3.

To quantify the reduction in platform motion we have estimated the displacement spectra of the third signal recycling mirror, SR3 on HAM5 in length, as this is the representative coupling into DARM from 10 mHz to 100 Hz, though we have truncated the frequency axis for clarity. The

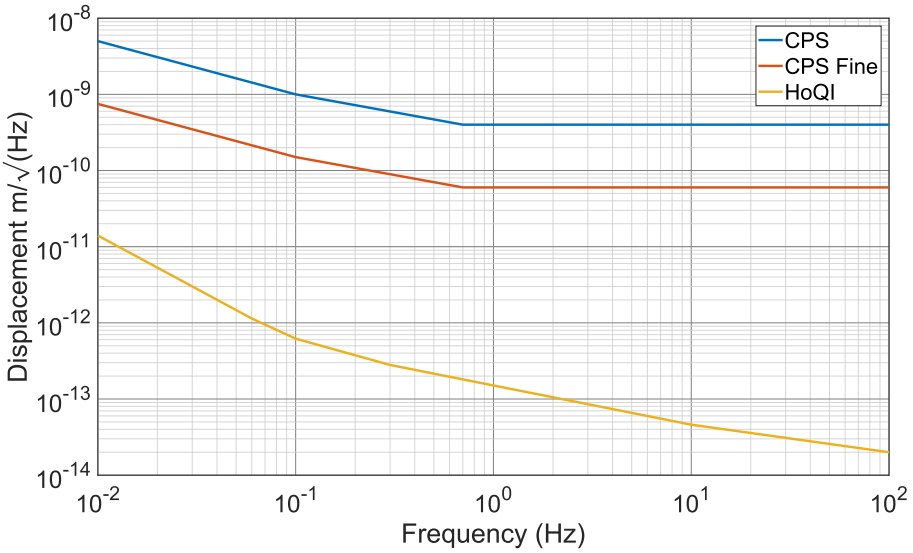

Figure 5.19: Figure showing the resolution of different sensor noises used in this comparison. CPS traces taken from [12], HoQI sensor noise taken from [5].

model was run in two stages, first replacing just the RX sensors and a second time replacing all the CPS sensor paths with sensors described above. The results of this are shown in FIG 5.20. The measured and estimated platform motion using current sensors is shown in blue and red, respectively. The potential improvement due to replacing only the RX blending filters is shown in yellow. This filter change yields a factor of 5 reduction in the length suspension point motion between 0.7 and 4 Hz. This has the side effect of increasing the motion between 0.4 and 0.7 Hz by around 60%, due to the altered blend frequency. Replacing the displacement sensors in RX prevents the increase in motion in this frequency range, while reducing motion by a further factor of two between 0.7 and 4 Hz. Fine CPSs (the ADE_0.25mm) are just as effective as compact interferometers when placed on the ISI using these blending filters, as the platform is no longer limited by sensor noise and rather by residual ground motion injection.

Table 5.3: Parameters that were used to generate the noise projections in the HAM model

Parameter	Value
IFO	H1
Chamber	HAM5
Suspension Name	SR3 - L
Start Time	1185876018
Duration	3600 (s)
Data Rate	256 Hz
Control loops activated	All
Tilt estimation	On
Calculate suspoint	On
Ground L4C's	Stage 0

5.5 Isolation Improvements Due to New Sensors

The model can be used to investigate potential improvements to the platform's isolation performance by replacing the currently used sensors, namely the L-4C, GS13 and the CPS, with state of the art interferometric ones, described in chapters 4 and 2 respectively. The damping and isolation filters, and therefore their corresponding loop gains, will be kept the same as what is currently used on site. Doing so will give a more realistic prediction, as increasing the loop gain would further suppress the motion than what is possible in reality. The high and low pass blending filters will have to be changed to more effectively utilise the new sensors. For this comparison we shall replace the L-4C and GS13s with their interferometrically readout counterparts using projected sensitivities, taking the noise floor measured in chapter 2 and modifying the response of the L-4C and GS13 and summing this in quadrature with the suspension thermal noise to create predicted sensitivities for these devices.

We shall vary the displacement sensors that are used, creating a set of blend filters for use with

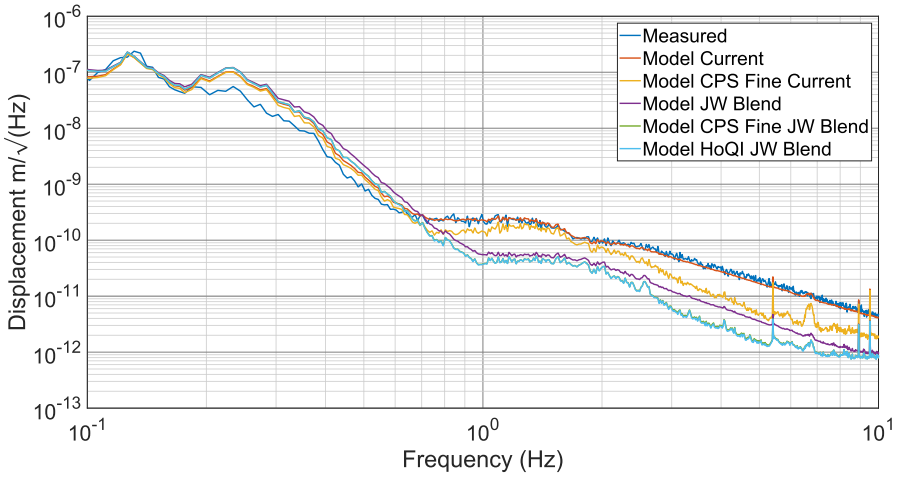

Figure 5.20: Figure showing the estimated reduction in platform motion by changing first the blend filters and then CPS sensors on HAM5 and propagating these changes to calculate the expected suspension point motion with these changes. The blue and red trace show the measured (from site) and estimated (from the model) suspension point motion using current blending filters and sensors. The yellow trace is the result of changing the blending filter alone, while the green and purple traces are from changing both the blending filter and the displacement sensors used.

the CPS and another with HoQI. To create the blending filters, the velocity RMS of the sensor noise contributions when passed through the respective blending filters has been minimised. In doing this we assume that the isolation loops used in LIGO are not gain limited. For simplicity we use the Z degree of freedom to eliminate the need to account for tilt cross coupling. Equation 5.7 shows how the total noise injection term is calculated by summing the inertial (GS13) and position sensors (CPS) when multiplied by the high and low pass blend filters respectively.

$$\text{TotalNoise} \quad = \text{LPB} * \text{CPS} + \text{HPB} * \text{GS13}. \tag{5.7}$$

5.5.1 Creating Blending Filters

To simplify the process of creating blends, especially if we want to test sensors with different responses and noise floors, binomial blends are used as an input to equation 5.7. Here the CPS and GS13 denote the displacement and interferometric inertial sensors noises that are used to generate the filters. The binomial filters are complementary, such that the summation of the two blending filters across the whole frequency band is equal to one, and have a customisable roll-off and corner frequency.

We then convert this to velocity by multiplying by $2\pi f$, calculate the RMS and use the combination of high and low pass filters that produce the lowest overall sensor noise. We can afford to brute force this problem, as the sensor noise curves are not very complex, nor are they large vectors; this is not the case with seismometer data. In this process we assume that LIGO's seismic isolation is not gain limited in at least some part of the frequency band. If this were not the case, even with new sensors, there would be little improvement to the predicted isolation performance. This process is repeated for a range of different sensors and blend filters. To ease computation, the range of corner frequencies has been bounded between 10 mHz and 1 Hz.

An example of the blends used in this process are displayed in FIG 5.21, which shows the low and high pass blending filters used to estimate the benefit of using optical GS13s instead of the coil GS13s in the Z degree of freedom. If we were to do this process manually we would vary the blending, or corner frequency of the blending filters, and choose a roll off such that the sensor noise contribution would be small compared to the ground motion. For example, if we were using L-4Cs with CPS sensors, our high pass filter would need a roll off as $\frac{1}{f^3}$ to avoid injecting sensor noise at low frequency. Conversely the low pass blend would need to have a roll off of at least $\frac{1}{f^2}$ to avoid injecting CPS noise at high frequency. The corner frequency would be chosen to ensure that the amplitude of the filter in question was sufficiently low to suppress sensor noise.

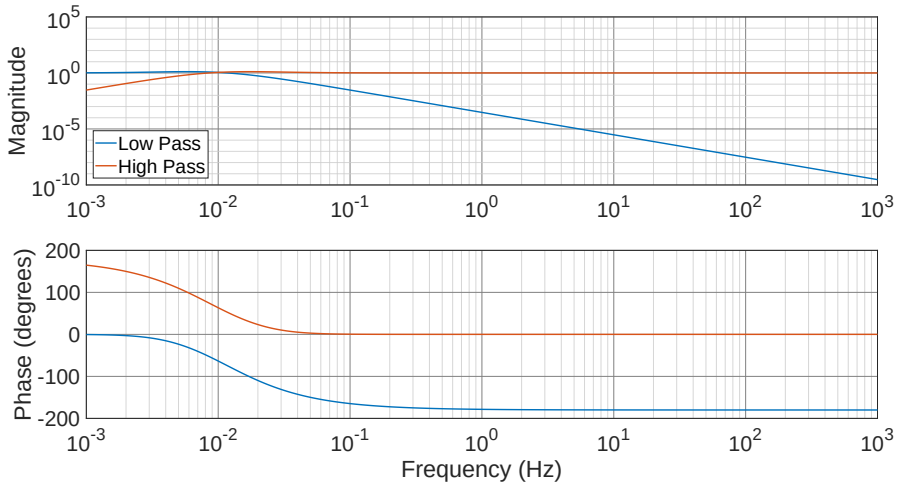

Figure 5.21: Figure showing the binomial blends that were used to estimate the benefit of using optically readout GS13s instead of the conventional coil read-out sensors.

5.5.2 Replacing Inertial Sensors

We shall use the Z degree of freedom as an example. This is the simplest degree of freedom to consider as there is no tilt cross coupling and therefore can be considered as a standalone case. The model is run on an hour long stretch of data during a locked segment in the second observing run on August 6th 2017 8AM UTC. The high and low pass filters in the Z degree of freedom are replaced with the newly designed binomial blend filters. All other filters, including the damping and isolation filters, are left at their default values. For sensor noise inputs we used the 1 mm CPS sensor as our displacement sensor and an optical GS13 as our inertial sensor, taking the predicted noise from the sensor shown in chapter 2.

To compare the ISI performance using different sensors, we have elected to show the predicted motion using the HAM model with the current sensors instead of using the measured motion. This is due to the sensor correction path is over-predicting the amount of subtraction that can be gained from using sensor correction, despite being weighted by the coherence between the STS-2 and the GS13. Using the estimated motion gives a clearer indication in the potential isolation gain that can be achieved by comparing the two predicted traces. The red trace shows the estimated motion in Z using the current sensor noises and filters as inputs. The blue trace shows the estimated motion

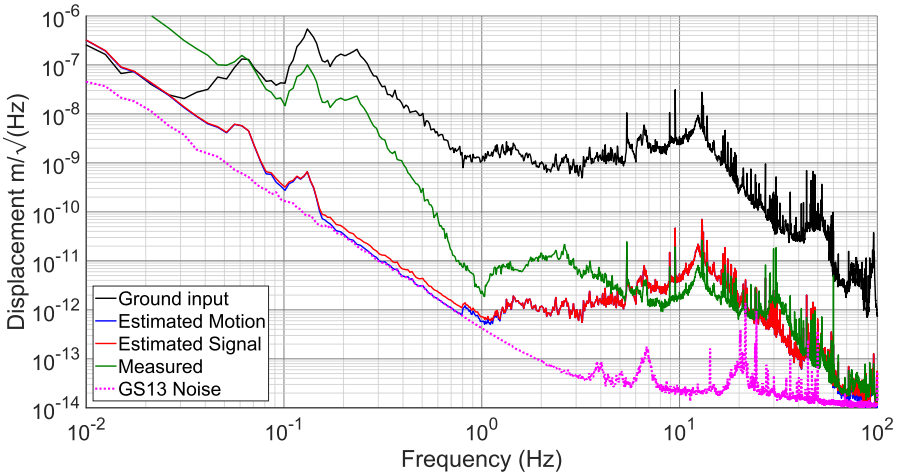

Figure 5.22: Figure showing the predicted HAM5 Z motion when replacing the coil GS13 with optically readout GS13s described in chapter 4. Here we find a significant reduction in the microseismic motion and RMS when replacing these sensors.

by replacing the L-4Cs, GS13s and CPS sensors with the interferometric counterparts with their respective RMSs shown in the appropriate dashed color. By doing this we manage to reduce the motion caused by the microseismic peak at around 0.2 Hz by a factor of 100, reducing the RMS significantly in this range. Shown in magenta is the readout noise of the optical GS13 which is now the limiting factor between 0.2 Hz and 0.8 Hz. Unsurprisingly replacing the inertial sensors did little to improve the ISI performance at 10 Hz and upwards where the isolation loop gain is low.

This reduction in the Z motion as seen by the ISI will likely improve the expected isolation in RX and RY if optical levers are used on the HAM ISIs, as optical levers are limited by differential Z motion below 0.5 Hz [104].

5.6 Improvements to Suspension Isolation

If gravitational waves are to be detected below 10 Hz, sensors on the suspensions will need to be replaced with sensors with significantly lower noise [47]. Currently, shadow sensors, named

BOSEMs are used throughout the detector to measure the positions of each suspension stage, with the exception of the lowest stage of the quadruple suspensions. In addition to applying local control, the BOSEMs are used to damp suspension resonances, reducing the RMS in the 0.5 Hz to 5 Hz frequency band, significantly improving controllability of the detector. Due to the noise of these sensors at 10 Hz and above, damping can only be applied to the uppermost stage of the suspension [105]. Sensors with a noise floor 100 times lower than the BOSEMs can be used to damp not only the top stage of the suspension, but also the UIM [79].

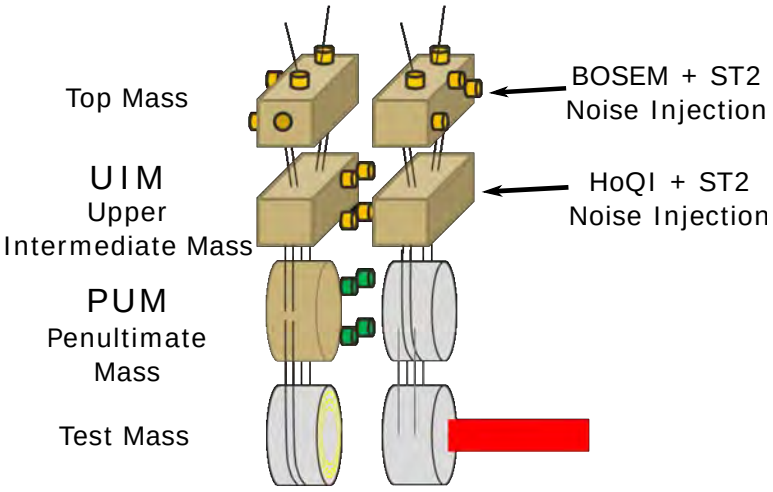

Figure 5.23: Figure showing the layout of the QUAD suspension and where sensor noise is injected with the different position sensors to compare the effect of sensor noise injection on the test mass. Green and yellow sensors are where the BOSEMs, or BOSEM-like sensors are located. Figure adapted from [13].

To examine the benefits of using interferometers to sense the suspension motion, we shall use the quad suspension model as presented by Shapiro and Bonilla [94] and compare the noise injection caused by BOSEMs and HoQIs detailed in chapter 2. FIG 5.23 shows how sensor noise is combined with stage two suspension point motion on the chamber ITMY. In the case of the BOSEMs, stage two motion suspension point motion is summed in quadrature with the BOSEM noise, shown in [8], and injected into the top stage of the quadruple suspension. In the case of the interferometer, stage two motion is summed in quadrature with measured HoQI noise, shown in [5] and injected

at the upper intermediate mass and propagated through the quad model. Our comparison point will be the noise injected in both cases at 10 Hz as seen by the test mass, in the length and pitch degrees of freedom, as this is a critical frequency for noise injection.

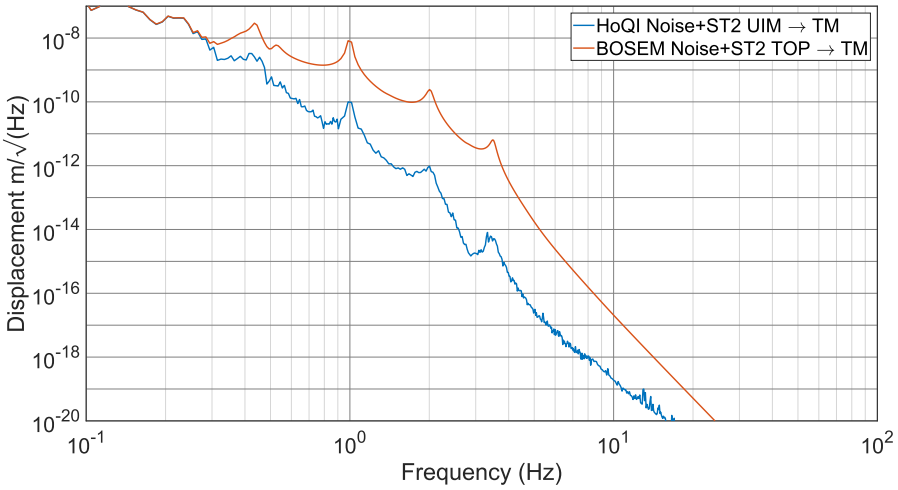

Figure 5.24: Figure showing the noise injection by summing HoQI noise in quadrature with stage 2 length suspension point motion from the UIM (blue) and injecting BOSEM noise from the top mass and projecting the effects to the test mass length (red). The BOSEM noise is a stick-figure noise curve based on the values taken from [8], the HoQI noise is taken from chapter 2.

FIG 5.24 shows the current projected noise seen at the test mass caused by BOSEM noise being injected at the top stages of the suspension, along with stage 2 motion of the BSC ISI. Damping is provided by the current BOSEM damping filters. This is compared with HoQI noise which is summed in quadrature with the stage 2 motion of the BSC ISI from the UIM and projected onto the test mass in length, again with BOSEM damping turned on. From this we see that interferometers, despite their one stage fewer of passive isolation, would cause a factor of 100 less noise injection compared with the BOSEMs at 10 Hz in the length degree of freedom as seen by the test mass.

