Interferometric Experiments Towards Advanced Gravitational Wave Detectors

John Robert Taylor BSc. (Hons)

Department of Physics and Astronomy
University of Glasgow

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

for Papa,

my favourite wee man . . .

and Orla,

my favourite wee girl

Contents

List of Figures

List of Tables

Acknowledgements

That's that then.

How many roads must a man walk down...? "Bloody loads" would appear to be the answer. This whole shenanigans has taken some time, and no small amount of effort. There are simply screeds of people I'd like to thank. As ever, I'll aim for brevity, but will no doubt miss and end up in over-loquaciousness. Here goes....

My mother, step-Dad, and sister deserve the first and foremost credit, not just for being there throughout this thesis, but for encouraging and supporting me my whole life. Orla – the newest member of the family – also merits thanks. Although she might not know it yet, her burbles and pops and winning smile have proved more than a little inspirational of late.

Mum, Derek, Joanna, Orla – my deepest thanks and love to you all.

Of course, there are also many work collegues to mention. My supervisor Ken Strain has provided a mind-bogglingly deep reservoir of knowledge and advice over the course of my Ph.D; and Mike Plissi, Sabina Huttner, Borja Sorazu and Nicolas Crouzet were all closely involved with my work. I'm indebted to their efforts.

My gratitude also goes to a host of other folk who either assisted work-wise (Stephen Craig, Colin Craig, Allan Latta, Neil Robertson, Mike Perreur-Lloyd, Russell Jones, Vickie Livingstone) or had the dubious honour of sharing an office with me (Christian Killow, Eoín Elliffe, Stuart Reid, Peter Murray, Alastair Grant). I'm also indebted to the STFC (formerly PPARC) and the University of Glasgow for funding my escapades, and to Jim Hough and the Institute for Gravitational Research in general.

And, oh yes – let's not forget Bryan Barr....

Bryan I think, deserves his own wee paragraph: after putting up with me for all these years, I think it's only fair. As a colleague, Bryan's knowledge of the subject and admirable suppression of his laughter reflex when fielding my silly questions were invaluable. As an office-mate, he joined in with my endless knuckle cracking and whistling, and as a friend ... well, he was always ready with a bottle of port and a packet of *HobNobs*. Cheers dude.

Above and beyond all other friends, I'm grateful to the *Old Skool*. John "*Columbo*" Veitch and Kenny "$13\,C^{\star\star}$" Walaron have been rattling off quality rancid chat (a.k.a *The Shite*) since we first rolled into *Room 257* almost a decade ago. Hilarity, insanity and borderline debauchery have never been too far away.

Likewise, Neil "*The Mole*" Thompson, RAWK "*Richard Codling*" CINEMA and Helen McGlone thoroughly deserve some long service medals.

And now to the *New Skool*:

Jen Toher was my bestest thesis-buddy, having the good sense to write up at the same time as me, while starting a year later (genius!). She was always ready to have some tea and cake, and share the pain. This was *very much* appreciated: I hope I helped her a little in return.

Ellie Chalkley put in a stonking shift as both flat-mate and friend, sharing my fondness for pretentious foodstuffs, Sunday afternoon telly, tea, and the odd bit of amateur scat jazz. Long may she continue to knit furry teal things.

Matt "*Stella Tops*" Pitkin's superb drunken ramblings, incessant air-drumming and affection for pop music provided much amusement. Thanks for being slightly more camp than me sometimes, old chap.

Siong "*Ballsack!*" Heng has given me hope, in the sense that I can now imagine a future in academia that contains actual amounts of fun, not to mention riotous Halloween parties.

Morag "*I'll have the salmon*" Casey , Dave "*Biblos!*" Crooks and Bryan "*Bryan*" Barr (again) – or *The Harbinger Three* as I like to think of them – expounded their own brand of caustic humour and sarcastic pedantry, which – combined with a penchant for rosé – livened up many an evening. Cheers!

Ross Galloway has been a constant source of erudite observation and Wildean wit – I wish him a hearty "*all the best*" in his future vocation.

Despite being *annoyingly young*, Hazel Bain has always been a much-needed force for positivity and non-apathy, although I never did quite understand her need to fling herself into the sea. She is silly....

Chris "*assumption is the mother of all fuck-ups*" Messenger gets a high-five for going out of his way to improve my time in Hannover over the past year. I'd probably have flipped out completely otherwise. Vielen Dank ya bass!

Thanks also to James Clark, for being a brother in curry, coffee, and whisky.

I'll also give a shout-out to Dave "*Snide Rat*" Petrie, Peter Sneddon, Fiona Speirits, Jennifer Pollock, Peter "*Crazy Pete*" Athron, and Samir Ferrag for either living with me without killing me, and/or just generally being ace.

Also, brief thanks to *Peckham's* (for midnight cakes), *M&S* (for daylight cakes), tea (for being awesome), LaTeX and BibTeX (contrary yet compelling), *I. J. Mellis* (oh the cheese!), and *Little Italy* (for tasty caffeinated sustenance).

Really though, it's the people that matter. You know who you are....

Thanks,

Bob

Preface

Chapter 1 introduces the phenomenon of gravitational waves and outlines the principles that lie behind interferometric gravitational wave detectors. Potential sources – of astronomical interest – are discussed and the response of a simple Michelson interferometer to gravitational waves given. Current detectors are detailed and the need for future detectors presented, along with a selection of the astrophysical results already produced by current generation detectors.

Chapter 2 details the myriad noise sources that inhibit the detection of gravitational waves with interferometric detectors. Noise due to seismic effects, thermal noise and photon noise is discussed, and suitable expressions for quantifying these processes given. In particular, the thermal noise from various heat-driven effects in the mirror substrates and coatings is discussed. Additionally, the control theory that underpins the successful operation of all gravitational wave interferometers is outlined.

Chapter 3 outlines an experiment for the direct interferometric measurement of the thermal noise spectrum of a test-mass mirror, of the sort typically used in gravitational wave interferometers. The measurement scheme of the experiment is introduced, the relevant limiting noise effects are described and estimated, and the final design sensitivity of the experiment presented.

Chapter 4 details the implementation of the experiment described in *Chapter 3*. The design and installation of the suspension system necessary to seismically isolate the test optics is described, as is the design and construction of a flame fibre pulling machine – necessary to produce suitable fused silica fibres for the test optic suspensions. The stabilisation of the frequency of the measurement laser is also detailed, as well as all the electronics, calibration and experiments performed; before the final sensitivity level is presented. The work presented in this chapter was performed in collaboration with Dr Mike Plissi, Dr Borja Sorazu and Dr Bryan Barr.

Chapter 5 recounts the experiments performed on the three-mirror coupled cavity system in the JIF lab. Different modulation schemes and their relative merits are discussed, modelled and measured; a novel and flexible optical modulation scheme is presented; and methods for optimising the relevant control signals outlined. The work presented in this chapter was performed in collaboration with Dr Bryan Barr and Dr Sabina Huttner.

Chapter 6 summarises the salient points of the previous chapters, and gives perspectives on possible implications and future related work.

Appendix A lists numerous materials properties and constants, primarily related to estimating the levels of thermal noise due to the various thermal dissipation mechanisms detailed in *Chapter 2*.

Appendix B contains schematics and discussion of the electronics used in both the thermal noise measurement, and coupled cavity control experiments.

Appendix C briefly summarises how to effectively measure the open and closed-loop transfer functions of interferometric control systems.

Summary

In 1905, Einstein postulated that the speed of light is not only finite, but that its speed in vacuum is a universal limit that *no* process can exceed. The *Theory of General Relativity* later extended this concept to include gravitational interactions, and Eddington's timely measurements of stellar positions during a solar eclipse in 1919 confirmed that gravity's effect on spacetime is both real and entirely physical – not merely a mathematical curiosity.

With the death of Newton's notions of universal time and instantaneous gravity came the idea of *gravitational waves* as distortions in space-time that propagate the gravitational interaction at the speed of light. These gravitational waves are emitted from any object undergoing a non-axi-symmetric acceleration of mass, but – due to the exceptionally weak coupling between gravitational waves and matter – are expected to induce displacements of the order of 10^{-18} m in kilometre-scale detectors: the extraordinary diminutiveness of this effect has thus far precluded any direct detection of the phenomenon.

Numerous gravitational wave detectors have been built since the 1960s, in the form of both interferometric detectors and resonant mass devices. Interferometric detectors currently represent the most promising form of detector, due to their relatively wide-band response to gravitational wave signals and promising levels of sensitivity. In recent years a worldwide network of these interferometric detectors (LIGO, GEO600, Virgo and TAMA300) have begun to approach (or indeed reach) their design sensitivities.

Although these detectors have started to provide upper limit results for gravitational wave emission that are of astrophysical significance, there have as yet been no direct detections. As such, work is underway to upgrade and improve these detectors. However, increasing the signal sensitivity necessarily leads to an increase in their sensitivity to their limiting noise sources. Two critical noise limits that must be char-

acterised, understood, and hopefully reduced for the benefit of future detectors, are thermal noise (from mirror substrates, reflective coatings and suspension systems) and photon noise – associated with the intrinsic shot noise of light and the noise due to light's radiation pressure.

Two interferometric experiments designed to help inform on these phenomena were constructed at the University of Glasgow's *Institute for Gravitational Research.*

The first experiment compared the relative displacement noise spectra of two specially constructed optical cavities, to extract the thermal noise spectrum of a single test mirror. In future experiments, this optic could be changed and the thermal noise spectrum for any suitable combination of mirror substrate and reflective coating evaluated.

The second experiment involved the investigation of suitable control schemes for a three-mirror coupled optical cavity. As the resonant light power in interferometers increases in future devices (in order to decrease the photon shot noise) the need to de-couple the control schemes that govern the respective cavities so that they can be controlled independently, becomes more important. As a three-mirror cavity effectively represents a simple coupled system, it provides a suitable test-bed for characterising suitable control schemes for more advanced interferometers.

Together, these experiments may provide information useful to the design of future gravitational wave interferometers.