FIG 5.25 shows the same noise injection as FIG 5.24 but projecting the pitch motion of the top and UIM mass to the test mass using BOSEMs and HoQI respectively. From this we see that interferometers would cause 35 times less noise injection at 10 Hz in pitch. This improvement is smaller than in the length degree of freedom as the stage 2 pitch motion is more significant at

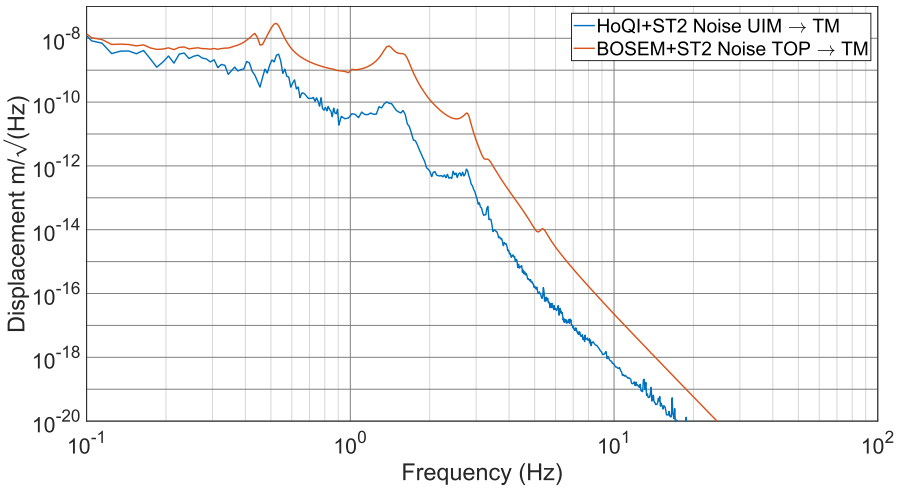

Figure 5.25: Figure showing the noise injection by summing HoQI in quadrature with stage 2 pitch suspension point from the UIM (blue) and injecting BOSEM noise from the top mass and projecting the effects to the test mass pitch (red).

10 Hz, reducing the benefit of using interferometers at the test mass. The design of filters for damping at M2 of the beamsplitter suspension has been undertaken by Mow-Lowry et.al [106] for the A+ upgrade of Advanced LIGO. The same method can be applied to designing damping filters to reduce the resonances of the quad suspension.

5.7 Summary

To quantify the performance improvement that the sensors in chapter 2 and 4 will have on gravitational wave detectors, a modular, predictive model of a LIGO HAM ISI has been created in MATLAB. This model uses real seismic data and filters that are downloaded from each of the sites to produce a live noise budget of stage 1 platform motion. This model can be run on either of the two LIGO sites and generates an estimate of the closed loop motion and measured signal using only out of loop signals and sensor noises. As the model uses only out of loop sensors to estimate this motion, control filters, sensor noises and loop gain can be changed offline to estimate the impact on platform motion.

This model has been used to estimate the improvements that could be possible if seismometers such as the L-4C and GS13 are upgraded to have an interferometric readout. By changing the high and low pass blending filters, while keeping all other filters at their default values, the performance enabled by the interferometric sensors can be predicted. When the model is run using these parameters, the motion caused by the microseismic peak at approximately 0.2 Hz is reduced by a factor of 100, reducing the RMS significantly in this range. At higher frequencies, the expected performance improvement is lower as the loop gain decreases, such that at 2 Hz the estimated motion is a factor 10 lower than the present sensors.

Due to the modular nature of the model, the estimated suspension point motion of each of the ISI's can be calculated. During the commissioning time before O3 the SRCL signal was close to limiting the performance of DARM [3] and the RMS motion of HAM4 and HAM5 needed to be reduced. Using the noise budgeting ability of the HAM ISI model, the platform tilt was found to be the underlying noise source that was casing this issue. This rotational motion was found to be limited by CPS noise and as such the blend frequencies in the rotational degrees of freedom have been altered to reduce the CPS noise coupling by a factor of 5 between 0.7 and 4 Hz. This change has been implemented for the third observing run.

The use of HoQI as a sensor to measure the motion of stages of the quad suspension has also been evaluated. Here HoQIs on the second stage, the upper intermediate mass, are compared against BOSEMs on the top stage of the suspension. This is possible due to the increased resolution of when compared to BOSEMs. To ensure a fair comparison, the stage 2 motion that would couple through to each stage has been summed in quadrature with sensor noise of that stage. E.g. HoQI noise has been summed in quadrature with residual platform motion that couples into the UIM. Using HoQI as a sensor to measure the motion of the quad Suspension at the UIM results in approximately a factor 60 less noise injection at 10 Hz, providing a significant reduction when compared to the BOSEMs. This reduction in noise injection will allow for increased damping of the suspension resonances as shown in [106].

Chapter 6

Particle Swarming of Sensor Correction Filters

I'll be honest, we're throwing science at the wall here to see what sticks.

No idea what it'll do.

Cave Johnson

This chapter describes the development of a set of scripts used to perform particle swarm optimisation of sensor correction filters that are used in isolation control loops on the LIGO sites. This section is based on information published in a technical note [107] of which I was second author. My contribution to this work was in developing the infrastructure to the swarming scripts, allowing them to run on each LIGO interferometer across all chambers and degrees of freedom. Text from this tech note has been copied, re-written and exanded upon for further analysis and discussion, figures have been re-made.

6.1 Overview of Sensor Correction

Sensor correction is a feed forward technique used as part of the ISI control loops, where the signal from the STS-2 seismometer located on the floor is added to the CPS, in the aim of removing ground motion from the CPS. As discussed in chapter 5, the CPS sensor measures the difference between the platform motion and the ground motion, such that, in the ideal case with no tilt coupling, the sensor corrected CPS, (CPS_{SC}) is written as,

$$CPS_{SC} = \mathrm{CPS} + \mathrm{STS}, \tag{6.1}$$

$$CPS_{SC} = x_p - x_g + \mathrm{SC} * \mathrm{x_g}, \tag{6.2}$$

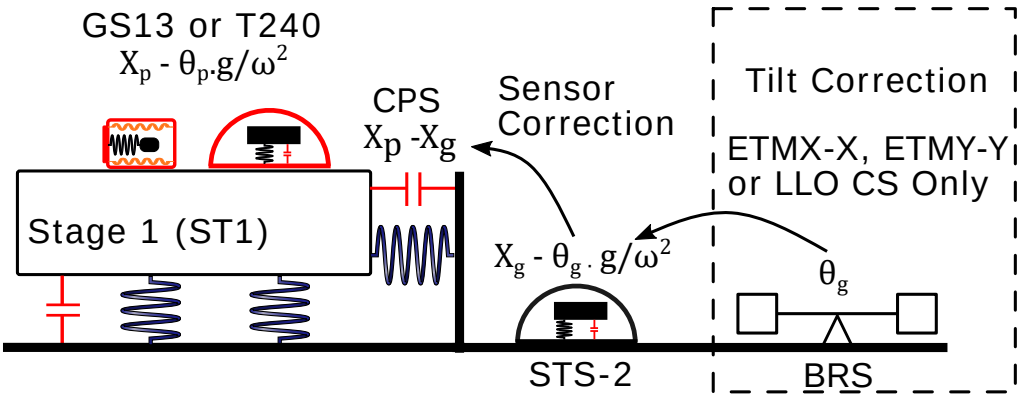

Figure 6.1: Figure showing the signal path of tilt and sensor correction. Tilt correction is used at LIGO Livingston (LLO) on all chambers, while it is only used at LIGO Hanford (LHO) on the end stations.

where CPS_{SC} is the sensor corrected CPS sensor, x_p is the measured platform motion and x_g is the measured ground motion, $*$ represent convolutions between objects. At low frequencies, as discussed in chapter 5, the ground tilt is indistinguishable from horizontal movement and as such we must modify the equation above to include this ground tilt term. This means that by adding in ground translation, we must also inject ground tilt, thus increasing the platform tilt at low frequency. The ground tilt injection couples in with a factor of $\frac{g}{\omega^2}$.

$$CPS_{SC} \quad = \quad x_p - x_g + \mathrm{SC} * (\mathrm{x_g} - \frac{\mathrm{g}}{\omega^2}\theta_\mathrm{g}). \tag{6.3}$$

As the primary ground translation occurs between 0.1 and 0.3 Hz and ground tilt is responsible for the majority of the measured STS-2 signal below 0.05 Hz [75, 108], the contribution of the tilt measured in this frequency band is strongly dependent on the local wind speed. A set of complementary blending filters can be created to minimise the injected ground tilt and maximise the ground translation that is added to the CPS, suppressing overall platform translational motion. In the ideal case, the designed filter would be infinitely steep, to roll off the tilt injection term quickly. However, as the filter has to run on a live system, we are limited to the roll off that is achievable to generate a causal filter. The equation to calculate the sensor corrected CPS signal, or the residual ground motion injection is therefore[1],

$$CPS_{SC} = x_p - \underbrace{(1 - SC) * x_g}_{Ground\ Injection} + \underbrace{SC * \frac{g}{\omega^2}\theta_g}_{Tilt\ Injection} . \tag{6.4}$$

At frequencies above 1 Hz the coherence between the ground STS-2 and inertial sensor on top of the ISI decreases, as such the translation subtraction contribution must be attenuated at high frequency to avoid injecting additional incoherent ground translation into the platform. The self noise of the STS-2 will also couple in, though, as this is much smaller than the residual CPS ground injection shown in FIG 5.13 in chapter 5, it is ignored for this purpose.

6.1.1 Motivation Behind Optimising and Automating Filter Design

Design of these filters is mainly done by people working at each one of the two LIGO sites and can take many hours to design and test by tweaking the filters pole/zero frequencies and the associated

[1] In the swarming process the tilt contribution term is calculated exclusively in the frequency domain and summed in quadrature with the translation terms which are evaluated in the time domain and then converted into the frequency domain when evaluating the cost of a particular filter.

quality factors to reduce noise injection and increase suppression of motion. While these filters are good enough to detect gravitational waves, these filters may be improved, allowing for improved controllability of the detectors or reducing the ground motion injected into the ISIs.

These filters are designed around single sets of environmental conditions, namely the ground motion and local wind conditions. The microseismic peak, located around 0.1 to 0.3 Hz is responsible for the largest contribution in the translation RMS, where as the ground tilt is strongly linked to the local wind speed that acts on the buildings tilting the floors and seismometers. Therefore it is clear to imagine conditions where one of these effects would need to be prioritised over the other. For example, at LHO, it is not uncommon for the wind speed to exceed 20 mph causing a significant tilt contribution measured by the STS-2 [109]. In such conditions a sensor correction filter, designed around high levels of microseismic motion could risk injecting tilt motion into the platform.

With automated sensor correction design, and a suitably defined cost function, filters can be designed specifically for a range of environmental conditions. In the future these filters could be switched between, ensuring that the optimum filter that most closely matches the current environmental conditions is always active, increasing the controllability of the detector and thus potentially increasing its uptime.

6.2 Particle Swarm Optimisation

Particle swarming is a optimisation routine used for exploring parameter spaces with a large number of dimensions to find a global minimum based on a cost function describing the problem. The method works by testing a large number of potential solutions, or particles, these are scattered randomly throughout the parameter space. Each particle then evaluates the given mathematical cost function and returns its cost to the main program. After each iteration the particles are given a velocity, moving them towards the previous global minima and a random velocity to search more of the parameter space. This method of biasing the search area around known areas containing minima excludes large areas of the space allowing the swarm to converge on a solution faster.

6.2.1 Input Data

To evaluate the performance of these filters, the RMS velocity of equation 6.4 is calculated from 10 mHz to 4 Hz. The upper frequency is set to this value due to the low pass filter attenuating the sensor correction filter above around 0.5 Hz, at 10 mHz the contribution to the overall RMS velocity will be small. During observation times the ISI is in an isolated state, meaning the isolation loops and all feed forward control paths are activated. To optimise the sensor correction filters, the ISIs need to be in a damped state, meaning that isolation loops and all feed forward loops are disabled. This results in the inertial sensor on the isolated platform measuring x_p, the measuring CPS $x_p - x_g$ and the STS-2 being out of loop measuring only the platform and ground motion respectively. Data was chosen from the winter break of the second observing run between 23rd December 2016 and January 2nd 2017. Data was checked to ensure the interferometer was in the right control state, by checking the state of the ISI control loop over the chosen data stretch and to check that no larges spikes, caused by earthquakes, were present in the data. Not performing these checks could have resulted in measurement sensors saturating, in the case of earthquakes and thus training the swarm sub-optimal data.

FIG 6.2 shows how data is passed around the particle swarming scripts. First, the input data is filtered to account for the response of the sensors where applicable and converting all data into velocity. Any filter transients are removed after this process by cutting the first 400 seconds of data, afterwards the data is down sampled via Fourier transform to 8 Hz to increase the speed of the optimisation process. This is done outside of the swarm for speed and is loaded only once in the swarming process for the requested degree of freedom and chamber. If no data or current sensor correction filters exist, these are downloaded from the requested site and processed accordingly.

The main swarming script then calls swarmPrep.m, this creates a matrix of boundary conditions limiting the parameter space the swarm can access. A ground cost weighting is calculated to take into account the response of the quad suspension at high frequencies and sets the ground weight to zero in regions where the ground tilt is the dominant cause of motion. This weighting

is passed to the heart of the swarm, which generates sensor correction filters to be evaluated and calculates their associated cost. The particle swarm continues to generate new filters until one of two exit conditions are satisfied: either the maximum time is exceeded or the improvement over a set number of iterations is less than the user defined function tolerance. The maximum time for the optimum filter to be designed depends strongly on the number of CPU cores allocated to the problem. For a dual core laptop an optimum filter will be designed after 1-2 hours, on a 64 core workstation, filters can be designed after 5-10 minutes, highlighting the parallelism of the particle swarming process.

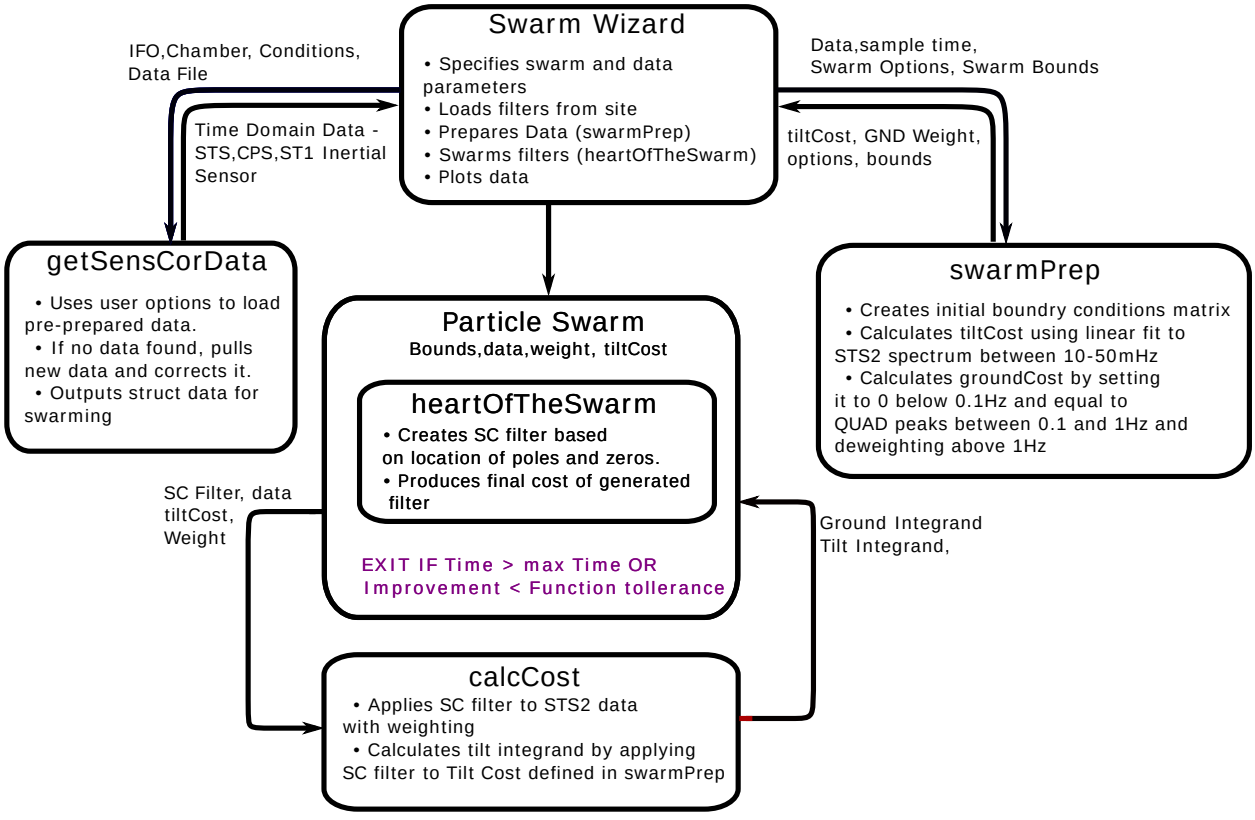

Figure 6.2: Figure showing the layout of the particle swarm optimisation script.

6.2.2 Generating a Cost Function

Optimising equation 6.4 would not be representative of motion seen by the ISI, for example in some chambers and degrees of freedom there is no accurate measure of ground rotation, and thus the

tilt injection is hard to quantify. Moreover, in all degrees of freedom the swarm has no information about the increase in RMS caused by the quad resonances, both of these issues must be addressed to increase the accuracy of the cost function.

In some cases, such as the end stations of LHO and the entirety of LLO, a BRS described by Venkateswara [75] can be used to directly measure the ground rotation. These BRSs are located as close to the ground STS-2 as possible to maximise the coherence between the tilt seen by both devices. The BRS signal can be subtracted from the STS-2 reducing the tilt injection that needs to be minimised by the sensor correction filter, shown by FIG 6.1.

In chambers and degrees of freedom where there is no ground rotation sensor, we must make some estimation of the local ground tilt to correctly weight the swarm from injecting ground tilt into the platform. By looking at STS-2 data during times with different wind speed, there is a correlation between wind speed and motion in the frequency band 50 mHz to 10 mHz. Therefore the tilt cost is calculated by generating a linear fit of the STS-2 data in this frequency range, an example of this fit is shown in FIG 6.3.