When you are face to face with a difficulty,

you are up against a discovery

William Thomson

Baron Kelvin of Largs

Chapter 1

Gravitational Waves & Interferometry

1.1 Introduction

Wave phenomena permeate almost every aspect of science. Observations of oscillations in media provide information on both the nature of the medium and the properties of whatever process is occurring; and can be made, from a distance, without disturbing the source itself.

Seismic waves have been used to infer the internal structures of the Earth and the Sun, Alfvén waves transport energy in plasmas, longitudinal pressure waves in gas are interpreted as sound; and electromagnetic waves provide information on all processes that result in the emission of light.

The twentieth century saw the measurable electromagnetic spectrum widen from the purely optical to include infra-red, ultra-violet, radio, X-ray and γ-ray observa-

tions. Our picture of the universe improved accordingly; measurement of the cosmic microwave background [1] allowed the age of the universe to be quantified, and X-ray observations have proven the existence of black holes at the centre of our galaxy [2].

The prospect of examining the sky, in an entirely *new* and novel spectrum then, is an exciting one. According to the *Theory of General Relativity*, certain astrophysical events will channel significant amounts of energy, not into electromagnetic events, but into gravitational interactions. These astrophysical *gravitational wave* phenomena should have unique properties that distinguish them from their electromagnetic counterparts, and make them an invaluable tool for studying the universe in the twenty-first century and beyond.

1.2 Relativity & Gravitational Waves

Einstein's theory of General Relativity [3] recast the familiar notions of Newtonian gravity; classifying gravity as a geometrical distortion of a universal medium (termed *spacetime*) acting with finite speed, rather than an instantaneous force. Gravitational waves – oscillations in the background medium – are a physical consequence of this interpretation.

This viewpoint was initially explored by Einstein in 1905 in the *Theory of Special Relativity* [4], to allow co-ordinate frame transformations to agree with the Maxwellian view of light. An event in spacetime can be characterised by four co-ordinates (one temporal and three spatial); x^0, x^1, x^2, x^3.

The *invariant interval* between two events is then defined as

$$ds^2 = -(dx^0)^2 + \sum_{n=1}^{3} (dx^n)^2, \tag{1.1}$$

where ds is the invariant interval, dx^n is the difference in the nth co-ordinate between the two events $(x_2^n - x_1^n)$ and the time co-ordinate x^0 is measured in metres.

Compare this to an analogous expression for the distance (R) squared between two

events in Newtonian mechanics:

$$R^2 = \sum_{n=1}^{3} (dx^n)^2. \tag{1.2}$$

Observers at different points in spacetime obtain the same value for the invariant interval between any two events, in the same way that observers at different points in Newtonian space would measure the same physical distance between two events.

This concept can be expressed by rewriting *Equation 1.1* as

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu, \tag{1.3}$$

where $n_{\mu\nu}$ is the *Minkowski* metric for describing flat spacetime with no gravitational influence. In Cartesian co-ordinates this is given by,

$$n_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

In the presence of gravity, *Equation 1.3* can then be written as,

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu \tag{1.4}$$

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \tag{1.5}$$

where $h_{\mu\nu}$ is a perturbation from the flat Minkowski metric caused by e.g. gravitational waves and $g_{\mu\nu}$ is the resultant metric – or curvature – of spacetime.

According to General Relativity, gravitational radiation will be highly non-linear close to the source, but will conform to a simple wave equation in the weak-field limit

[5] :

$$\left(\nabla^2 - \frac{\partial^2}{(\partial x^0)^2}\right) h_{\mu\nu} = 0\,. \tag{1.6}$$

Noting that the x^0 co-ordinate can be expressed as $x^0 = ct$ leads to the more intuitive expression

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right) h_{\mu\nu} = 0\,, \tag{1.7}$$

delineating gravitational radiation as a wave phenomenon, propagating at the speed of light.

1.3 Properties of Gravitational Waves

Gravitational radiation from astrophysical sources has many features that distinguish it from electromagnetic radiation, and as such its observation may bring unique insights.

- The small coupling constant of gravity implies that gravitational radiation will interact weakly with surrounding matter. As such, any radiation will arrive at Earth essentially unhindered from its source, unlike electromagnetic information which is easily scattered and absorbed by intervening matter.

- Gravitational radiation emanates from *within* the host system, allowing a potentially unique view of the processes involved. Any electromagnetic radiation detected is typically emitted from the *outermost* active region.

- Gravitational radiation may allow us to see events that are highly energetic, but relatively weak electromagnetic emitters. These may include the mergers of black holes and neutron stars, and the motions of dark compact bodies – and may also include new phenomena.

No monopole or dipole gravitational emission is possible [6], with quadrupole being the minimum order of excitation required to generate gravitational waves. Thus, grav-

itational waves can only be generated when there is a non-axisymmetric acceleration of mass or energy [7].

Gravitational radiation has two separate polarisation components, termed h_+ and $h_\times$, rotated 45° with respect to each other: an arbitrary gravitational wave will consist of a superposition of these two components. A wave incident on an idealised ring of test-mass particles will alternately stretch and squeeze the system in perpendicular directions returning the ring to its original form after intervals of π and 2π in every 2π period cycle. See *Figure 1.1*.

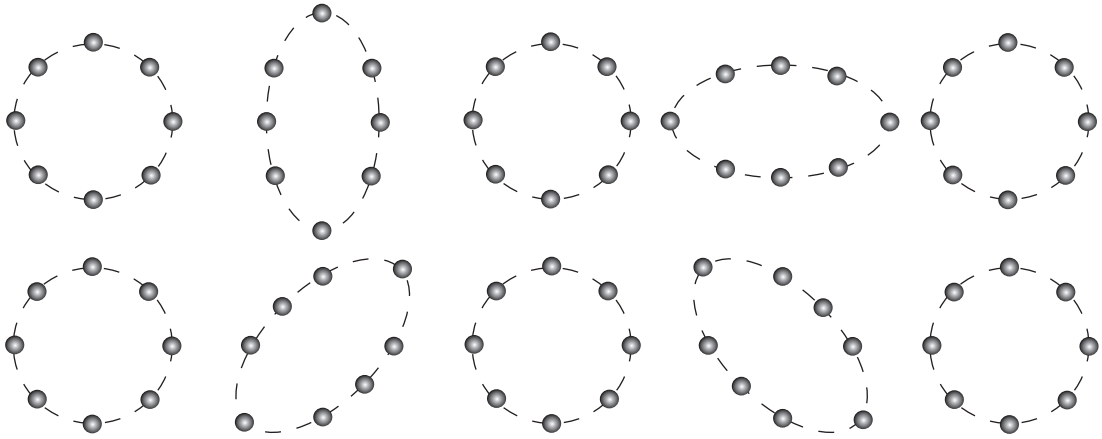

Figure 1.1: The effects of an h_+ (top) and $h_\times$ (lower) polarised gravitational wave, incident normal to the page, passing through two rings of test particles. The five stages depicted are 0, $\frac{\pi}{2}$, π, $\frac{3\pi}{2}$, 2π in the gravitational wave cycle, and the distortion effect is strongly exaggerated for visual clarity.

The amplitude of a gravitational wave, commonly termed the *strain* is defined (see *Section 1.5.2*) as,

$$h = \frac{2\Delta L}{L} \tag{1.8}$$

where h is the dimensionless strain amplitude, L is the the length considered (e.g. the diameter of the idealised ring of test-masses in *Figure 1.1*), and ΔL is the effective length change imparted by the passing gravitational wave, e.g. the major axis of the distorted ring in *Figure 1.1* minus the undistorted diameter.

1.4 Sources of Gravitational Waves

In summary, the presence of matter or energy curves spacetime, with any localised disturbances in the medium caused by non-axisymmetric accelerations of mass propagating as gravitational waves at the speed of light – see *Equation 1.7*.

1.4.1 Gravitational Radiation in the Laboratory

Although all massive objects radiate gravitational energy while undergoing non-axisymmetric accelerations, even extreme laboratory based experiments will not produce measurable effects. A massive iron cylinder of length l, rotating about its centre (with its ends tracing out a circle of diameter $d = l$) will emit gravitational radiation with a luminosity [6]

$$L_{\mathrm{gw}} = \left(\frac{ml^2\omega^3}{24}\right)^2 \left(\frac{G}{c^5}\right), \tag{1.9}$$

where G is the Newtonian gravitational constant, m is the mass of the cylinder, l is the length of the cylinder, ω the angular rotation frequency and c the speed of light. For a 5 m iron cylinder with a radius of 0.5 m rotating at 1 kHz (taking $\rho_{\mathrm{Fe}} = 7783$ kgm^{-3} [8]) this corresponds to an emitted luminosity of just $L_{\mathrm{gw}} = 2.86 \times 10^{-26}$ Js^{-1}.

Additionally, comparing this emission process to the rotational energy of the system yields the fractional rate of rotational energy lost to gravitational emission, neglecting decay of rotational frequency:

$$\frac{L_{\mathrm{gw}}}{E_{\mathrm{rot}}} = \frac{L_{\mathrm{gw}}}{\frac{ml^2\omega^2}{24}} \quad \simeq 10^{-36}\ \mathrm{s}^{-1}\,. \tag{1.10}$$

The creation of suitable observable radiation in the laboratory is unfeasible due to both the small magnitude of the energy emitted and the small transduction coefficient between the mechanical and gravitational processes.

However, astrophysical sources where high-energy accelerations of mass are involved may provide events of suitable luminosity to be observable on Earth. These sources may include supernovae, black hole and/or neutron star coalescences, oblate neutron

stars, and the Big Bang itself.

1.4.2 Burst Sources

Burst sources are short-lived, transient phenomena caused by catastrophic stellar events. These events may include supernovae and the coalescence of dense compact objects e.g. neutron stars, black holes and white dwarfs.

Supernovae

Supernovae are extremely luminous stellar events, which can exceed the brightness of their host galaxy. The absolute (electromagnetic) magnitude of a Type Ia supernova event can be [9] $M \simeq -19.5$, corresponding to a bolometric luminosity of the order of $5\times10^9\ L_\odot$.