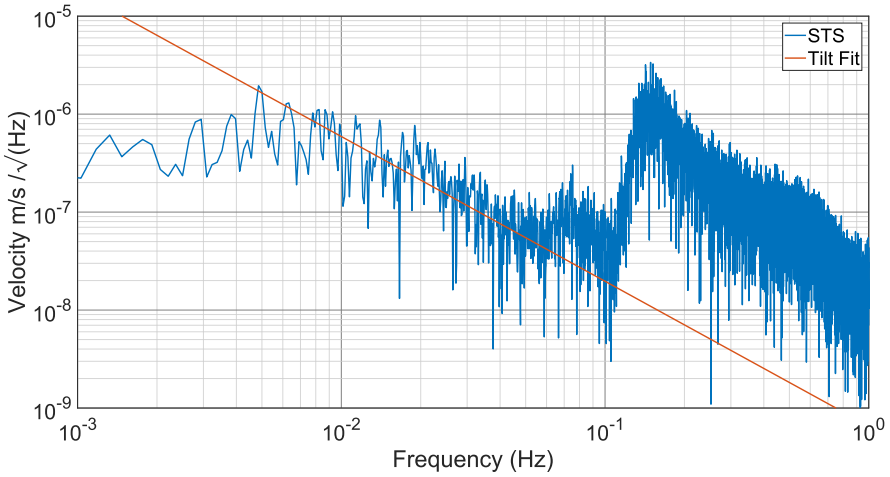

Figure 6.3: Figure showing the tilt estimation on the ground STS. Here we assume that the motion between 0.01 and 0.05 Hz is caused entirely by tilt coupling into the seismometer. Data taken from a 2 hour period where the wind speed was 7 m/s at LHO.

At high frequencies the swarm needs to be weighted to encourage subtraction of ground translation,

rather than focusing solely on reducing ground tilt injection. To do this, at high frequencies the swarm is weighted by the transfer function of the quad suspension, to prevent the optimisation routine from inadvertently amplifying resonances of the quad. This has been simplified from the full quad model, used in chapter 5 to improve the speed of the optimisation routine. The ground injection weighting is shown in FIG 6.4.

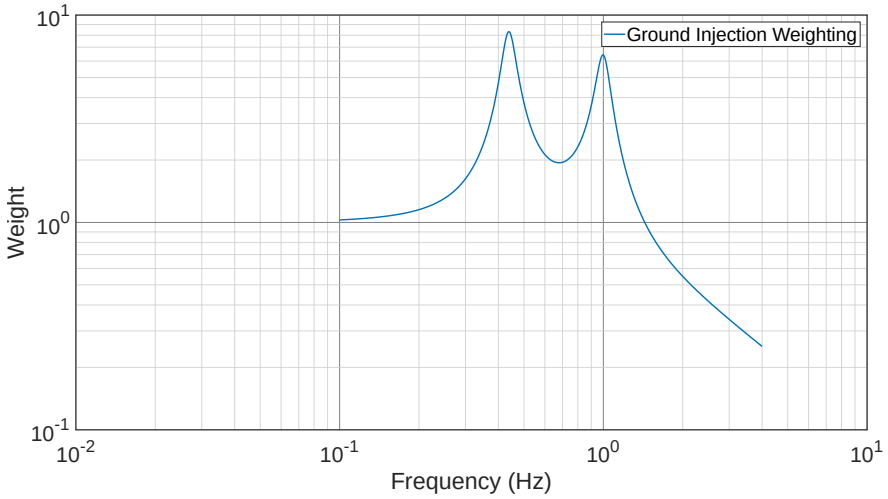

Figure 6.4: Figure showing the ground injection weighting, at frequencies below 0.1 Hz the ground injection cost is set to 0, as any injection is thought to be caused entirely by ground tilt injection.

The ground injection cost is calculated by equation 6.5, once this has been evaluated this is converted to an amplitude spectral density where it is multiplied by the ground injection weighting shown in FIG 6.4. This is summed in quadrature with the estimated tilt cost, comprised of the STS-2 signal between 0 and 50 mHz and the tilt fit, shown in FIG 6.3 above 50 mHz. The ground injection is given by the equation,

$$GND_{\mathrm{inj}} = \mathrm{STS}_{\mathrm{sc}} + \mathrm{CPS}_{\mathrm{ST1}} - \mathrm{IS}_{\mathrm{ST1}}, \tag{6.5}$$

where the STS_{SC} is the ground STS-2 output filtered by the sensor-correction filter, CPS_{ST1} is the stage 1 CPS that measures the difference between the platform and the ground and IS_{ST1} is the inertial sensor stage 1 of the isolated platform. Another consideration to be made is the number

of poles and zeros that can be used in the optimisation process. For simplicity of installation the designed sensor correction filter should not use more than 4 complex poles and 3 complex zeros in order to fit into a single LIGO filter bank, a single real pole and two real zeros are also used to shape the filter. The designed filters can then be evaluated using other sets of input motion to ensure the filters perform well in a range of different environmental conditions.

6.2.3 Optimising the Swarm

To reduce the dynamic range required in the optimisation process, the final cost function is calculated in units of velocity, as such the CPS and GS13, the latter only present on the HAM ISI chambers, need to be converted into velocity. This presents a problem with the GS13 signal, as its response requires it to be integrated twice below 1 Hz leading to large filter transients, spoiling the quality of the input data. As such, the GS13 response is filtered by a ZPK system in MATLAB, the filter used is shown in FIG 6.5. This filter has to balance getting a gain of 10^4 at 10 mHz along with the appropriate phase and minimising any filter transients produced - an example of this is shown in FIG 6.6. To remove the remaining filter transient, the first 400 seconds of data is cut from the input data to remove the effects caused by these transients. After this process, the data is downsampled via Fast Fourier Transform (FFT) to 8 Hz to further increase the speed of the optimisation. Every downsampling process used in the swarming scripts, including acquiring data from site is done via FFT, eliminating errors close to the Nyquist frequency when decimating and downsampling using other methods.

To improve computational efficiency of the optimisation routine a number of tweaks were made to the initial parameterisation of the swarm to avoid degenerate solutions to the problem. This arises as the order of the poles (or zeros) can be swapped, resulting in the same shape of filter and the same RMS value. To break this degeneracy, instead of specifying poles and zeros with frequencies, only the highest value and the log of the differences from this value are specified, reducing the degeneracy by a factor of $num_{poles}! \times num_{zeros}!$. In the final optimisation, 4 poles and 3 zeroes were used increasing the speed of the simulation by a factor of 144.

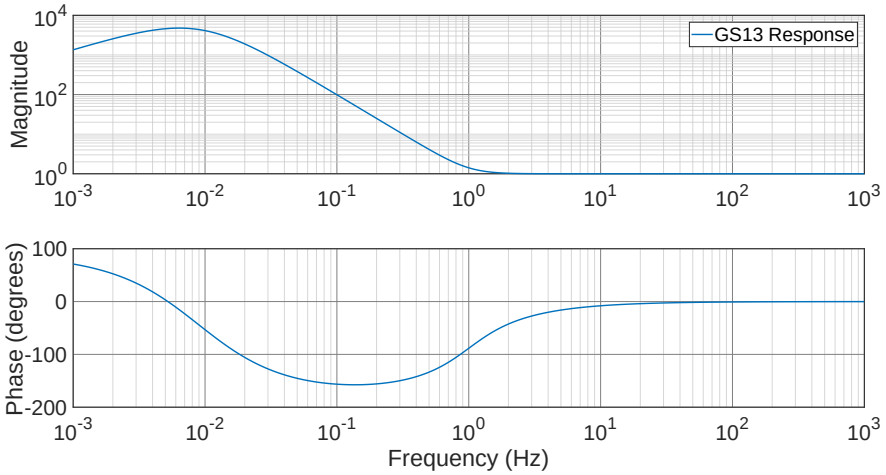

Figure 6.5: A bode plot of the ZPK system used to correct for the GS13s response in velocity, the filters response was balanced between its gain at 10 mHz and the filter transients caused by the data, the latter is shown in FIG 6.6

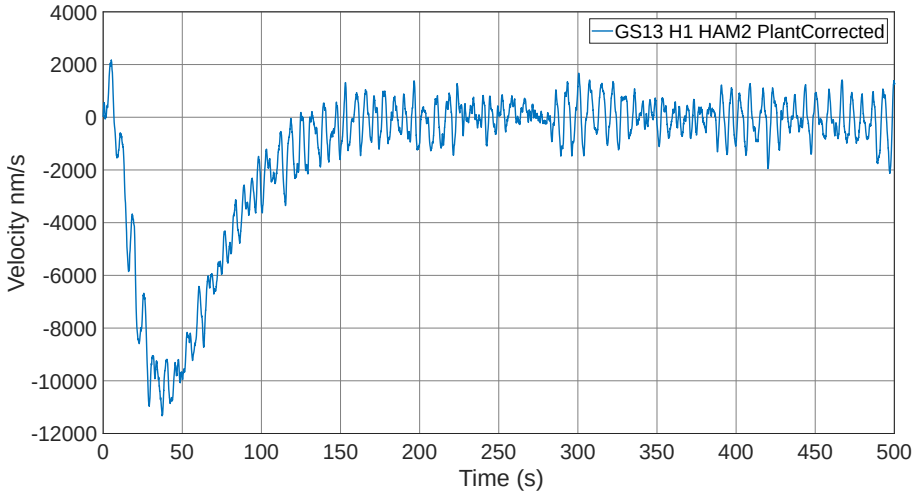

Figure 6.6: Plot showing the effects of the filter transient on the input data, caused by applying the GS13 inverse response filter shown in FIG 6.5.

Table 6.1 shows a typical output of the swarming process, at the time the swarm was designing the corner station sensor correction filter in the Y degree of freedom. Highlighted in bold text is the average cost of the swarm for that iteration, along with the best cost for each iteration. During the optimisation the best cost converges to a minimum cost, though the average cost will

not necessarily do the same as this is dependent on the shape of the parameter space. For instance if there are minima located close to areas of large cost, as the swarm explores the minima of the parameter space, some of the particles will evaluate areas of extremely high cost due to the random velocity each particle is given after each iteration.

6.3 Swarming Results

The optimisation routine was left to run on a 64 core workstation, with the number of particles set to 3000, the maximum time set to 1800 seconds and the function tolerance set to 10^{-12}. The maximum number of stall iterations where, the global improvement can be below the function tolerance is set to 10. Each filter bank in CDS can hold a maximum of ten second order sections and as such the total number of poles and zeros is limited. The number of complex poles was set to 4 and the number of complex zeros set to 3, with a single real pole and two real zeros used in addition to shape the filter. The swarm was looped over all chambers in the X and Y degrees of freedom. While the scripts functioned as expected for Z, the sensor correction signal is fed through HEPI rather than the ISI on the BSCs, complicating the problem, as the response of HEPI needs to be accounted for. The ground conditions used for this optimisation was a ground motion of 400 nm/s, measured by taking the Band Limited Root Mean Square (BLRMS) of the ground motion, measured by the STS-2 in the microseismic peak band of 0.1 to 0.3 Hz. To attempt to get a 'best fit' sensor correction filter, data with a local wind speed of 7 m/s or 16 mph was used, which equates to the 75 th percentile of wind speeds experienced at LHO [109].

To evaluate the performance of the swarmed filter, the current sensor correction filter is evaluated through the same cost weighting and cost function and is plotted alongside the input ground motion as measured by the STS-2 and the best residual possible. This is calculated by taking the coherent subtraction using the method described by Allen [98], previously applied in [95], between the ground STS-2 and the platforms inertial sensor. The residual injected tilt and ground costs are displayed for the current and swarmed filters along with their respective total RMS cost.

Table 6.1: Table showing a typical output of the swarming process. Here the total number of iterations as well as the total number of filters designed are shown. The best cost for that iteration is shown as well as the average cost for that iteration. Stall iterations indicates the number of iterations that the best cost has not improved within an error of the function tolerance, in this case was 10^{-12}. Iterations where the average cost increased significantly are highlighted in bold text.

Iteration	Number of Swarmed filters	Best Cost	Average Cost	Stall Iterations
271	816000	3.318e-07	1.483e+08	0
272	819000	3.318e-07	3.513e+07	1
273	822000	3.318e-07	212	0
274	825000	3.318e-07	**7.541e+07**	1
275	828000	3.318e-07	**2.54e+13**	2
276	831000	3.318e-07	2.586e+07	3
277	834000	3.317e-07	1.48e+07	0
278	837000	3.317e-07	36.61	1
279	840000	3.317e-07	2076	2
280	843000	3.317e-07	**3.847e+08**	3
281	846000	3.317e-07	**1.246e+19**	4
282	849000	3.317e-07	1.906e+07	5

In total, six filters need to be designed, the end stations (ETMX and ETMY) in X and Y, as well as filters for the corner station in the X and Y degrees of freedom. Due to the close proximity of chambers in the corner station, a single sensor correction filter is used on all chambers, to maximize the common tilt motion seen by all chambers. The end stations can be evaluated using solely data from that chamber, with a cost function given by equation 6.4 and are discussed first.

6.3.1 End Stations

The end station sensor correction filters are the simplest filters to design, as the designed filters only have to work on a single chamber in a single degree of freedom. The on axis degrees of freedom, e.g. ETMX-X and ETMY-Y have BRSs to measure the ground tilt and are placed as close to the STS-2 as possible to maximise the coherence between the two sensors. The off axis degrees of freedom, e.g. ETMX-Y and ETMY-X do not have any ground tilt subtraction from the STS-2 and thus the sensor correction filters will look different.

FIG 6.7 shows the swarmed filter SC in equation 6.4 and its complement $(1 - \mathrm{SC})$ plotted against the sensor correction filter that is currently in use at LHO. The swarmed filter is broadly the same shape as the previous filter, it contains a strong roll off at low frequency to minimise tilt injection and is flat at high frequency, while its complement rolls off at high frequency. The differences between the two filters are subtle, the swarmed filter is not as aggressive at low frequency as the current filter, the AC coupling of the STS-2 reduces the response of the seismometer, reducing the tilt injection before sensor correction. The microseismic suppression at 0.1 Hz appears not to be as strong as the current filter, though this results in significantly lower gain peaking between 10 and 100 mHz. This reduction in gain peaking will reduce the amplification of ground motion in the earthquake band, between 30 and 100 mHz [110].

The performance of the ETMX-X filter is shown in FIG 6.8, here the swarmed filter (red) and the current filter (blue) are compared against the coherent subtraction of the STS before sensor correction and the T240 inertial sensor located on stage 1 of the BSC ISI. This coherent subtraction

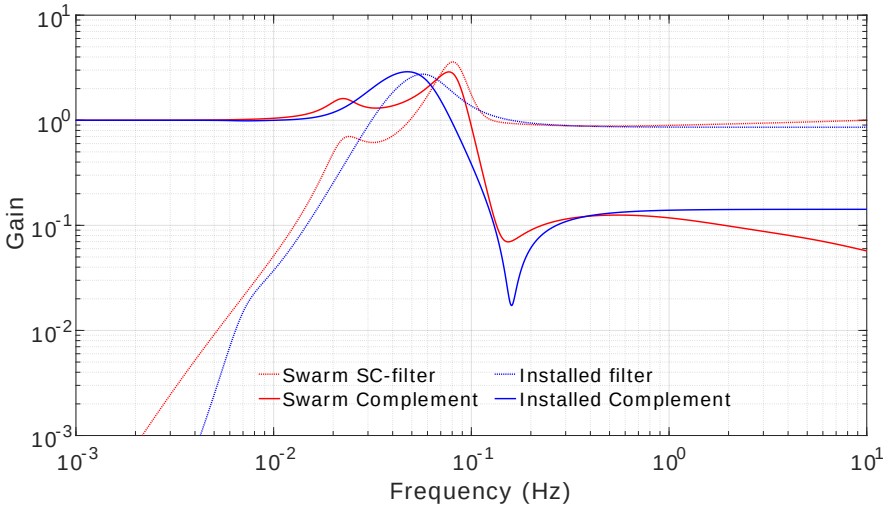

Figure 6.7: Figure showing the response of the swarmed sensor correction filter (dashed red) with its complement (red) compared against the response of the current sensor correction filter (dashed blue) and its complement (blue) for ETMX in the X degree of freedom at LHO

represents the best possible subtraction available by the sensor correction filter. The swarmed filter subtracts the coherent motion from the platform from 4 Hz to 0.1 Hz, gaining its RMS improvement by more aggressive damping of the microseismic peak while keeping the tilt contribution (shown by the dashed curve) at similar RMS to the current filter down to 10 mHz. Below this frequency the swarmed filter injects more tilt than the current filter, though the tilt contribution still rolls off towards low frequency. This is important as a $\frac{1}{f}$ slope in velocity will look flat in displacement, thus to avoid low frequency drift the filter should roll off faster than $\frac{1}{f}$.

FIG 6.9 shows the designed filter in the off axis degree of freedom for the ETMX chamber at LHO. Compared to the current filter (blue), the swarmed filter (red) has a large notch at the microseismic peak, which will suppress motion at the microseismic peak. The swarmed filter causes some small gain peaking below 0.1 Hz, though this is quite small and shouldn't cause problems when in observation time, as other current sensor correction filters, shown by FIG 6.7 and FIG A.4 show more gain peaking over a wider frequency range. To achieve this increased suppression at the microseismic peak, the filter injects more tilt at low frequency, though has the same roll off below 6 mHz as the currently installed filter.

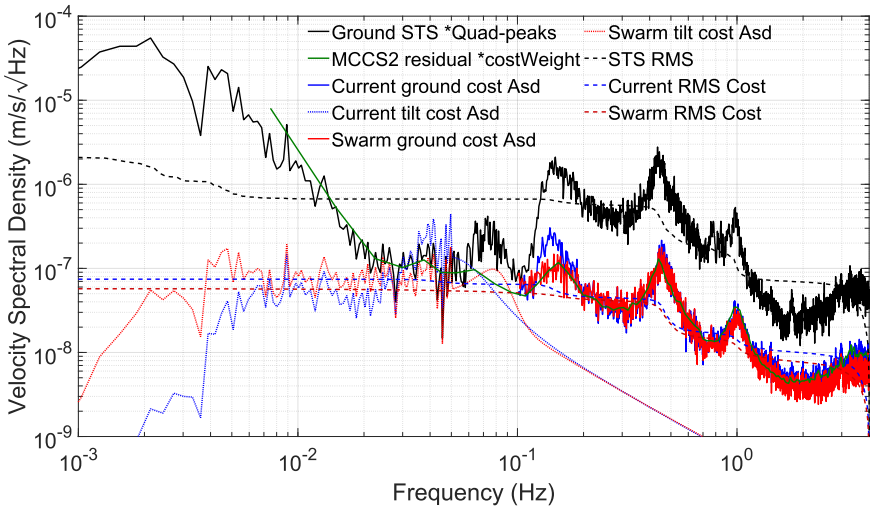

Figure 6.8: Figure comparing the performance of the current sensor correction filter (blue) with the swarmed sensor correction filter (red) on ETMX at LHO in the X degree of freedom. The ground motion (black) and the Multi Channel Coherent Subtraction (MCCS2) residual (green) showing maximum possible subtraction from the CPS.