Although extremely powerful, the emitted electromagnetic radiation is heavily scattered by surrounding envelopes of gas and dust: a gravitational observation could illuminate the poorly understood mechanisms behind these events and their evolution.

Type Ia supernovae consist of a white dwarf star accreting mass from a partner object, until it exceeds the Chandrasekhar limit ($M_{\mathrm{Ch}} = 1.38\,M_\odot$), causing the core to collapse – potentially forming a neutron star. These are not thought to produce significant amounts of gravitational radiation.

Type II supernovae however are caused by the radiation pressure in a large aging star failing to balance the gravitational force, causing the star to collapse, leaving a dense object; either a white dwarf or a neutron star. If this process is non-axisymmetric, observable gravitational waves may be emitted. Also, due to conservation of angular momentum, a rapidly spinning compact object will remain. If this resultant object has a suitably asymmetric distribution of mass, it may continue to emit gravitationally, "ringing down" until completely de-excited.

The gravitational wave strain h at Earth from such a source is estimated [10] to be,

$$h \simeq 5\times10^{-22}\left(\frac{E}{10^{-3}M_\odot c^2}\right)\left(\frac{15\ \mathrm{Mpc}}{r}\right)\left(\frac{1\ \mathrm{kHz}}{f}\right)\left(\frac{1\ \mathrm{ms}}{t}\right), \tag{1.11}$$

where E is the amount of energy emitted from the event gravitationally at frequency f in an observation time t, and r is the distance to the source in Mpc.

Coalescences

Two dense astronomical objects (e.g. black holes, neutron stars, white dwarfs) orbiting each other in a binary system will lose energy via gravitational wave emission. This situation is similar to the *gedanken* experiment outlined in *Section 1.4*, where rotational energy is lost in the system due to gravitational emission. The key difference here is that the energy loss causes an orbital decay, bringing the objects physically closer and increasing the orbital frequency, thus increasing the energy lost to gravitational waves and accelerating the orbital decay.

This process results in an energetic collision and eventual coalescence of the two objects. During the preceding inspiral phase, the emission will be of lesser amplitude but at an essentially constant frequency. When the objects approach, the gravitational wave emission level will rise, as will the frequency, culminating in a so-called "chirp" waveform. This is then thought to be followed by a ring-down phase as the objects settle, then possibly an ongoing periodic level of emission if the final object is notably asymmetric and rotating.

The detectable strain at Earth from two coalescing neutron stars at frequency f, and distance r from the Earth, is estimated to be [11],

$$h \simeq 1\times10^{-23}\left(\frac{100\ \text{Mpc}}{r}\right)\left(\frac{M_\text{b}}{1.2\ M_\odot}\right)^{\frac{5}{3}}\left(\frac{f}{200\ \text{Hz}}\right)^{\frac{2}{3}}, \tag{1.12}$$

where $M_\text{b} = (M_1M_2)^{\frac{3}{5}}/(M_1+M_2)^{\frac{1}{5}}$ is termed the *mass parameter* of the binary system, and the two objects have masses M_1 and M_2 respectively.

The gravitational luminosity of two solar mass objects in decaying circular orbits around each other is depicted in *Figure 1.2*.

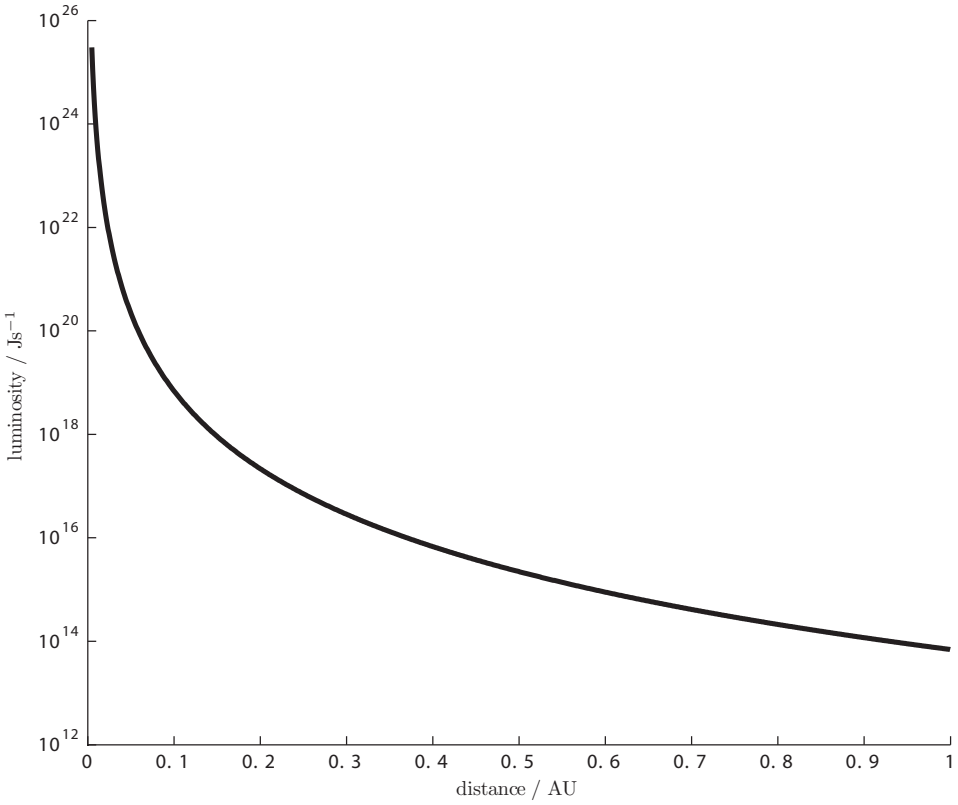

Figure 1.2: The gravitational wave luminosity emitted from a hypothetical source consisting of two stellar mass objects with orbital radii varying from 1 AU to 1 $R_\odot$. The gravitational wave luminosity becomes comparable to the electromagnetic luminosity ($L_\odot = 3\times10^{26}$ Js^{-1}) at small orbital radii.

1.4.3 Rotating Sources

Rotating sources are persistent and periodic sources such as rapidly rotating neutron stars and pulsars, characterised by a particular rotation frequency – where the rotation frequency of the object is half that of the emitted gravitational waves. Typical rotation periods for pulsars are in the region $\sim 2\,\text{ms} \rightarrow 5\,\text{s}$. To generate observable gravitational waves, the object must deviate from pure sphericity either through localised "mountains" or global oblateness, given by the equatorial ellipticity ϵ.

The strain amplitude at Earth from such an object is estimated [12] to be,

$$h \simeq 6\times10^{-25}\left(\frac{f}{500\text{ Hz}}\right)^2\left(\frac{1\text{ kpc}}{r}\right)\left(\frac{\epsilon}{10^{-6}}\right), \tag{1.13}$$

1.4.4 Binary Inspiral Sources

Binary inspirals occur when two dense objects (black holes and neutron stars being the most sensible candidates) lose energy in the form of gravitational radiation. The subsequent orbital decay leads to a greater loss of energy in the form of gravitational waves, until the bodies eventually collide and coalesce into a single entity, which may then continue to emit gravitationally as it de-excites and emits a ring-down waveform: see *Section 1.4.2*. The initial period of this process is a binary system, with a relatively stable orbital frequency. Although the gravitational waves emitted will be weaker than at the coalescence stage, the continuous orbit and predictable waveform allow for long signal integration times, increased signal to noise ratio and repeated verifiable observations.

1.4.5 Stochastic Sources

A stochastic gravitational wave background – analagous to the COBE and WMAP[1] results for electromagnetic radiation – should exist, consisting of low-level signals from unresolvable binaries in the galaxy and relic cosmological effects from processes occurring in the Planck epoch, shortly after the Big Bang. The gravitational force is thought to have de-coupled from the other forces at around this time ($t_{\text{Planck}} \simeq 10^{-43}$ s), so a measurement of this effect and any associated inhomogeneities may inform cosmological theories on how the universe began and developed, as well as provide an insight as to how the fundamental forces of nature interact. Most predictions state that any stochastic background will be present at a very low level ($h \sim 10^{-25}$) and low frequency. It will require multiple space-based detectors to detect, due to the incoherent superposition of signals, and low frequencies involved.

1.5 Detection Methods

Recall from *Figure 1.1* that a gravitational wave will cause an alternating perpendicular oscillation of the metric. To detect such a wave, this must be converted into an electrical

or mechanical signal that we can amplify and analyse.

1.5.1 Resonant Bars

In the 1960s, Joseph Weber began research on the use of resonant bar detectors. This method exploits the natural resonance modes of large metallic masses, typically aluminium cylinders. The relatively high quality factor – or Q – of aluminium gives it a particularly sharp mechanical resonance – responding strongly in a narrow frequency band. Bars were constructed with resonances in the stellar inspiral/burst range of $0.8 \rightarrow 1\,\mathrm{kHz}$, and fitted with piezo-electric transducers, converting the mechanical motion into an electrical signal that could then be amplified and processed.

Despite the apparent simplicity of the idea, complications and lack of sensitivity hampered Weber's efforts. Claims of coincident detections were not later verified by other experiments. Many other bar detectors now exist around the world (NIOBE[13], NAUTILUS[14], AURIGA[15], ALLEGRO[16], EXPLORER[14]) with peak sensitivites from the most sensitive of these instruments being of the order of $10^{-21}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$ [17][18].

Resonant mass detectors are hindered by several factors. It is difficult to construct, maintain and suitably isolate such a mass seismically, and cryogenic cooling is required to limit thermal noise effects. However, the most fundamental issue with resonant mass detection is one of bandwidth – resonant masses are intrinsically sensitive only to gravitational waves close to their mechanical resonant frequencies. Widening this bandwidth necessarily requires reducing the peak sensitivity.