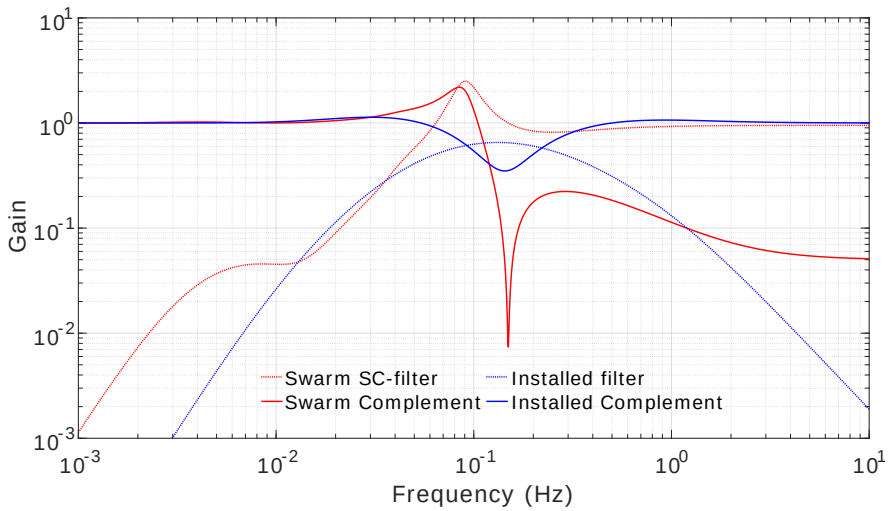

Figure 6.9: Figure showing the response of the swarmed sensor correction filter (dashed red) with its complement (red) compared against the response of the current sensor correction filter (dashed blue) and its complement (blue) for ETMX in the Y degree of freedom at LHO

FIG 6.10 shows the velocity RMS comparing the swarmed filter against the current filter. We find that the swarmed filter has substantially greater primary and secondary microseism suppression

down to the coherence level of the ground seismometer and the platforms inertial sensor. This suppression leads to a factor of 5 reduction in the velocity RMS between 100 and 400 mHz. At 100 mHz, the swarmed filter couples in more tilt to the injected motion, reducing its overall RMS improvement. Below 10 mHz the filter injects around a factor of three more tilt, though this contributes little to the final RMS, like ETMX-X the velocity rolls of like $\frac{1}{f^2}$ below 50 mHz.

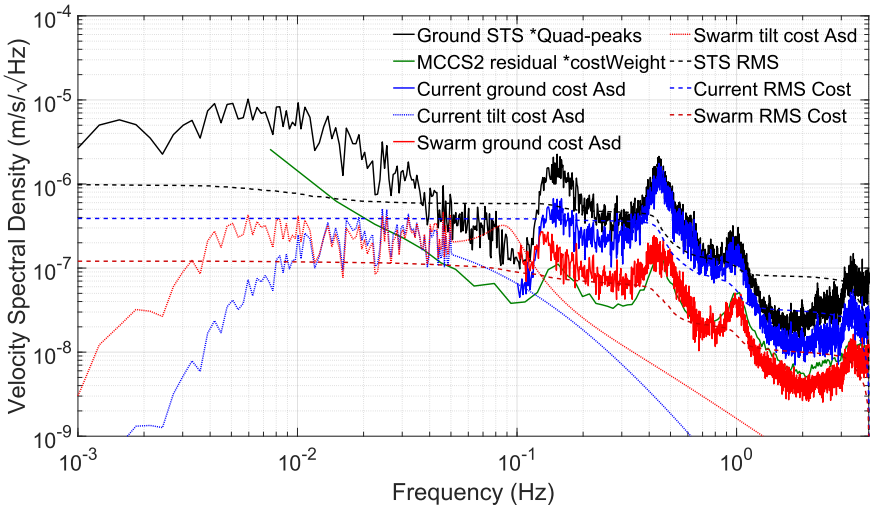

Figure 6.10: Figure comparing the performance of the current sensor correction filter (blue) with the swarmed sensor correction filter (red) on ETMX at LHO in the Y degree of freedom. The ground motion (black) and the MCSS2 residual (green) showing maximum possible subtraction from the CPS.

The filter design process was repeated for the ETMY, the filters and filter performance can be found in Appendix A.4. Any designed sensor correction filter must be able to provide good microseism suppression and minimal tilt injection in a range of different environmental conditions. The filters designed for both ETMs in the X and Y degrees of freedom were tested on a range of different ground motion and wind conditions, these are summarised in Table 6.2. From this we find that the swarmed filters provide modest reductions in ground injection shown in the on axis degrees of freedom, ETMX-X and ETMY-Y. In the off axis degrees of freedom, the performance increase is much more substantial, in ETMX-Y the reduction in ground injection is at least 60% less.

Table 6.2: Table showing the percentage improvement of the velocity RMS of the swarmed filters compared to the currently installed sensor correction filters at LHO. G represents the average BLRMS microseismic motion in the 0.1 to 0.3 Hz band in nm/s, and w represents the average wind speed in m/s over the data stretch as measured at the corner station of LHO. In principle, filters designed for high microseism or high wind could be switched to in real time, though this is not currently planned. Bold indicates that filters were designed on this stretch of data.

Data	ETMX-X	ETMX-Y	ETMY-X	ETMY-Y
G396 w7	16.86	60.18	38.21	20.1
G432 w7	23.09	68.96	37.99	24.12
G538 w1	26.76	70.69	18.89	10.62
G614 w6	17.22	65.58	27.15	0.39
G628 w3	27.37	72.38	27.81	21.87
G714 w5	23.06	55	21.8	12.21
Mean	22.39	65.47	28.64	14.89

6.3.2 Corner Station

In the corner station at LHO and LLO, a single sensor correction filter is used for each degree of freedom to maximise the common tilt motion seen by each chamber. Using FIG 5.4 from chapter 5 as a guide and taking the HAM ISI chambers as an example, tilt couples into the stage 1 motion through two separate paths. First is the tilt coupling through the so called ground plant, the transmission of ground tilt motion through the blade springs of the ISI, as we cannot control the ground tilt and the spectra of ground tilt is likely different depending on the chamber that is being investigated. The second cause is the tilt caused by feedback loops and tilt coupling into sensors on the ISI. Examples of this are the tilt to translation coupling of the horizontal GS13s on the ISI and tilt injection through the sensor correction filter.

FIG 6.11 shows the measured motion of the stage 1 GS13s compared against the common and differential signals between the GS13s. This data was taken from the same data stretch as that used to design sensor correction filters, i.e. with the isolation and feed forward loops turned off. This shows the motion between the GS13s on HAM2 and HAM3 is dominated by common motion between the platforms at 0.1 Hz and above, this is true for the motion between the BSCs and the HAMs, shown by FIG 6.12 which shows the same plot but for HAM3 and the beamsplitter ISI.

Below 0.1 Hz in both FIG 6.11 and 6.12 the motion becomes dominated by differential motion between the chambers. By looking at FIG 5.12 in chapter 5 we find that this motion is likely due to the differential tilt being injected by the stage 1 GS13s on each platform. To minimise the differential tilt that the platforms will measure, a single sensor correction filter should be used on each platform. This will ensure that the tilt injected by the sensor correction filter will remain common to all platforms in the corner station, minimising the differential tilt injection across the corner station.

As such, to optimise this filter, rather than using the inertial, CPS and STS-2 from a single chamber, data from multiple chambers must be combined to design a filter to optimise a wider range of platform translational motion and rotation. To simplify the optimisation process only

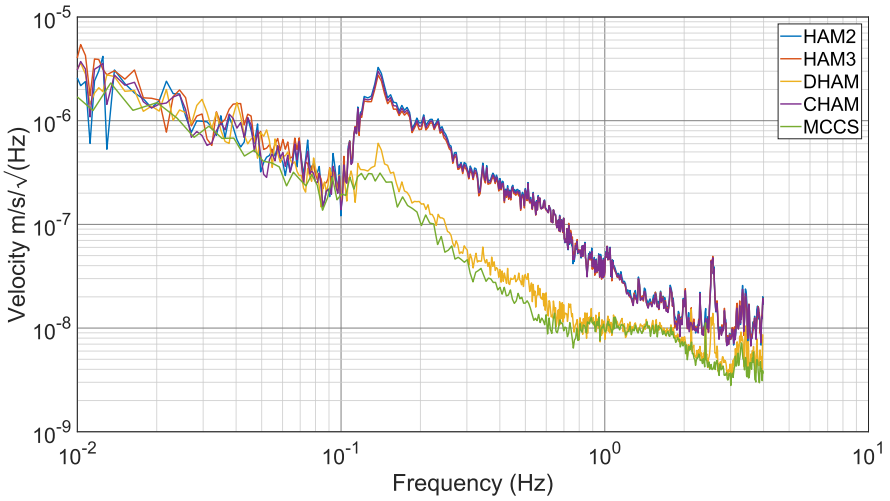

Figure 6.11: Figure comparing the common (CHAM), shown in purple and differential (DHAM) motion, shown in yellow of the HAM ISI chambers at LHO during damped only time. The common motion between HAM2 (blue) and HAM3 (red) is dominated by the common motion between 0.1 and 1 Hz, the coherent subtraction is shown in green for comparison. The differential motion at frequencies below 0.1 Hz becomes dominant due to platform tilt coupling into the measurement.

data from the sensitive degrees of freedom was added together to use as an input. The sensitive degree of freedom is any degree of freedom that is in the beam axis, meaning that for the X degree of freedom: ITMX-X, HAM2-X HAM3-X and the BS-X were used. Similarly for the Y degree of freedom the chambers: ITMY-Y, BS-Y,HAM5-Y and HAM6-Y are used. Once these filters have been swarmed using a single set of input conditions they can be tested over multiple sets of ground conditions to evaluate their RMS improvement over the current filters. Due to their differing sensors, the HAM ISIs will likely measure GS13 noise below 0.3 Hz, while the T240s, due to their superior noise performance will not be subjected to this limitation. This means that the T240s should be able to measure platform tilt to lower frequencies than the GS13s and this should result in better subtraction on the BSC ISIs compared with the HAM ISIs.

The swarming process was set to run using the same cost as outlined previously in FIG 6.4, however this resulted in a mixed sets of filters, that offered only benefits on certain chambers, the summary of this test is shown in Table 6.3.

Table 6.3: Table showing the performance of the swarmed sensor correction filter compared to the currently installed filter in terms of percentage improvement. Bold text indicates sensitive degrees of freedom, red cells indicate that swarmed filter has worse performance than the current filter, green indicates the swarmed filter performs better.

Data	**ITMY**	ITMX	HAM2	HAM3	**HAM4**	**HAM5**	**HAM6**	**BS**	Mean
G396 w7	**-3.84**	-6.89	10.70	6.71	**4.53**	**13.26**	**14.04**	**-3.72**	4.35
G432 w7	**9.63**	11.03	15.36	14.77	**15.03**	**19.44**	**20.53**	**11.99**	14.72
G538 w1	**-7.23**	-7.88	2.33	2.31	**-0.75**	**2.49**	**3.64**	**-6.86**	-1.49
G614 w6	**-5.34**	-6.76	9.92	7.53	**6.87**	**12.16**	**11.17**	**-2.80**	4.09
G628 w3	**2.28**	0.25	13.89	17.55	**15.37**	**22.72**	**23.53**	**3.94**	12.44
G714 w5	**4.69**	4.03	6.77	13.44	**17.10**	**18.79**	**16.41**	**6.77**	11.00
Mean	**0.03**	-1.04	9.83	10.39	**9.69**	**14.81**	**14.89**	**1.55**	7.52
Min	**-7.23**	-7.88	2.33	2.31	**-0.75**	**2.49**	**3.64**	**-6.86**	-1.49
Max	**9.63**	11.03	15.36	17.55	**17.10**	**22.72**	**23.53**	**11.99**	16.11

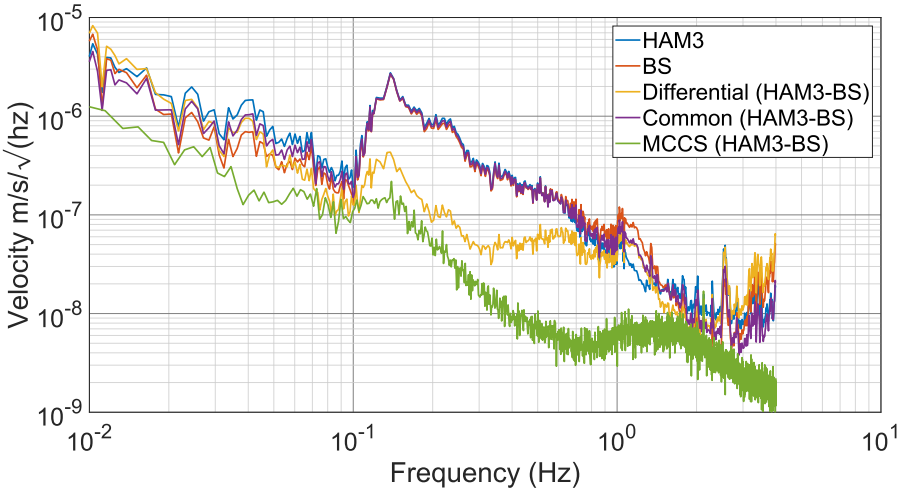

Figure 6.12: Figure comparing the common (purple) and differential (yellow) motion of the HAM ISI and beamsplitter chambers at LHO during damped only time. The common motion between HAM3 (blue) and beamsplitter (red) is dominated by the common motion between 0.1 and 1 Hz, the coherent subtraction is shown in green for comparison. The differential motion at frequencies below 0.1 Hz becomes dominant due to platform tilt coupling into the measurement.

This set of filters were not subtracting enough ground motion around the microseismic peak on the BSC chambers. To encourage the swarm to design a filter to improve the microseismic suppression, the ground cost was modified and is shown in FIG 6.13. Instead of cutting off the ground injection cost at 0.1 Hz and setting the cost below this frequency to zero, the ground injection cost at low frequency is rolled off with a 'f' slope down to 10 mHz. Early filters using this cost resulted in the velocity tilt injection being flat with frequency below 10 mHz, this would result in a sensor correction filter that would slowly drift with time. To counteract this problem the minimum pole frequency was changed from 4 mHz to 6 mHz. As the name would suggest, this prevents the swarm from placing its first pole below 6 mHz, changing the frequency at which the two zeros at DC can be counteracted. This small change in pole frequency likely eliminates a minimum that resulted in the slowly drifting sensor correction filters, in previous corner station tests.

The swarm was set to run across the corner stations in the X degree of freedom using two hours of data taken under 700 nm/s ground motion and 5 m/s wind speed, a comparison between the

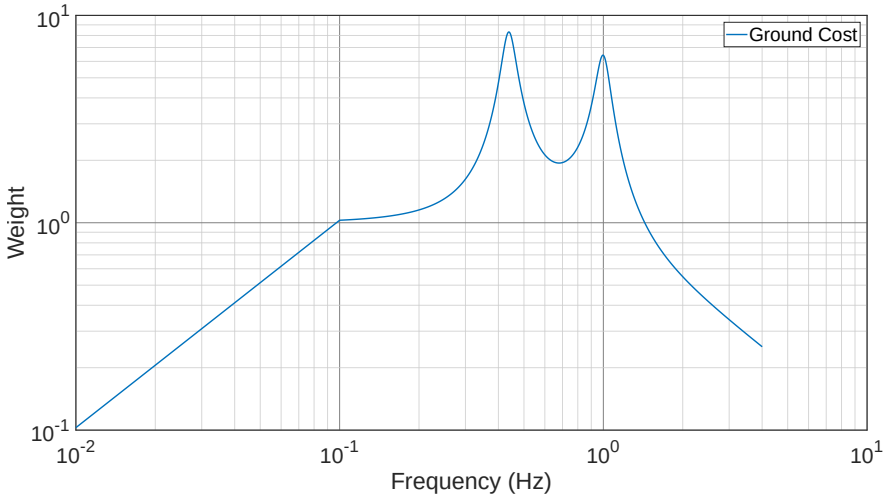

Figure 6.13: Figure showing the modified cost used to swarm multiple chambers in the corner station at LHO. Below 0.1 Hz the cost is decreased like 'f' down to 10 mHz to encourage more suppression of the microseism compared to the previous cost.

designed filter and the swarmed filter is shown in FIG 6.14. The swarmed filter has similar levels of microseismic suppression at a slightly lower frequency than the current filter. The gain peaking, shown between 20 and 100 mHz is a factor of two lower than the current filter at its highest value, and spans a narrower frequency band. The tilt suppression starts at a higher frequency and is a factor of two lower at 40 mHz. The swarmed filter injects more tilt below 20 mHz than the current filter though overall contribution to the RMS should be negligible, the filter has a less severe roll off at very low frequencies, though this should be sufficient to eliminate any drifts caused by the filter.

FIG 6.15 shows the performance of the swarmed filter on the data used to produce the filter with the modified cost function shown in FIG 6.13. The optimised filter has greater ground motion suppression from 4 Hz down to 50 mHz than the current filter, at the microseismic peak the RMS is a factor of 1.5 lower with the new filter. This reduction in ground injection causes more tilt motion to be injected below 40 mHz though this has a negligible effect on the overall velocity RMS of the injected ground motion injection. Below 10 mHz the swarmed filter rolls off with a slope of f^3, due to the three zeros at DC, compared to the current filter's roll-off of f^5, which is

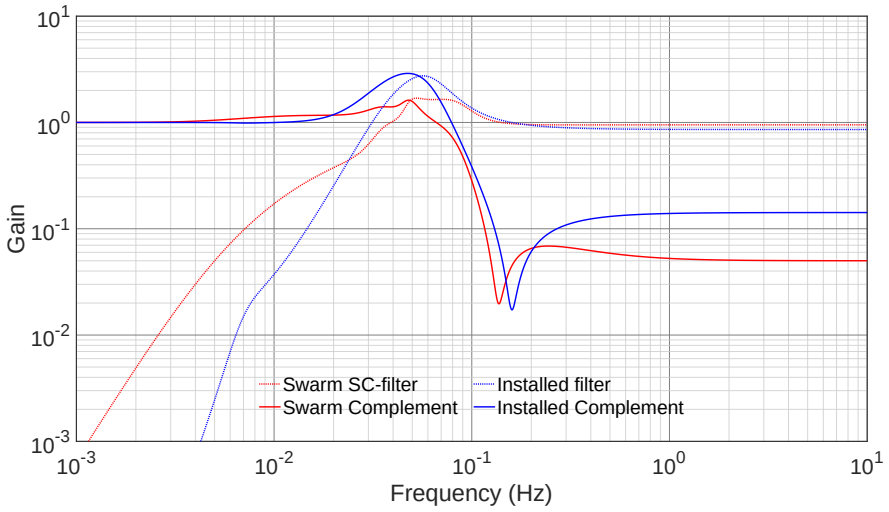

Figure 6.14: Figure showing the response of the swarmed sensor correction filter (dashed red) with its complement (red) compared against the response of the current sensor correction filter (dashed blue) and its complement (blue) for the corner station in the X degree of freedom at LHO

likely excessive. The difference between the coherent subtraction trace and the predicted motion injection is likely due to the different number of averages that each trace underwent when plotting.