1.5.2 Interferometry

In 1887, Michelson and Morley used what became known as a Michelson interferometer to ultimately disprove the existence of a universally pervasive etheric medium for the transmission of electromagnetic waves. Schematically, this device consisted of a light source, a 50% beam-splitter optic, and two mirrors at the ends of two equal length paths, perpendicular to each other. The light, having travelled down both – nominally

identical arms – was then recombined and the position of the resultant interference fringes examined to give a measure of the change in relative position of the two end mirrors. Michelson and Morley showed that there was no measurable drift in their measurement of this phase shift. The diurnal and annual motions of the Earth around the Sun would have caused such an effect, had an etheric medium existed.

The Michelson Interferometer

The inherent L shape of a Michelson type interferometer can be utilised in the detection of gravitational waves: a Michelson interferometer works by comparing the lengths of two perpendicular optical paths, and a gravitational wave alters the metric in just such a fashion; as one optical path is shortened, the other is lengthened, and vice-versa.

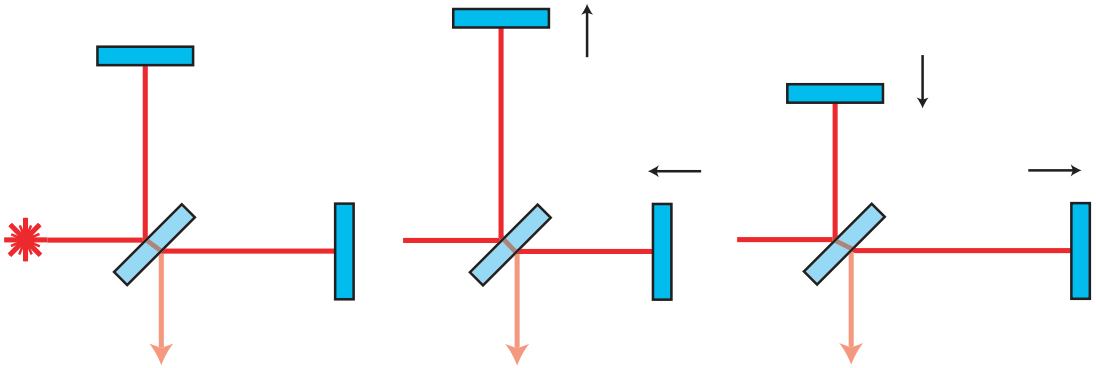

Figure 1.3: The effect of a suitably polarised gravitational wave on a Michelson interferometer at $0, \frac{\pi}{2}, \frac{3\pi}{2}$ stages of a wave cycle. Here, the components of the interferometer replace the idealised test-masses of *Figure 1.1*, with the beam-splitter optic at the centre. As one arm of the interferometer is contracted, the other is lengthened, with the reverse occurring later in the cycle.

If we consider an interferometric detector orientated such that the arms lie along the x and y axes, and a suitably orientated, linearly polarised gravitational wave is incident on the instrument along the z axis, we have (when the waveform $h_+(t)$ representing the gravitational wave is a maximum),

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1+|h_+| & 0 & 0 \\ 0 & 0 & 1-|h_+| & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

When events are measured by propagating photons, the invariant interval is $ds^2 = 0$. We then find,

$$(dx^0)^2 = (1+|h_+|)(dx^1)^2 + (1-|h_+|)(dx^2)^2 + (dx^3)^2\,. \tag{1.14}$$

Choosing the photons travelling in the x^1 direction ($dx^2 = dx^3 = 0$) we see that,

$$\frac{dx^1}{dx^0} = \frac{1}{\sqrt{1+|h_+|}} \quad \simeq \quad \frac{1}{1+\frac{|h_+|}{2}}\,. \tag{1.15}$$

Rearranging this expression then gives,

$$dx^0 = dx^1 + \left(\frac{|h_+|}{2}\right) dx^1\,, \tag{1.16}$$

where $|h_+| \ll 1$.

Defining $\Delta L = dx^0 - dx^1$ (i.e. the proper distance traversed by the photon in the presence of a gravitational wave, minus the length of the arm in the x direction) and setting $dx^1 = L$ we arrive at a definition of the dimensionless gravitational wave strain:

$$|h_+| = \frac{2\Delta L}{L}\,. \tag{1.17}$$

The same expression can be derived by following the same treatment for photons propagating in the perpendicular arm. So, the minimum gravitational wave strain detectable is linearly proportional to the arm length L of the Michelson interferometer.

If the light entering the interferometer along the x axis is a monochromatic field of

amplitude A and angular frequency ω then

$$E_{\text{in}} = Ae^{i(\omega t - kx + \phi)}, \tag{1.18}$$

where $k = \frac{2\pi}{\lambda}$ is termed the *wavenumber*, and λ is the wavelength of the light.

Setting the phase at the beam-splitter ($\phi = 0$) and noting that power $P = EE^*$, it can be shown [5] that the power at the output port is

$$P_{\text{out}}(\Phi) = \frac{P_{\text{in}}}{2}(1 + \cos(\Phi)), \tag{1.19}$$

where Φ is the relative phase difference between the two beams on recombination. A Michelson interferometer then acts as a transducer between gravitational induced phase shift, and detected light power.