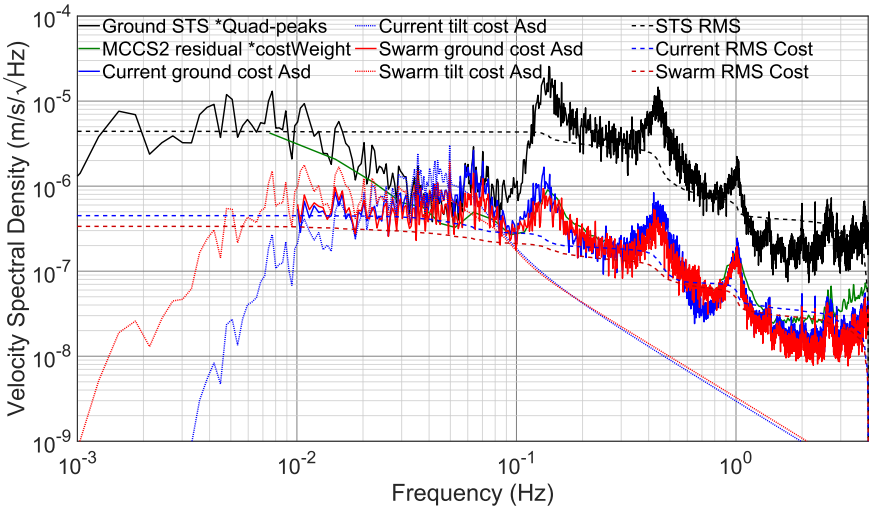

Figure 6.15: Figure showing the performance of the particle swarmed sensor correction filter (red), compared against the current filter (blue), the coherent residual of the ST1 inertial sensor and the ground seismometer (green) and the input ground motion (black)

Table 6.4: Table showing the percentage improvement by using the swarmed filters in X over the current filters in the X degree of freedom at LHO in the corner station. Bold text indicates the on axis, or most sensitive chambers for this degree of freedom. The average, as well as the maximum and minimum improvement expected for the swarmed filter, is shown. This filter is shown to improve the isolation performance of each chamber across a range of different environmental conditions

Data	ITMY	**ITMX**	**HAM2**	**HAM3**	HAM4	HAM5	HAM6	**BS**	Mean
G396 w7	5.45	**12.87**	**15.32**	**19.03**	9.40	7.22	13.75	**9.25**	11.54
G432 w7	7.71	**17.23**	**16.54**	**21.45**	10.79	10.71	14.08	**12.00**	13.81
G538 w1	6.44	**5.70**	**8.88**	**14.28**	16.40	8.62	20.04	**7.69**	11.01
G614 w6	9.49	**17.47**	**15.50**	**21.50**	10.66	10.95	14.54	**12.30**	14.05
G628 w3	11.00	**15.98**	**12.47**	**21.76**	17.41	11.34	22.57	**15.40**	15.99
G714 w5	8.85	**16.35**	**11.44**	**20.32**	12.41	11.88	16.57	**12.72**	13.82
Mean	8.16	**14.27**	**13.36**	**19.72**	12.85	10.12	16.93	**11.56**	13.37
Min	5.45	**5.70**	**8.88**	**14.28**	9.40	7.22	13.75	**7.69**	9.05
Max	11.00	**17.47**	**16.54**	**21.76**	17.41	11.88	22.57	**15.40**	16.75

Table 6.4 shows the relative performance of the optimised sensor correction filter vs the current sensor correction filter in terms of a percentage improvement in the velocity RMS of ground motion injection to the isolated platform through the sensor correction path in the X degree of freedom. The swarmed filter shown in FIG 6.14 reduces the ground injection by an average of 13.37% over a range of environmental conditions across all chambers in the corner station. For the on axis degrees of freedom, we find that this improvement increases to 14.72%. Some degrees of freedom, such as a couple of the HAM ISI chambers experience as much as a 20% improvement in RMS motion using this filter during either high wind conditions (432 nm/s ground motion, 7 m/s wind speed) or high microseismic conditions (714 nm/s ground motion, 5 m/s wind speed).

A similar process was repeated for the Y degree of freedom, however the results were not as successful. The filters, according to the metric of reducing the RMS velocity performed better than the current filter, though failed 'sense checking' when examining the filter. In total six optimisations, with tweaks to the minimum pole frequency and number of poles were conducted, the designed filters either resulted in large tilt injection in the 10-100 mHz band, or had insufficient roll off at low frequency to counteract low frequency drift. As such, the filter designed for the X degree of freedom was evaluated on the corner station for the Y degree of freedom. The performance of the filter, when evaluated on the Y degree of freedom input data, i.e. the summed on axis chambers in the corner station is shown in FIG 6.16.

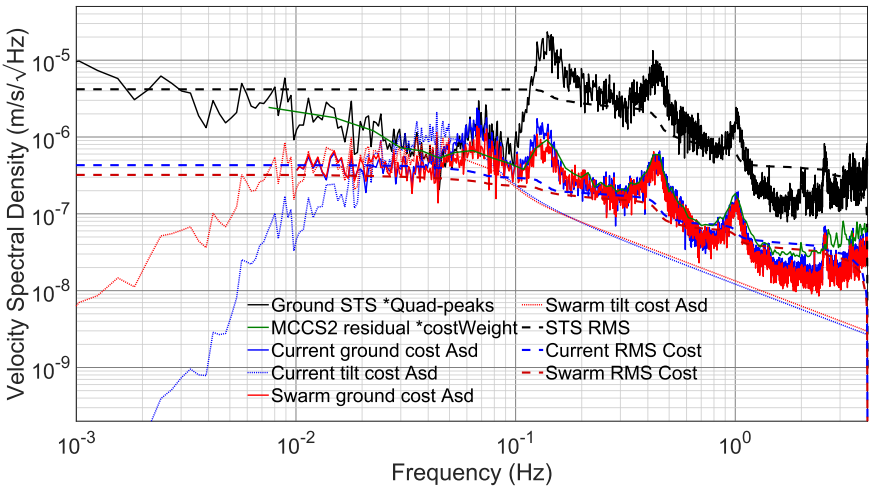

Figure 6.16: Figure showing the performance of the particle swarmed sensor correction filter (red), compared against the current filter (blue), the coherent residual of the ST1 inertial sensor and the ground seismometer (green) and the input ground motion (black) in the Y degree of freedom at LHO. This filter was originally designed for the X degree of freedom at the corner station of LHO, though is evaluated on chambers in the Y degree of freedom in the corner station.

The swarmed sensor correction filter offered very similar performance from 4 Hz down to 200 mHz, as both the current and swarmed filter are sitting on the coherent subtraction trace. From 200 mHz the swarmed filter injects 25% less ground motion at the primary microseismic peak and makes additional gains compared to the current filter in the 50 to 80 mHz region due to the reduction in

Table 6.5: Table showing the percentage improvement by using the swarmed filters over the current filters in the Y degree of freedom at LHO. The average of the combined RMS for the current and swarmed filters are then calculated to determine the best sensor correction filter to run.

Data	**ITMY**	ITMX	HAM2	HAM3	**HAM4**	**HAM5**	**HAM6**	**BS**	Mean
G396 w7	**4.27**	6.19	8.76	10.74	**13.76**	**11.14**	**12.34**	**6.7**	9.24
G432 w7	**12.89**	17.32	11.97	15.35	**19.64**	**14.9**	**15.94**	**16.1**	15.51
G538 w1	**-2.78**	-3.31	1.59	6.4	**4.58**	**3.16**	**3.23**	**-2.47**	1.3
G614 w6	**8.83**	11.91	8.61	9.34	**16.32**	**13.38**	**12.88**	**9.82**	11.39
G628 w3	**10.57**	11.15	7.21	15.41	**16.32**	**14.95**	**20.87**	**12.83**	13.66
G714 w5	**12.82**	16.76	4.49	11.95	**18.26**	**14.2**	**14.36**	**15.17**	13.50
Mean	**7.77**	10	7.105	11.53	**14.81**	**11.96**	**13.27**	**9.69**	10.77
Min	**-2.78**	-3.31	1.59	6.4	**4.58**	**3.16**	**3.23**	**-2.47**	1.3
Max	**12.89**	17.32	11.97	15.41	**19.64**	**14.95**	**20.87**	**16.1**	16.14

gain peaking of the optimised filter. In the 30-50 mHz frequency range the swarmed filter injects a factor of two less tilt into platform below 30 mHz the tilt contributions equalize. Due to the sharp roll-off of the current filter, at 10 mHz the current filter has a factor of 4 times less tilt injection though the contributions to the RMS in this frequency band are minimal.

Table 6.5 shows the percentage improvements to the velocity RMS of the ground motion that is injected into the chambers in the corner stations of LHO, in a variety of different input conditions. The swarmed filter, across all chambers in all degrees of freedom injects 10% less ground motion and tilt into the platform. A few chambers, notably the BSCs (ITMX, ITMY, BS) when wind speed is 1 m/s injects 2% more motion into the platform, in all other conditions there is at least a 4% decrease in the injected motion. Despite these small increases in motion, every chamber under all ground conditions, or a single chamber across each set of ground conditions should experience less ground motion injection due to these filters.

6.4 Summary

Sensor correction is a feed forward technique used to subtract ground motion from the isolated platform by adding the measured ground seismometer signal to CPS sensor. This is complicated as the ground seismometer is limited by ground tilt at frequencies lower than approximately 50 mHz, which should not be injected into the platform motion. This problem is further complicated, as the main contribution to ground translational motion spectrally close to where the ground seismometer is limited by tilt. This results in the design of the sensor correction filter being a balancing act between limiting ground tilt injection and suppressing the microseismic peak. Typically such filters take tens or even hundreds of hours to be designed by on site commissioners and around a single set of environmental conditions.

The design of these filters can be automated and optimised with particle swarm optimisation. This uses a suitably designed and physically motivated cost function to evaluate thousands of test sensor correction filters in each iteration of the optimisation technique. When comparing the CPS injected ground motion, the 'swarmed' sensor correction filters result in the RMS velocity of the end stations being reduced by up to 70%. When designing the filters for the corner station, in order to reduce the differential tilt between each platform, the input data for each chamber in the on-axis degrees of freedom has to be added together. While the reduction in RMS motion is lower for the corner station when compared with the end stations, running the swarmed filters results in the RMS motion of the platforms being reduced by up to 24% when compared to the current filters using the same input ground motion and wind speed. The swarmed filters are being evaluated for use in O3 in the commissioning break between O3a and O3b.

6.4.1 Future Work

Now that a method of swarming control filters has been shown to produce good results with the 'simple' problem of optimising sensor correction filters, optimisation of harder problems can be developed. Notable examples of these would be swarming blending filters for the HAM ISIs and

suspension damping filters, the latter designed to limit the injected sensor noise into the suspension chain, while simultaneously minimising the resonances of the triple and quadruple suspensions on the ISI.

In the case of swarming blending filters for optimising the ISIs, a modified version of the HAM model, that uses FFTs instead of ASDs to estimate the final motion of the chambers can be used as a base for the cost function. This can be weighted by the resonances of the suspensions present on the HAM ISIs, and weighted by the HAM ISI requirements as given by [111]. For example any designed filter that exceeds these limits can be given a large cost weight to avoid designing filters that do not meet the requirement. Particle swarming would be an ideal use case for designing these filters due to the dimensionality of the problem. For instance, the blend filters in X and Y are influenced by the tilt to translation coupling through RY and RX respectively. This cross coupling is hard for humans to design, due to the number of compromises that need to be made in the same degree of freedom or indeed between multiple degrees of freedom. The particle swarm, with a sufficiently well designed cost function and weightings should be able to find a set of filters that balances out the requirements, as has been demonstrated with the design of sensor correction filters.

Chapter 7

Summary and Future Work

Well, that was easy

Tony Stark

Reduction of seismic noise in gravitational wave detectors remains an important area of investigation despite not directly limiting the current sensitvity of the current Advanced LIGO detectors. Technical noises, largely driven by residual isolation platform motion dominate the measured strain sensitivity of current generation detectors [49] and must be reduced to detect gravitational waves between 10 and 30 Hz [47]. Reducing low frequency motion is the central theme of this thesis.

At the beginning of this thesis, I outlined the principles behind the operation of interferometers and phasemeters. I reported on the design and the construction of a high sensitivity and large working-range homodyne phasemeter, HoQI. It has a footprint compact enough for use in LIGO's' suspension systems and as the readout for inertial sensors. I have demonstrated the sensitivity of these devices, as shown in the paper [5], which have a 1000 times finer resolution compared with similar displacement sensors in use at LIGO. With this increase in sensor resolution, it is possible to further damp the resonances of suspensions in use at LIGO, reducing the RMS motion, with no additional noise injection [106]. HoQI is compact enough and has sufficient resolution to increase the sensitivity of geophones, down to the suspension thermal noise of the mechanical

springs, substantially improving the resolution of the sensors currently in use as part of the ISI tables.

Due to the imperfections in optics and potential to measure displacement over a scale of several mm, phasemeters such as HoQI are prone to errors between the measured phase and the real phase, this is caused by non-linearities. In order to quantify the effect these non-linearities will have on future use cases of HoQI, I have built a model to simulate these effects and compared it with real data. This model can be used to determine whether these non-linear effects will limit the resolution of the device and to test improvements in linearity gained by using ellipse fitting based on the technique shown by Rosin [91]. By conducting these tests, I have demonstrated the ability to quantify the scale of non-linear effects and compare them directly with measured data, which show excellent agreement. Moreover, I show that it is crucial to look at the frequency content of these non-linear effects, rather than just using the RMS error. I have calculated the effect that non-linearities in HoQI will have on high Q geophones and the LIGO quad suspension system and have shown that the non-linear effects can be satisfactorily mitigated by ellipse fitting.

The interferometric inertial sensors proposed at the start of this thesis have been constructed and demonstrated, albeit in early prototypes. I have shown that by using HoQI, as the readout mechanism for an inertial sensor, the resolution can be increased by a factor of 60 at 10 mHz. This result shows that self-noise of the interferometric inertial sensor is only factor of 3 above the suspension thermal noise of the L-4C geophone at 10 mHz. The initial prototype, while showing encouraging low frequency performance was beset by mechanical cross couplings and lack of coherence between the devices two readout mechanisms. This lack of coherence between multiple devices prevented a better measurement of the noise floor. Extensive re-design work was undertaken that successfully increased the coherence between multiple devices. I have detailed the increase in isolation performance using a combination of active and passive control that must be achieved in the lab to achieve this goal. Future work, to build more interferometric sensors and the design of an active control loop, to suppress input motion is underway and should allow the noise floor of the inertial sensor to be observed over a wide frequency span.

The second half of my thesis focuses on improving the control systems and seismic isolation performance of the LIGO ISI. To do this I have constructed a predictive, modular model of the HAM ISI to demonstrate the improvements that the sensors described in the first part of my thesis will have if they are used as part of the ISI control loop. As the model uses no in-loop sensors to calculate the expected ISI platform motion, the model can highlight the cause of limitations in the control loop in-situ and be run in any configuration the two LIGO sites are in. The former is useful for noise hunting and diagnosis, as I have shown in a tech-note on reducing SRCL motion [100], where through the use of new filters, the ISI motion can be reduced by a factor of 4 in the region of 0.7 to 10 Hz, while sacrificing only a factor of two increase in motion between 0.3 and 0.7 Hz. By using higher resolution displacement sensors, I have shown that the small increase in noise can also be eliminated. By designing new blending filters, I have used the model to evaluate the isolation performance of a ISI that uses the interferometric sensors described in chapters 2 and 4 in place of current sensors. The isolation performance can be increased by a factor of 70 at 0.1 Hz and a factor of 10 at 2 Hz when looking at the vertical degree of freedom. In reality, the improvement in isolation performance will likely not be as high, due to cross couplings between different degrees of freedom, which the model doesn't account for. In addition, I have shown that the sensor noise injection from sensors used to control the quad suspension is a factor of 60 lower when using HoQI at the UIM at 10 Hz compared with using BOSEMs at the top stage, by using a model written by Shapiro and Bonilla [94].

The final chapter in my thesis focuses on using the particle-swarming optimisation technique to improve the sensor correction filters by reducing the RMS velocity of the isolation performance. I have demonstrated that by using this technique, it is possible to create new sensor correction filters, that when evaluated through a cost function, show a reduction in RMS velocity of the injected ground motion as measured by the CPS. Once a suitable cost function has been designed, these filters can be designed automatically for a range of different environmental conditions, saving future design time. The downside to this technique is that some of the time saved in creating a filter is offset by the creation of the initial cost function. Nevertheless, filters using this technique

show as much as a 70% reduction in the velocity RMS in the end stations and 20% reduction in the corner station, where the motion of multiple chambers must be summed together to minimise the differential motion of the chambers. A natural expansion to this work would be to use the particle swarming technique on other control systems, such optimising filters on the HAM ISI or using it to design improved damping filters for use in the QUAD suspensions.

The work on the interferometric inertial sensors is currently being expanded upon - two more sensors have been constructed as have new electronics boxes allowing for the simultaneous measurement of L-4Cs with coil and optical readouts. Work has begun on calibrating the active isolation platform as has the construction of a box to acoustically and thermally shield the sensors from environmental noise. Once these steps have been completed the sensors can be 'huddle' tested in a shielded environment which should allow for the self noise of the inertial sensor to be measured. Moreover, a natural expansion of this work would be to integrate HoQI with horizontal sensors, such as Watts linkages and perform more 'huddle' tests. Once completed, these sensors could be packaged to be integrated into the control system of a prototype gravitational wave detector, such as the 10 m prototype at the Albert Einstein Institute in Hanover, Germany. This would allow for the expected improvements to seismic isolation systems to be validated in real world conditions.

HoQI can be further developed by packaging the interferometer into a complete unit, full with an enclosure to protect against stray light coupling into the readout standard electronics connector, to allow for HoQI to be used in any experiment. To become immune from angular misalignments HoQI will need use corner cubes instead of mirrors, this will require a slight re-design as a different fibre coupler with a larger collimated beam-spot will be required. This should allow for the working range of the device to be extended, as HoQI will no longer be susceptible to tilt. For use in LIGO HoQI will need to be vacuum qualified and tested accordingly. This testing should show an improvement to the reported sensor resolution at low frequency, where noise associated with air currents should be reduced.

Finally, work on the modeling aspects of this thesis can be expanded in two ways, either by modeling of the BSC ISIs or by modeling global control signals. The former of these will quantify

the performance improvement that interferometric sensors can provide to both types of ISI present in LIGO. The modeling of global controls should allow for the RMS motion reduction of the ISIs to be propagated throughout the whole interferometer control loop, doing so would allow for new filters to be designed with the aim of reducing bandwidths of global controls, this in turn should directly reduce technical noise coupling into the detector.

Appendix A

Appendicies

A.1 Ellipse Fitting

The equation of a simple ellipse, with semi-major and minor axes, a and b is given by,

$$\frac{x'^2}{a^2} + \frac{y'^2}{b^2} = 1. \tag{A.1}$$

In principle though this ellipse may be rotated, by an arbitrary angle α, so we must apply a rotation matrix on the ellipse.

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{pmatrix} x' - x_0 \\ y' - y_0 \end{pmatrix}, \tag{A.2}$$

where x' and y' are the un-rotated co-ordinates, x_0 and y_0 are the offsets. x' and y' can be represented by the equations,

$$x' = a\cos(\phi), \tag{A.3}$$

$$y' = b\sin(\phi), \tag{A.4}$$

where ϕ is the optical phase. We then see that the general form for our output ellipse is,

$$x = (a\cos(\phi) - x_0)\cos(\alpha) - (b\sin(\phi) - y_0)\sin(\alpha), \tag{A.5}$$

$$y = (a\cos(\phi) - x_0)\sin(\alpha) + (b\sin(\phi) - y_0)\cos(\alpha). \tag{A.6}$$

We can rearrange these values for x and y and substitute them into the general form of an ellipse and expand all the terms out. The ellipse fitting routine uses the standard quadratic form of a rotated ellipse, given by the equation,

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0. \tag{A.7}$$

The cross terms arise from the fact that our ellipse is rotated thus each new coordinate on the rotated axes is comprised of a combination of each of the old coordinates. To put the outputs the form of the rotated ellipse, we must substitute in equations A.5 and A.6 into equation A.1. This gives,

$$\begin{aligned}
1 &= \left[\frac{((x'-x_0)\cos\alpha + (y'-y_0)\sin(\alpha))^2}{a^2},\right]\left[\frac{(-(x'-x_0)\sin(\alpha) + (y'-y_0)\cos(\alpha))^2}{b^2}\right], \\
&= \frac{(x'-x_0)^2\cos^2(\alpha) + (y'-y_0)^2\sin^2(\alpha) + 2(x'-x_0)(y'-y_0)\cos(\alpha)\sin(\alpha)}{a^2}, \\
&+ \frac{(x'-x_0)^2\sin^2(\alpha) + (y'-y_0)^2\cos^2(\alpha) - 2(x'-x_0)(y'-y_0)\sin(\alpha)\cos(\alpha)}{b^2}, \\
&= \frac{(x'^2 + x_0^2 - 2x'x_0)\cos^2(\alpha) + (y'^2 + y_0^2 - 2y'y_0)\sin^2(\alpha)}{a^2}, \\
&+ \frac{(y'^2 + y_0^2 - 2y'y_0)\cos^2(\alpha) + (x'^2 + x_0^2 - 2x'x_0)\sin^2(\alpha)}{b^2}, \\
&+ \frac{2(x'-x_0)(y'-y_0)\cos(\alpha)\sin(\alpha)}{a^2} - \frac{2(x'-x_0)(y'-y_0)\cos(\alpha)\sin(\alpha)}{b^2}.
\end{aligned} \tag{A.8}$$

Multiplying by a^2b^2,

$$\begin{aligned}
a^2b^2 &= 2(x-x_0)(y-y_0)\cos(\alpha)\sin(\alpha)(b^2-a^2), \\
&+ b^2\left[(x^2 + x_0^2 - 2xx_0)\cos^2(\alpha) + (y^2 + y_0^2 - 2yy_0)\sin^2(\alpha)\right], \\
&+ a^2\left[(x^2 + x_0^2 - 2xx_0)\sin^2(\alpha) + (y^2 + y_0^2 - 2yy_0)\cos^2(\alpha)\right].
\end{aligned} \tag{A.9}$$

Subtracting a^2b^2,

$$\begin{aligned} 0 &= x'^2\left[b^2\cos^2(\alpha)+a^2\sin^2(\alpha)\right]+x'y'\left[2\cos(\alpha)\sin(\alpha)(b^2-a^2)\right], \\ &+ y'^2\left[a^2\cos(\alpha)+b^2\sin(\alpha)\right], \\ &+ x'\left[-2b^2x_0\cos^2(\alpha)-2a^2x_0sin^2(\alpha)-2y_0(b^2-a^2)\sin(\alpha)\cos(\alpha)\right], \\ &+ y'\left[-2b^2y_0\sin^2(\alpha)-2a^2y_0\cos^2(\alpha)-2x_0(b^2-a^2)\cos(\alpha)\sin(\alpha)\right], \\ &+ x_0^2(a^2\sin^2(\alpha)+b^2\cos^2(\alpha))+x_0y_0\sin(2\alpha)(b^2-a^2), \\ &+ y_0^2(a^2\cos^2(\alpha)+b^2\sin^2(\alpha))-(ab)^2. \end{aligned} \tag{A.10}$$

We then collect all terms that have x'^2, $x'y'$, y'^2, x', y' and finally the constant terms, and assign their coefficients labels A through F. We now need to calculate the ellipse parameters based on these coefficients.

These coefficients are,

$$A = b^2\cos^2(\alpha)+a^2\sin^2(\alpha), \tag{A.11}$$

$$B = \sin(2\alpha)(b^2-a^2), \tag{A.12}$$

$$C = a^2\cos^2(\alpha)+b^2\sin^2(\alpha), \tag{A.13}$$

$$D = -2b^2x_0\cos^2(\alpha)-2a^2x_0sin^2(\alpha)-y_0(b^2-a^2)\sin(2\alpha), \tag{A.14}$$

$$E = -2b^2y_0\sin^2(\alpha)-2a^2y_0\cos^2(\alpha)-x_0(b^2-a^2)\sin(2\alpha), \tag{A.15}$$

$$\begin{aligned} F &= x_0^2(a^2\sin^2(\alpha)+b^2\cos^2(\alpha))+x_0y_0\sin(2\alpha)(b^2-a^2) \\ &+ y_0^2(a^2\cos^2(\alpha)+b^2\sin^2(\alpha))-(ab)^2. \end{aligned} \tag{A.16}$$

These coefficients can be then written in terms of each other,

$$
\begin{aligned}
A &= a^2 \sin^2(\alpha) + b^2 \cos^2(\alpha), && \text{(A.17)}\\
B &= (b^2 - a^2) \sin(2\alpha), && \text{(A.18)}\\
C &= a^2 \cos^2(\alpha) + b^2 \sin^2(\alpha), && \text{(A.19)}\\
D &= -2x_0 A - y_0 B, && \text{(A.20)}\\
E &= -2y_0 C - x_0 B, && \text{(A.21)}\\
F &= x_0 y_0 B + x_0^2 A + y_0^2 C - (ab)^2. && \text{(A.22)}
\end{aligned}
$$

A.1.1 Calculating the Ellipse Parameters

As we have these coefficients, the ellipse parameters themselves, a, b, x_0, y_0, α need to be calculated in terms of these coefficients. Firstly to find α, noting that,

$$
\begin{aligned}
(A - C) &= a^2 \sin^2(\alpha) + b^2 \cos^2(\alpha) - a^2 \cos^2(\alpha) - b^2 \sin^2(\alpha),\\
(A - C) &= (b^2 - a^2)(\cos^2(\alpha) - \sin^2(\alpha)),\\
(A - C) &= (b^2 - a^2) \cos(2\alpha),\\
\frac{B}{A - C} &= \frac{(b^2 - a^2) \sin(2\alpha)}{(b^2 - a^2) \cos(2\alpha)}. && \text{(A.23)}
\end{aligned}
$$

and thus,

$$
\begin{aligned}
\tan(2\alpha) &= \frac{B}{A - C},\\
\alpha &= \frac{\arctan(\frac{B}{A-C})}{2}. && \text{(A.24)}
\end{aligned}
$$

This result can be used to write $\tan^2(2\alpha)$ as,

$$\begin{aligned}
\frac{\sin^2(2\alpha)}{1-\sin^2(2\alpha))} &= \frac{B^2}{(A-C)^2}, \\
(A-C)^2\sin^2(2\alpha) &= B^2(1-\sin^2(2\alpha)), \\
\sin^2(2\alpha) &= \frac{B^2}{(A-C)^2+B^2},
\end{aligned} \tag{A.25}$$

Using the definition of $\sin^2(2\alpha)$ from equation A.12, we substitute for $\sin^2(2\alpha)$ and get,

$$\begin{aligned}
\frac{B}{b^2-a^2} &= \sin(2\alpha), \\
\frac{B^2}{(b^2-a^2)^2} &= \frac{B^2}{(A-C)^2+B^2}, \\
(A-C)^2+B^2 &= (b^2-a^2)^2, \\
b^2-a^2 &= \sqrt{(A-C)^2+B^2}, \\
b^2 &= a^2+\sqrt{(A-C)^2+B^2}, \\
a^2 &= b^2-\sqrt{(A-C)^2+B^2}.
\end{aligned} \tag{A.26}$$

Noticing that $(A+C) = a^2+b^2$, we can use this relation, and the expressions for both a^2 and b^2 to recover the semi major and minor axes in terms of A, B and C.

$$a = \sqrt{\frac{(A+C)-\sqrt{(A-C)^2+B^2}}{2}} \tag{A.27}$$

$$b = \sqrt{\frac{(A+C)+\sqrt{(A-C)^2+B^2}}{2}} \tag{A.28}$$

Now we need to find expressions for the offsets for the ellipse, to do this we shall rearrange expressions for D, E, and then solve for the offsets.

$$D = -2x_0A + y_0B \tag{A.29}$$

$$E = -2y_0C - x_0B \tag{A.30}$$

Solving for x_0,

$$\begin{aligned} E &= -2\left(\frac{-2x_0A - D}{B}C - Bx_0\right) \\ EB &= 4x_0AC + 2DC - B^2x_0 \\ x_0 &= \frac{2DC - EB}{B^2 - 4AC} \end{aligned} \tag{A.31}$$

Similarly for y_0,

$$y_0 = \frac{2AE - BD}{B^2 - 4AC} \tag{A.32}$$

A.1.2 Correcting for the Ellipse Parameters and Rescaling

Now that we know the ellipse parameters, α, a, b, x_0, y_0 we can correct for the rotation and calibrate the x and y axes accordingly, transforming the ellipse into a circle - reducing the amount of up-conversion in the output data.

First, we subtract off the offsets, x_0 and y_0 from x and y and then apply a rotation matrix, using the angle $-\alpha$ to correct for the rotation of the ellipse.

First we apply a rotation matrix with an angle of $-\alpha$,

$$\begin{pmatrix} x' - x_0 \\ y' - y_0 \end{pmatrix} = \begin{bmatrix} \cos(\alpha) & \sin(\alpha) \\ -\sin(\alpha) & \cos(\alpha) \end{bmatrix} \begin{pmatrix} (x' - x_0)\cos(\alpha) - (y' - y_0)\sin(\alpha) \\ (x' - x_0)\sin(\alpha) + (y' - y_0)\cos(\alpha) \end{pmatrix} \tag{A.33}$$

This gives the coordinates of the un-rotated ellipse according in terms of our original parameters, $x' - x_0$ and $y' - y_0$.

The offsets can now be removed by adding x_0 and y_0 to equation A.33 The final step is to scale the ellipse, we use the values for the parameter F. Remembering that, for this ellipse fitting method to work, we require that F, our constant term to equal 1 and that, F, for our fitted ellipse is equal to,

$$F = A'x_0^2 + B'x_0y_0 + C'y_0^2 - (a'b')^2. \tag{A.34}$$

We can define our scaling factor, K, as $a' = \frac{a}{K}$ and $b' = \frac{b}{K}$, where a, b are the returned parameters from the fitting routine, therefore,

$$K^4 = \frac{A'x_0^2 + B'x_0y_0 + C'y_0^2 - 1}{(a'b')^2}. \tag{A.35}$$

With this we can scale the ellipse, and recover the phase by taking the arctangent,

$$\phi = \arctan\left(\frac{a'y'}{b'x'}\right). \tag{A.36}$$

A.2 Transfer Function Derivations

To derive the transfer functions of inertial sensors, the damping and spring coefficients can be represented in terms of the resonant frequency, ω_0 and the quality factor, Q, by the equations,

$$\omega_0^2 = \frac{k}{m} \tag{A.37}$$

$$\zeta = \frac{1}{2Q} = \frac{b}{2m\omega_0} \tag{A.38}$$

$$\frac{b}{m} = \frac{\omega_0}{Q} \tag{A.39}$$

Where k is the spring constant, m is the mass, ζ is the damping ratio and b is the damping in the system.

A.2.1 Ground to Platform Motion

Referring to diagram a) in FIG 4.1 in chapter 4 the equation of motion for this system is,

$$\begin{aligned} m\ddot{x_p} + b\dot{x_p} + kx_p &= b\dot{x_g} + kx_g, \\ \ddot{x_p} + \frac{\omega_0}{Q}\dot{x_p} + \omega_0^2 x_p &= \frac{\omega_0}{Q}\dot{x_g} + \omega_0 x_g. \end{aligned} \tag{A.40}$$

We now take the fourier transform of this and pick up a factor of $i\omega$ for each differentiation, this gives,

$$-\omega^2 x_p + \frac{i\omega\omega_0 x_p}{Q} + \omega_0^2 x_p = \frac{i\omega_0\omega x_g}{Q} + \omega_0^2 x_g. \tag{A.41}$$

We can rearange this to get the transfer function describing how to map ground to platform motion, this is given by,

$$\begin{aligned} x_p\left(-\omega^2+\frac{i\omega\omega_0}{Q}+\omega_0^2\right) &= x_g\left(\frac{i\omega_0\omega}{Q}+\omega_0^2\right), \\ \frac{x_p}{x_g} &= \frac{\frac{i\omega\omega_0}{Q}+\omega_0^2}{-\omega^2+\frac{i\omega_0\omega}{Q}+\omega_0^2}. \end{aligned} \tag{A.42}$$

A.2.2 External Force to Platform Motion

Referring to diagram b) in FIG 4.1 in chapter 4 the equation of motion for this system is,

$$F_{ext} = m\ddot{x_p}+b\dot{x_p}+kx_p \tag{A.43}$$

$$\begin{aligned} \frac{F_{ext}}{m} &= \ddot{x_p}+\frac{\omega_0}{Q}\dot{x_p}+\omega_0^2x_p, \\ \frac{F_{ext}}{m} &= \ddot{x_p}+\frac{\omega_0}{Q}\dot{x_p}+\omega_0^2x_p. \end{aligned} \tag{A.44}$$

we now take the fourier transform of this,

$$\begin{aligned} \frac{F(\omega)}{m} &= \left(-\omega^2x_p+\frac{i\omega_0\omega x_p}{Q}+\omega_0^2x_p\right), \\ \frac{x_p}{F_{ext}} &= \frac{1}{m(-\omega^2+\frac{i\omega_0\omega}{Q}+\omega_0^2)} \end{aligned} \tag{A.45}$$

A.2.3 Delta X to Ground Motion

Referring to diagram c) in figure 4.1 the equation of motion for this system is,

$$\begin{aligned} m\ddot{x_p}+b\dot{x_p}+kx_p &= kx_g+b\dot{x_g}, \\ \ddot{x_p}+\frac{\omega_0}{Q}\dot{x_p}+\omega_0^2x_p &= \omega_0^2x_g+\frac{\omega_0}{Q}\dot{x_g}. \end{aligned} \tag{A.46}$$

Moving terms onto one side of the equation and adding $-\ddot{x_g}$ to both sides, we can write everything in terms of $\ddot{x_p}-\ddot{x_g}$),

$$(\ddot{x_p} - \ddot{x_g}) + \frac{\omega_0}{Q}(\dot{x_p} - \dot{x_g}) + \omega_0^2(x_p - x_g) \quad = \quad -\dot{x_g}. \tag{A.47}$$

We now say that $\Delta X = x_p - x_g$, this gives,

$$\ddot{\Delta x} + \frac{\omega_0}{Q}\dot{\Delta x} + \omega_0^2 \Delta x = \quad -\ddot{x_g}. \tag{A.48}$$

We now take the fourier transform of this and rearrange to get the transfer function.

$$\begin{aligned} -\omega^2 \Delta x + \frac{i\omega\omega_0}{Q}\Delta x + \omega_0^2 \Delta x &= \omega^2 x_g, \\ \Delta x \left(-\omega^2 + \frac{i\omega_0\omega}{Q} + \omega_0^2 \right) &= \omega^2 x_g. \end{aligned} \tag{A.49}$$

$$\frac{x_g}{\Delta x} \quad = \quad \frac{-\omega^2 + \frac{i\omega\omega_0}{Q} + \omega_0^2}{\omega^2} \tag{A.50}$$

A.3 Derivation of HAM-ISI platform motion

Using the X degree of freedom as an example, as it is the most complicated degree of freedom, we can calculate the expected stage 1 motion in X after it has passed through the control loop, as shown by Kissel in [11].

Working counter clockwise around the loop, shown by FIG A.1, the stage one motion is represented as,

$$\begin{aligned} x_{\mathrm{ST1}} &= P_x^{0-1}(x_{\mathrm{ST0}}) \\ &+ P_x^{1-1}\left(F_{\mathrm{x}}^{FF}[n_{\mathrm{L4C}} + x_{\mathrm{ST0}}]\right) \\ &+ K_x^D[x_{\mathrm{ST1}} + n_{\mathrm{GS13}} - \frac{g}{\omega^2}RY_{\mathrm{ST1}}] \\ &+ K_x^I\left[F_x^{\mathrm{HP}}(-\frac{g}{\omega^2}RY_{\mathrm{ST1}} + n_{\mathrm{GS13}} + x_{\mathrm{ST1}})\right) \\ &+ F_x^{\mathrm{LP}}(F_{\mathrm{GND}}^{\mathrm{SC}}[n_{\mathrm{STS}} + x_{\mathrm{GND}}] + [n_{\mathrm{CPS}} + x_{\mathrm{ST1}} - x_{\mathrm{ST0}}])\big]\big)\,. \end{aligned} \tag{A.51}$$

We re-arrange the loop in terms of the stage 1 motion x_{rmST1} and pull out a factor of $P_x^{1-1}K_x^D$. Using the convention presented by Kissel, we set the damping loop gain such that,

$$G_x^D = P_x^{1-1}K_x^D. \tag{A.52}$$

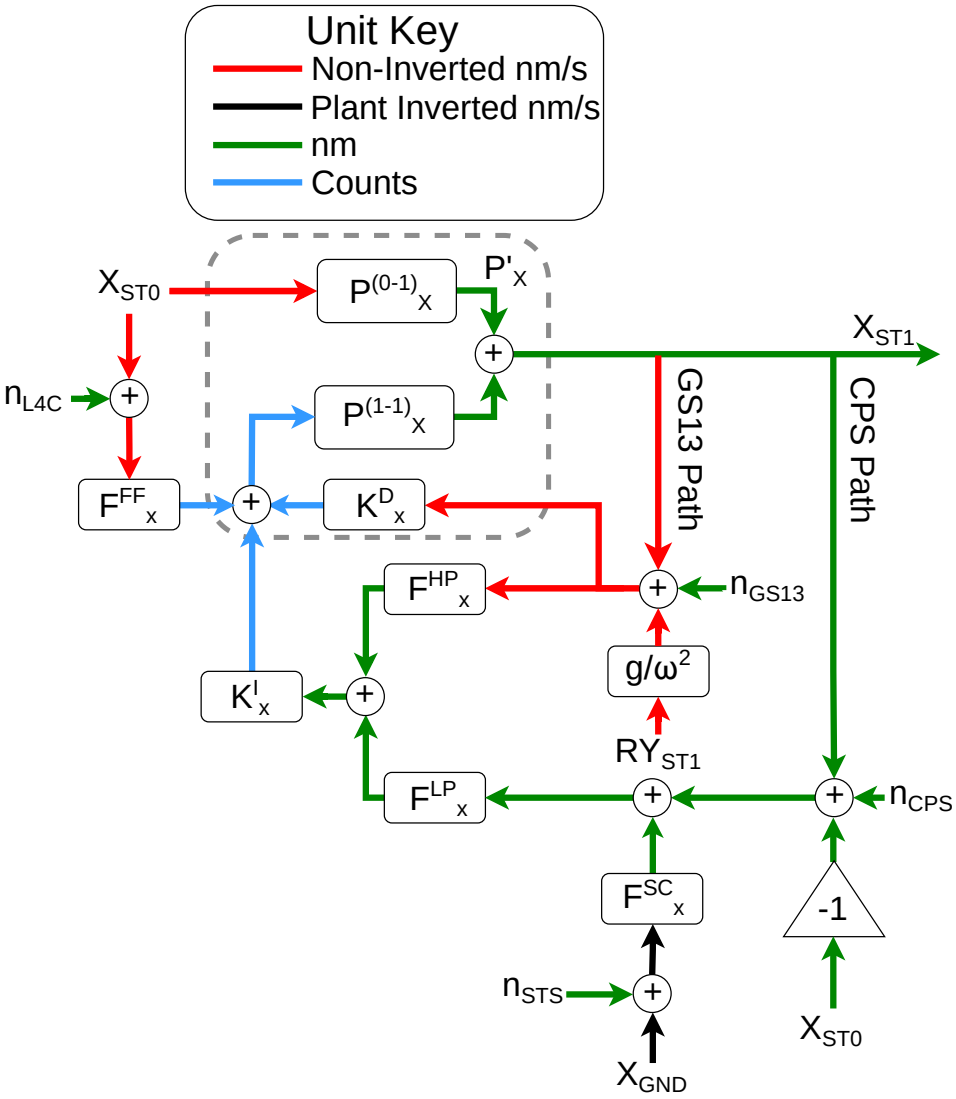

Figure A.1: A control loop diagram of a HAM ISI, showing filter modules e.g. the high pass filter $F_{\mathrm{x}}^{\mathrm{HP}}$, plants e.g. $P_{\mathrm{x}}^{(0-1)}$ describing the ground to platform transfer function, $P_{\mathrm{x}}^{(1-1)}$ describing the actuator to platform transfer function, and controllers $K_{\mathrm{x}}^{\mathrm{I}}$ e.g the isolation filter. Sensor noises are described with the notation n_{SN} while ground motion inputs are denoted as x_{ST0}, adapted from [11].

The stage 1 motion now becomes,

$$
\begin{aligned}
x_{\mathrm{ST1}}(1-G_x^D) &= P_x^{1-1}K_x^I(F^{\mathrm{HP}}+F^{\mathrm{LP}})x_{\mathrm{ST1}} \\
&+ (P_x^{0-1}+P_x^{1-1}F_x^{FF}+P_x^{1-1}K_x^IF_x^{\mathrm{LP}})x_{\mathrm{ST0}} \\
&- P_x^{1-1}(K_x^D\frac{g}{\omega^2}+K_x^IF^{\mathrm{HP}}\frac{g}{\omega^2})RY_{\mathrm{ST1}} \\
&+ P_x^{1-1}(K_x^IF^{\mathrm{LP}}F_{\mathrm{GND}}^{\mathrm{SC}})x_{\mathrm{GND}} \\
&+ P_x^{1-1}(K_x^D+K_x^IF^{\mathrm{HP}})n_{\mathrm{GS13}} \\
&+ P_x^{1-1}(F_x^{FF})n_{\mathrm{L4C}} \\
&+ P_x^{1-1}(K_x^IF^{\mathrm{LP}}F_{\mathrm{GND}}^{\mathrm{SC}})n_{\mathrm{STS}} \\
&+ P_x^{1-1}(K_x^IF^{\mathrm{LP}})n_{\mathrm{CPS}}. \qquad \text{(A.53)}
\end{aligned}
$$

From this, we can define the damped ground and actuation plants as

$$
\begin{aligned}
P_x^{'0-1} &= \frac{P_x^{0-1}}{1-G_x^D} \qquad \text{(A.54)} \\
P_x^{'1-1} &= \frac{P_x^{1-1}}{1-G_x^D}. \qquad \text{(A.55)}
\end{aligned}
$$

The measured stage one motion now becomes,

$$
\begin{aligned}
x_{\mathrm{ST1}} &= P_x^{'(1-1)}K_x^I(F^{\mathrm{HP}}+F^{\mathrm{LP}})x_{\mathrm{ST1}} \\
&+ (P_x^{'(0-1)}+P_x^{'(1-1)}F_x^{FF}+P_x^{'(1-1)}K_x^IF_x^{\mathrm{LP}})x_{\mathrm{ST0}} \\
&- P_x^{'(1-1)}(K_x^D\frac{g}{\omega^2}+K_x^IF^{\mathrm{HP}}\frac{g}{\omega^2})RY_{\mathrm{ST1}} \\
&+ P_x^{'(1-1)}(K_x^IF^{\mathrm{LP}}F_{\mathrm{GND}}^{\mathrm{SC}})x_{\mathrm{GND}} \\
&+ P_x^{'(1-1)}(K_x^D+K_x^IF^{\mathrm{HP}})n_{\mathrm{GS13}} \\
&+ P_x^{'(1-1)}(F_x^{FF})n_{\mathrm{L4C}} \\
&+ P_x^{'(1-1)}(K_x^IF^{\mathrm{LP}}F_{\mathrm{GND}}^{\mathrm{SC}})n_{\mathrm{STS}} \\
&+ P_x^{'(1-1)}(K_x^IF^{\mathrm{LP}})n_{\mathrm{CPS}}. \qquad \text{(A.56)}
\end{aligned}
$$

In the ideal case, the designed blending filters will be complementary, such that $F^{\mathrm{HP}} + F^{\mathrm{LP}} = 1$, however, this is not always the case, so we define this sum as $\epsilon = F^{\mathrm{HP}} + F^{\mathrm{LP}}$. From this, we can define the damped isolation gain as,

$$P_x^{'(1-1)} K_x^I = G_x^{'I}. \tag{A.57}$$

Re-writing this, the platform motion becomes,

$$\begin{aligned}
x_{\mathrm{ST1}}(1-\epsilon G_x^{'I}) &= (P_x^{'(0-1)} + P_x^{'(1-1)} F_x^{FF} + G_x^{'I} F_x^{\mathrm{LP}}) x_{\mathrm{ST0}} \\
&- (P_x^{'(1-1)} K_x^D \frac{g}{\omega^2} + G_x^{'I} F^{\mathrm{HP}} \frac{g}{\omega^2}) RY_{\mathrm{ST1}} \\
&+ G_x^{'I} F^{\mathrm{LP}} F_{\mathrm{GND}}^{\mathrm{SC}} x_{\mathrm{GND}} \\
&+ (P_x^{'(1-1)} K_x^D + G_x^{'I} F^{\mathrm{HP}}) n_{\mathrm{GS13}} \\
&+ P_x^{'(1-1)} (F_x^{FF}) n_{\mathrm{L4C}} \\
&+ G_x^{'I} F^{\mathrm{LP}} F_{\mathrm{GND}}^{\mathrm{SC}} n_{\mathrm{STS}} \\
&+ G_x^{'I} F^{\mathrm{LP}} n_{\mathrm{CPS}}.
\end{aligned} \tag{A.58}$$

Dividing through by $(1-\epsilon G_x^{'I})$, the final platform motion can be calculated, this is given by,

$$\begin{aligned}
x_{\mathrm{ST1}} &= \underbrace{\frac{P_x^{'(0-1)}}{1-\epsilon G_x^{'I}} x_{ST0}}_{Residual\ ground} + \underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} F_x^{FF} x_{ST0}}_{Feed\ forward} \\
&- \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} x_{ST0}}_{CPS\ ground\ injection} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} F_{gnd}^{SC} x_{GND}}_{Sensor\ correction} \\
&- \left(\underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} K_x^D}_{Damping\ tilt\ coupling} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F^{HP}}_{Isolation\ tilt\ coupling} \right) \frac{g}{\omega^2} RY_{ST1} \\
&+ \left(\underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} K_x^D}_{Damping\ GS13\ noise} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{HP}}_{Isolation\ GS13\ noise} \right) n_{GS13} \\
&+ \underbrace{\frac{P_x^{'(1-1)}}{1-\epsilon G_x^{'I}} F_{ST0}^{FF} n_{L4C}}_{L4C\ noise\ injection} + \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} n_{CPS}}_{CPS\ noise\ injection} \\
&+ \underbrace{\frac{G_x^{'I}}{1-\epsilon G_x^{'I}} F_x^{LP} F_{gnd}^{SC} n_{STS}}_{STS\ noise\ injection}
\end{aligned} \tag{A.59}$$

A.4 Particle Swarming Results

FIG A.2 shows a comparison between the swarmed filter and the current sensor correction filter for ETMY in the X degree of freedom. Unlike previous filters, the current and swarmed filters don't feature large notches at the microseism, instead the sensor correction filter rolls off at high frequency to suppress motion at frequencies higher than 0.1 Hz. Both filters have very little gain peaking, while the swarmed filter suppresses tilt between 0.1 and 0.01 Hz more aggressively than the current filter. The swarmed filter has a small notch at 0.5 Hz to suppress the first of the quad resonances.

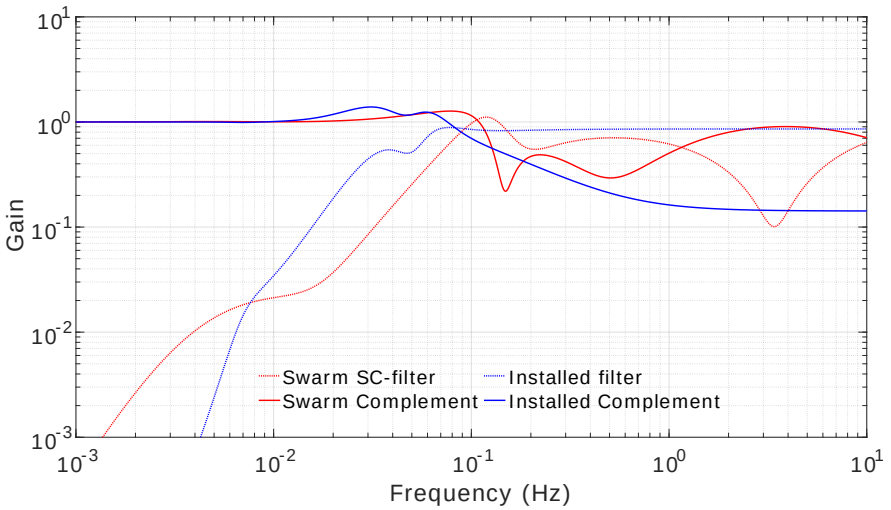

Figure A.2: Figure showing the response of the swarmed sensor correction filter (dashed red) with its complement (red) compared against the response of the current sensor correction filter (dashed blue) and its complement (blue) for ETMY in the X degree of freedom at LHO

FIG A.3 shows the velocity RMS comparing the swarmed filter(red) against the current filter (blue). Compared to the current filter, the swarmed filter performs slightly worse above 1 Hz than the current filter and is a factor of two higher than the best possible subtraction (green) in this frequency region. This shouldn't be much of an issue as the low pass blend attenuates the effect of the sensor correction substantially above the blending frequency; typically a few 100 mHz depending on the configuration of the ISI. The swarmed filter achieves greater suppression

of the microseismic peak and substantially less tilt injection between 10 mHz and 100 mHz where the majority of the RMS improvement is gained. Like previously designed filters the tilt injection term drops off below 10 mHz ensuring that the sensor correction filter will not cause any long term drifts.

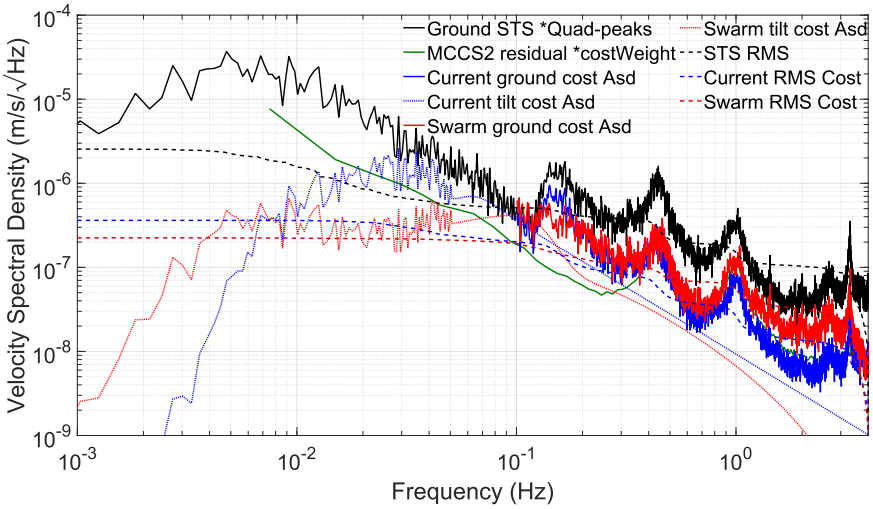

Figure A.3: Figure comparing the performance of the current sensor correction filter (blue) with the swarmed sensor correction filter (red) on ETMY at LHO in the X degree of freedom. The ground motion (black) and the MCSS2 residual (green) showing maximum possible subtraction from the CPS.

FIG A.4 shows the comparison between the current and swarmed sensor correction filter. The filter is broadly similar to the filter designed for ETMX-X, shown in FIG 6.7, it has a strong notch at the microseismic peak and strong tilt roll off below 10 mHz, while maintaining minimal gain peaking.

FIG A.5 shows the performance of the filter shown in FIG A.4 compared against the current sensor correction filter. The swarmed filter reduces the injected ground motion by around 20%. The filter offers around a factor of 1.5 better subtraction at the microseism and a factor of 2 less tilt injection between 30 and 60 mHz. At frequencies below 8 mHz the swarmed filter injects more tilt, though this is negligible to the overall RMS. The filter has a strong roll off of $\frac{1}{f^4}$ below 3 mHz, minimising the low frequency drift.

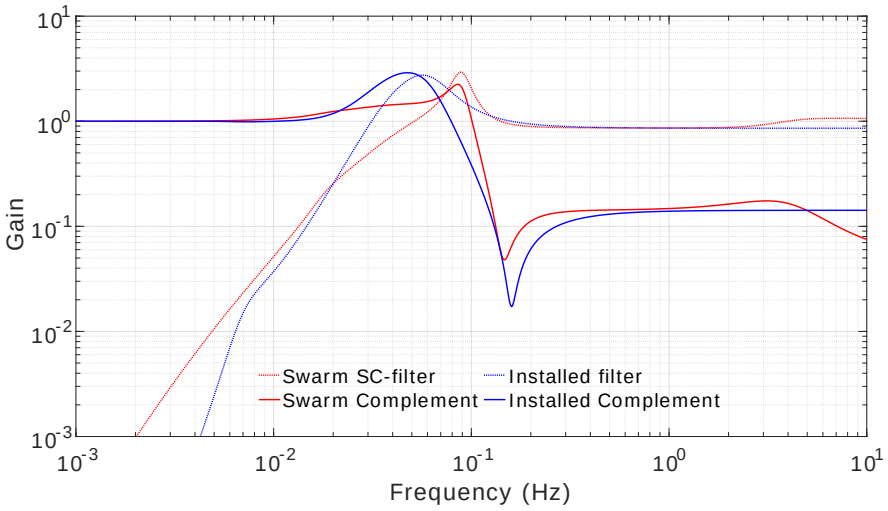

Figure A.4: Figure showing the response of the swarmed sensor correction filter (dashed red) with its complement (red) compared against the response of the current sensor correction filter (dashed blue) and its complement (blue) for ETMY in the Y degree of freedom at LHO

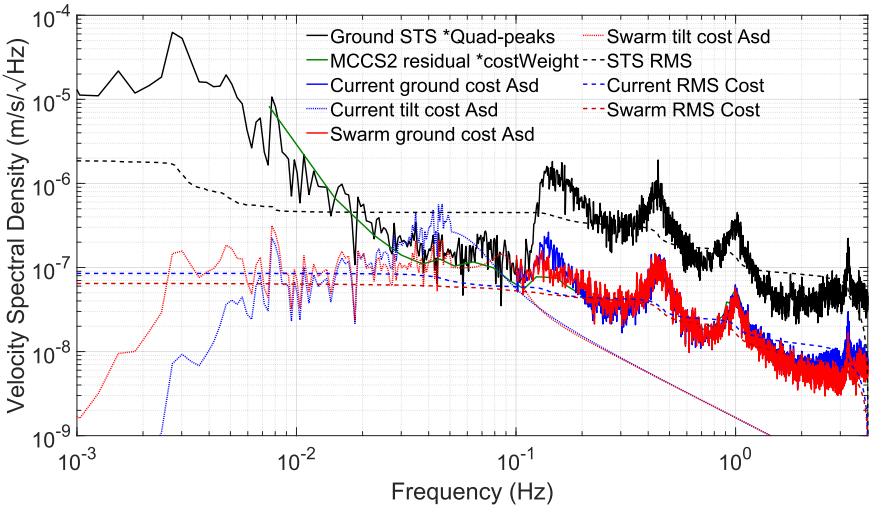

Figure A.5: Figure comparing the performance of the current sensor correction filter (blue) with the swarmed sensor correction filter (red) on ETMY at LHO in the Y degree of freedom. The ground motion (black) and the MCSS2 residual (green) showing maximum possible subtraction from the CPS.

Bibliography

[1] LIGO Scientific Collaboration et al. The sensitivity of the advanced ligo detectors at the beginning of gravitational wave astronomy. *arXiv preprint arXiv:1604.00439*, 2016.

[2] Conor Mow-Lowry and Denis Martynov. A 6d interferometric inertial isolation system. *Classical and Quantum Gravity*, 2019.

[3] Brian Lantz. Comments on the SRCL control bandwidth - HAM4 and HAM5 are guilty. `https://alog.ligo-la.caltech.edu/SEI/index.php?callRep=1393`, 2018.

[4] Conor Mow-Lowry, Sam Cooper, and Chris Collins. Interferometric local damping DCC: T1800049, 2017.

[5] Sam J Cooper, Chris J Collins, Anna C Green, et al. A compact, large-range interferometer for precision measurement and inertial sensing. *Classical and Quantum Gravity*, 35(9):095007, 2018.

[6] Rana Adhikari, Stefan Ballmer, Sam Finn, et al. http://lhocds.ligo-wa.caltech.edu:8000/advligo/GWINC.

[7] Fabián Erasmo Peña Arellano, Hasnain Panjwani, Ludovico Carbone, and Clive C. Speake. Interferometric measurement of angular motion. *Review of Scientific Instruments*, 84(4):043101, 2013.

[8] Stuart Aston. *Optical Read-out Techniques for the Control of Test-masses in Gravitational Wave Observatories*. PhD thesis, Unviersity of Birmingham, 2011.

[9] Daniel Clark and Brian Lantz. Replacement flexures for the GS-13 seismometer DCC: T0900089, 2009.

[10] LIGO Laboratory. Image of a prepped ISI, url =https://www.ligo.org/multimedia/gallery/sei.php.

[11] Jeff Kissel. Re-assessing HAM-ISI performance noise budget model for aLIGO DCC: T1300645, 2016.

[12] Brian Lantz and Jeffrey Kissel. Sensor noise estimates for advanced LIGO seismic isolation systems DCC: T0900450, 2009.

[13] Brett Shapiro. Brief introduction to the QUAD suspensions DCC: G1400964, 2014.

[14] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.

[15] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Tests of general relativity with gw150914. *Phys. Rev. Lett.*, 116:221101, May 2016.

[16] B. P. Abbott, R. Abbott, T. D. Abbott, et al. THE RATE OF BINARY BLACK HOLE MERGERS INFERRED FROM ADVANCED LIGO OBSERVATIONS SURROUNDING GW150914. *The Astrophysical Journal*, 833(1):L1, nov 2016.

[17] B. P. Abbott, R. Abbott, T. D. Abbott, et al. ASTROPHYSICAL IMPLICATIONS OF THE BINARY BLACK HOLE MERGER GW150914. *The Astrophysical Journal*, 818(2):L22, feb 2016.

[18] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.

[19] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Multi-messenger observations of a binary neutron star merger. *The Astrophysical Journal*, 848(2):L12, oct 2017.

[20] LIGO Scientific Collaboration, Virgo Collaboration, 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. *arXiv preprint arXiv:1811.12907*, 2018.

[21] LIGO Scientific Collaboration, Virgo Collaboration, 1M2H Collaboration, et al. A gravitational-wave standard siren measurement of the hubble constant. *Nature*, 551(7678):85–88, 2017.

[22] B. P. Abbott, R. Abbott, T. D. Abbott, et al. Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817a. *The Astrophysical Journal*, 848(2):L13, oct 2017.

[23] The LIGO Collaboration. GraceDB — gravitational wave candidate event database.

[24] Sam Cooper, Sam Morrell, Aaron Jones, Andreas Freise, and George Smetana. Chirp gravitational wave alerts. Nov 2019.

[25] Albert Einstein, Max Born, and Hedwig Born. *The Born-Einstein letters: friendship, politics, and physics in uncertain times. correspondence between Albert Einstein and Max and Hedwig Born from 1916 to 1955 with commentaries by Max Born.* Macmillan, 2005.

[26] A. Einstein and N. Rosen. On gravitational waves. *Journal of the Franklin Institute*, 223(1):43 – 54, 1937.

[27] J. Weber. Detection and generation of gravitational waves. *Phys. Rev.*, 117:306–313, Jan 1960.

[28] J. Weber. Evidence for discovery of gravitational radiation. *Phys. Rev. Lett.*, 22:1320–1324, Jun 1969.

[29] Russell A Hulse and Joseph H Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51–L53, 1975.

[30] LIGO Collaboration. A brief history of LIGO from www.ligo.org.

[31] Jorge L. Cervantes-Cota, Salvador Galindo-Uribarri, and George F. Smoot. A brief history of gravitational waves. *Universe*, 2(3), 2016.

[32] B Willke, P Aufmuth, C Aulbert, et al. The GEO 600 gravitational wave detector. *Classical and Quantum Gravity*, 19(7):1377–1387, mar 2002.

[33] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *Am. J.Sci. (3rd series)*, 34:333–345, 1887.

[34] J Aasi, B P Abbott, R Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015.

[35] Thomas Corbitt and Nergis Mavalvala. Review: Quantum noise in gravitational-wave interferometers. *Journal of Optics B: Quantum and Semiclassical Optics*, 6(8):S675–S683, jul 2004.

[36] J Abadie, BP Abbott, R Abbott, et al. A gravitational wave observatory operating beyond the quantum shot-noise limit. *Nature Physics*, 7(12):962, 2011.

[37] Gregory M. Harry, Helena Armandula, Eric Black, et al. Thermal noise from optical coatings in gravitational wave detectors. *Appl. Opt.*, 45(7):1569–1574, Mar 2006.

[38] Anna-Maria A van Veggel. Quasi-monolithic mirror suspensions in ground-based gravitational-wave detectors: an overview and look to the future. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2120):20170281, 2018.

[39] A V Cumming, A S Bell, L Barsotti, et al. Design and development of the advanced LIGO monolithic fused silica suspension. *Classical and Quantum Gravity*, 29(3):035003, jan 2012.

[40] Brett Shapiro, Rana X Adhikari, Odylio Aguiar, et al. Cryogenically cooled ultra low vibration silicon mirrors for gravitational wave observatories. *Cryogenics*, 81:83–92, 2017.

[41] G D Hammond, A V Cumming, J Hough, et al. Reducing the suspension thermal noise of advanced gravitational wave detectors. *Classical and Quantum Gravity*, 29(12):124009, jun 2012.

[42] Scott A. Hughes and Kip S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Phys. Rev. D*, 58:122002, Nov 1998.

[43] Jennifer C. Driggers, Jan Harms, and Rana X. Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Phys. Rev. D*, 86:102001, Nov 2012.

[44] Jan Harms, Edgard Bonilla, Michael Coughlin, et al. Linear ground-to-h(t) coupling at the corner station of the ligo hanford detector (manuscript under review). DCC: P1900275, 2019.

[45] Fabrice Matichard, Brian Lantz, Kenneth Mason, et al. Advanced ligo two-stage twelve-axis vibration isolation and positioning platform. part 1: Design and production overview. *Precision Engineering*, 40:273–286, 2015.

[46] Fabrice Matichard, Brian Lantz, Kenneth Mason, et al. Advanced ligo two-stage twelve-axis vibration isolation and positioning platform. part 2: Experimental investigation and tests results. *Precision engineering*, 40:287–297, 2015.

[47] Hang Yu, Denis Martynov, Salvatore Vitale, et al. Prospects for detecting gravitational waves at 5 hz with ground-based detectors. *Phys. Rev. Lett.*, 120:141102, Apr 2018.

[48] Denis V Martynov. *Lock acquisition and sensitivity analysis of advanced LIGO interferometers.* PhD thesis, California Institute of Technology, 2015.

[49] D. V. Martynov, E. D. Hall, B. P. Abbott, et al. Sensitivity of the advanced ligo detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93:112004, Jun 2016.

[50] Jennifer Watchi, Sam Cooper, Binlei Ding, Conor M. Mow-Lowry, and Christophe Collette. Contributed review: A review of compact interferometers. *Review of Scientific Instruments*, 89(12):121501, 2018.

[51] Charlotte Bond, Daniel Brown, Andreas Freise, and Kenneth A. Strain. Interferometer techniques for gravitational-wave detection. *Living Reviews in Relativity*, 19(1):3, Feb 2017.

[52] Charlotte Bond, Daniel Brown, Andreas Freise, and Kenneth A Strain. Interferometer techniques for gravitational-wave detection. *Living reviews in relativity*, 19(1):3, 2016.

[53] G. M. B. Bouricius and S. F. Clifford. An optical interferometer using polarization coding to obtain quadrature phase components. *Review of Scientific Instruments*, 41(12):1800–1803, 1970.

[54] C C Speake and S M Aston. An interferometric sensor for satellite drag-free control. *Classical and Quantum Gravity*, 22(10):S269, 2005.

[55] Fabián E. Pe na Arellano and Clive C. Speake. Mirror tilt immunity interferometry with a cat's eye retroreflector. *Appl. Opt.*, 50(7):981–991, Mar 2011.

[56] B. P. Abbott, R. Abbott, T. D. Abbott, et al. GW151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.*, 116:241103, Jun 2016.

[57] K. A. Strain and B. N. Shapiro. Damping and local control of mirror suspensions for laser interferometric gravitational wave detectors. *Review of Scientific Instruments*, 83(4):044501, 2012.

[58] L Carbone, S M Aston, R M Cutler, et al. Sensors and actuators for the advanced ligo mirror suspensions. *Classical and Quantum Gravity*, 29(11):115005, 2012.

[59] Alessandro Bertolini, Riccardo DeSalvo, Francesco Fidecaro, et al. Mechanical design of a single-axis monolithic accelerometer for advanced seismic attenuation systems. *Nuclear Instruments and Methods in Physics Research A*, 556, 2006.

[60] A. Miffre, M. Jacquey, M. Büchner, G. Trénec, and J. Vigué. Vibration-induced phase noise in mach–zehnder atom interferometers. *Applied Physics B*, 84(4):617–625, 2006.

[61] Min-Kang Zhou, Zhong-Kun Hu, Xiao-Chun Duan, et al. Performance of a cold-atom gravimeter with an active vibration isolator. *Phys. Rev. A*, 86:043630, Oct 2012.

[62] Min-Kang Zhou, Xin Xiong, Le-Le Chen, et al. Note: A three-dimension active vibration isolator for precision atom gravimeters. *Review of Scientific Instruments*, 86(4):046108, 2015.

[63] Christophe Collette, Stefan Janssens, Pablo Fernandez-Carmona, et al. Review: Inertial sensors for low frequency seismic vibration measurement. *Bulletin of the Seismological Society of America*, 102(4):1289–1300, 2012.

[64] Peter R. Saulson. Thermal noise in mechanical experiments. *Phys. Rev. D*, 42:2437–2445, Oct 1990.

[65] Aaron Barzilai, Tom VanZandt, and Tom Kenny. Technique for measurement of the noise of a sensor in the presence of large background signals. *Review of Scientific Instruments*, 69(7):2767–2772, 1998.

[66] Y. Dong, P. Zwahlen, A. M. Nguyen, R. Frosio, and F. Rudolf. Ultra-high precision MEMS accelerometer. *2011 16th International Solid-State Sensors, Actuators and Microsystems Conference, Beijing*, 2011.

[67] V. Josselin, P. Toubol, and R. Kielbasa. Capacitive detection scheme for space accelerometers applications. *Sensors and Actuators*, 78, 1999.

[68] Hareem Tariq, Akiteru Takamori, Flavio Vetrano, et al. The linear variable differential transformer (lvdt) position sensor for gravitational wave interferometer low-frequency controls. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 489(1):570 – 576, 2002.

[69] F Matichard, B Lantz, R Mittleman, et al. Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, 2015.

[70] D. M. Strayer, Ho Jung Paik, and M. Vol Moody. Short-range inverse-square law experiment in space. *Low Temperature Physics*, 29(9), 2003.

[71] Mark A. Zumberge, Jonathan Berger, Matthew A. Dzieciuch, and Robert L. Parker. Resolving quadrature fringes in real time. *Appl. Opt.*, 43(4):771–775, Feb 2004.

[72] Mark Zumberge, Jonathan Berger, Jose Otero, and Erhard Wielandt. An optical seismometer without force feedback. *Bulletin of the Seismological Society of America*, 100(2):598–605, 2010.

[73] Jennifer Watchi, Binlei Ding, Fabrice Matichard, and Christophe Collette. Development of a high-resolution optical inertial sensor for sub-hz seismic isolation. In *Proceedings of ISMA2016 USD2016*, pages 275–285. ISMA conference (Leuven, Belgium), 2016.

[74] T. B. Arp, C. A. Hagedorn, S. Schlamminger, and J. H. Gundlach. A reference-beam autocollimator with nanoradian sensitivity from mhz to khz and dynamic range of 107. *Review of Scientific Instruments*, 84(9):095007, 2013.

[75] Krishna Venkateswara, Charles A. Hagedorn, Matthew D. Turner, Trevor Arp, and Jens H. Gundlach. A high-precision mechanical absolute-rotation sensor. *Review of Scientific Instruments*, 85(1):015005, 2014.

[76] Nanometrics Inc. Trillium 240 seismometer user guide.

[77] B. Lantz, R. Schofield, B. O,ÄôReilly, D. E. Clark, and D. DeBra. Review: Requirements for a ground rotation sensor to improve advanced ligo. *Bulletin of the Seismological Society of America*, 99(2B):980–989, 2009.

[78] S M Aston, M A Barton, A S Bell, et al. Update on quadruple suspension design for advanced LIGO. *Classical and Quantum Gravity*, 29(23):235004, 2012.

[79] Brett Shapiro. Private Communication, 2017.

[80] Brian Lantz, Edgard Bonilla, and Conor Mow-Lowry. Thoughts on damping the testmass from the UIM DCC: T1800504, 2018.

[81] Patrick Kwee, Benno Willke, and Karsten Danzmann. Shot-noise-limited laser power stabilization with a high-power photodiode array. *Opt. Lett.*, 34(19):2912–2914, Oct 2009.

[82] Sam. J. Cooper, Jennifer Watchi, Christopher J Collins, et al. Quantifying cyclic nonlinearities of interferometers in the frequency domain DCC: P1900153, 2019.

[83] Michael J. Downs and K.W. Raine. An unmodulated bi-directional fringe-counting interferometer system for measuring displacement. *Precision Engineering*, 1(2):85 – 88, 1979.

[84] Norman Bobroff. Residual errors in laser interferometry from air turbulence and nonlinearity. *Appl. Opt.*, 26(13):2676–2682, Jul 1987.

[85] Christoph Weichert, Paul Köchert, Rainer Köning, et al. A heterodyne interferometer with periodic nonlinearities smaller than±10 pm. *Measurement Science and Technology*, 23(9):094005, 2012.

[86] Marco Pisani, Andrew Yacoot, Petr Balling, et al. Comparison of the performance of the next generation of optical interferometers. *Metrologia*, 49(4):455, 2012.

[87] Peter de Groot. Jones matrix analysis of high-precision displacement measuring interferometers. *Zygo*, 1(1):2, 1999.

[88] Jack A. Stone and Lowell P. Howard. A simple technique for observing periodic nonlinearities in michelson interferometers. *Precision Engineering*, 22(4):220 – 232, 1998.

[89] Peter L. M. Heydemann. Determination and correction of quadrature fringe measurement errors in interferometers. *Appl. Opt.*, 20(19):3382–3384, Oct 1981.

[90] NA Robertson, P Fritschel, B Shapiro, CI Torrie, and M Evans. Design of a tuned mass damper for high quality factor suspension modes in advanced ligo. *Review of Scientific Instruments*, 88(3):035117, 2017.

[91] Paul L. Rosin. A note on the least squares fitting of ellipses. *Pattern Recognition Letters*, 14(10):799 – 808, 1993.

[92] William A. Gardner. A unifying view of coherence in signal processing. *Signal Processing*, 29(2):113 – 140, 1992.

[93] Jonathan Berger, Peter Davis, Rudolf Widmer-Schnidrig, and Mark Zumberge. Performance of an optical seismometer from 1 μhz to 10 hz. *Bulletin of the Seismological Society of America*, 104(5):2422–2429, 2014.

[94] Brett Shapiro and Edgard Bonilla. Features in the quad state space model DCC: G1401132, 2014.

[95] R. Kirchhoff, C. M. Mow-Lowry, V. B. Adya, et al. Huddle test measurement of a near johnson noise limited geophone. *Review of Scientific Instruments*, 88(11):115008, 2017.

[96] Peter R. Saulson. Thermal noise in mechanical experiments. *Phys. Rev. D*, 42:2437–2445, Oct 1990.

[97] Sercel. L-4c brocure.

[98] Bruce Allen, Wensheng Hua, and Adrian Ottewill. Automatic cross-talk removal from multi-channel data. *arXiv preprint gr-qc/9909083*, 1999.

[99] Sam Cooper, Conor Mow-Lowry, Jim Warner, et al. Modeling of HAM ISIs DCC: T1800092, 2018.

[100] Sam Cooper, Brian Lantz, Jim Warner, Chiara Di Fronzo, and Conor Mow-Lowry. Some thoughts on controlling SRCL DCC: T1900107, 2019.

[101] Graham Clifford Goodwin, Stefan F Graebe, Mario E Salgado, et al. *Control system design*, volume 240. Prentice Hall New Jersey, 2001.

[102] Brett Shapiro. Intro to control theory for LIGO people. DCC: G1600726, 2016.

[103] Brian Lantz. Comments on the SRCL control bandwidth - HAM4 and HAM5 are guilty, 2018.

[104] Chiara DiFronzo, Sam Cooper, and Conor Mow-Lowry. Optical levers noise budget DCC: T1900105, 2019.

[105] B N Shapiro, R Adhikari, J Driggers, et al. Noise and control decoupling of Advanced LIGO suspensions. *Classical and Quantum Gravity*, 32(1):015004, 2015.

[106] C. M. Mow-Lowry, S. J. Cooper, C. Collins D. Hoyland, et al. Design review: compact interferometers for BS-SUS damping DCC: T1800448, 2018.

[107] C. M. Mow-Lowry, S. J. Cooper, and J. Briggs. Particle swarming sensor correction filters DCC: T1700541, 2017.

[108] Krishna Venkateswara. Livingston ground motion under high winds.

[109] Jim Warner Jeff Kissel. Mitigating wind at LHO DCC: G1600548, 2016.

[110] S Biscans, J Warner, R Mittleman, et al. Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors. *Classical and Quantum Gravity*, 35(5):055004, jan 2018.

[111] Brian Lantz. Revised HAM-ISI performance targets for advanced LIGO SRM chambers DCC: T1000216, 2015.