Differentiating with respect to Φ gives

$$dP_{\text{out}} = \frac{P_{\text{in}}}{2}\sin(\Phi)d\Phi \quad = \quad \frac{P_{\text{in}}\Phi}{2}d\Phi, \tag{1.20}$$

where it is assumed the interferometer is operated close to the *dark fringe* point ($\Phi \simeq 0$), i.e. where most of the light returns to the input.

A suitably orientated sinusoidal gravitational wave of angular frequency ω_{g} and amplitude h interacting with an interferometer will have the following form:

$$h\cos(\omega_{\text{g}}t) = \frac{2\Delta L}{L}. \tag{1.21}$$

The total phase shift imparted due to the gravitational wave (on recombination at the beam-splitter, i.e. after a single round trip) is

$$d\Phi = k\Delta L \quad = \quad \frac{2\pi}{\lambda}\frac{hL}{\tau}\int_{t-\tau}^{\tau}\cos(\omega_{\text{g}}t)dt\,, \tag{1.22}$$

giving, finally, an expression for the imparted phase shift in terms of the angular frequency of the input light ω, the dimensionless strain amplitude h, the measurement time t, the angular frequency of the gravitational wave ω_{g} and the length of time for a

photon to traverse an arm, τ:

$$d\Phi = \frac{4\pi hL}{\lambda\tau\omega_{\mathrm{g}}} \sin\left(\frac{\omega_{\mathrm{g}}\tau}{2}\right) \cos\left[\omega_{\mathrm{g}}\left(t - \frac{\tau}{2}\right)\right] \tag{1.23}$$

$$= \frac{h\omega}{\omega_{\mathrm{g}}} \sin\left(\frac{\omega_{\mathrm{g}}\tau}{2}\right) \cos\left[\omega_{\mathrm{g}}\left(t - \frac{\tau}{2}\right)\right]. \tag{1.24}$$

As longer arm lengths lead to greater sensitivity to gravitational waves, it is desirable to increase the optical paths in the arms of the interferometer. From *Equation 1.24* it can be seen that the optimal storage time for a photon is $\tau = \frac{\pi}{\omega_{\mathrm{g}}}$, which implies an ideal arm length of $L = 200$ km for detection of a 1 kHz gravitational wave signal.

This can be partially achieved through physically building the arms as long as possible – with a functional limit of a few km. Lengthening the time τ that each photon spends in the interferometer arms can also be achieved through the use of *Herriott delay lines*, which fold a long optical path in one arm by reflecting the beam off of each end optic several times (at different spatial locations) or – more commonly – through the use of integrated *Fabry-Perot cavities*, which spatially superimpose multiple beams: see *Figure 1.4*.

If the interferometer is constructed with appropriate cavities giving the $\tau = \frac{\pi}{\omega_{\mathrm{g}}}$ condition then the phase shift given at the output of the interferometer – taking both arms into account – is given by,

$$d\Phi_{\mathrm{max}} = \frac{2h\omega}{\omega_{\mathrm{g}}} \sin(\omega_{\mathrm{g}}t), \tag{1.25}$$

and the maximum change in output power is now given by (see *Equation 1.20* previously)

$$dP_{\mathrm{out}}^{\mathrm{max}} = \frac{P_{\mathrm{in}}\Phi h\omega}{\omega_{\mathrm{g}}} \sin(\omega_{\mathrm{g}}t). \tag{1.26}$$

Optical recycling techniques are also employed to increase the sensitivity and shape the frequency response of detectors.

Recycling

Power recycling makes use of the fact that most of the light in a detector operating on the dark fringe is reflected back towards the input. This light is effectively wasted, but can be reflected back into the system with the addition of a power recycling optic placed between the light source and the beam-splitter optic. As an interferometer operating on the dark fringe is effectively a mirror, the power recycling optic forms a resonant Fabry-Perot cavity with the rest of the system; increasing the circulating power and thus the output signal, as according to *Equation 1.26*.

The technique of *signal recycling* can also be used, and also requires an additional optic, this time placed between the beam-splitter optic and the detection photodiode at the output port. This forms a resonant optical cavity between the output – where the signal light is present – and the interferometer as a whole. In *broadband* signal recycling, the signal recycling cavity can be made to resonate at a frequency $\omega \pm \omega_g$ where ω is the laser's angular frequency and ω_g the angular frequency of a target gravitational wave. The storage time of the signal photons is then dominated by the storage time of the signal recycling cavity, and can be enhanced accordingly, at the expense of the detector's bandwidth. The signal recycling cavity can also be tuned to be resonant with the laser light, increasing bandwidth at the expense of signal amplitude, or *detuned* to a position between these two states.

This allows the sensitivity response of a detector to be increased at specific frequencies, allowing a detector to search for known signals with established periodicities. A schematic of a Fabry-Perot Michelson interferometer with power and signal recycling can be seen in *Figure 1.4*.

1.5.3 Detectors

A worldwide network of large-scale interferometric detectors now exists. The Laser Interferometer Gravitational wave Observatory or "LIGO" [19] installations (in Hanford, Washington and Livingston, Louisiana) consist of two 4 km detectors, and one additional 2 km instrument housed in the same enclosure as the 4 km Hanford machine.

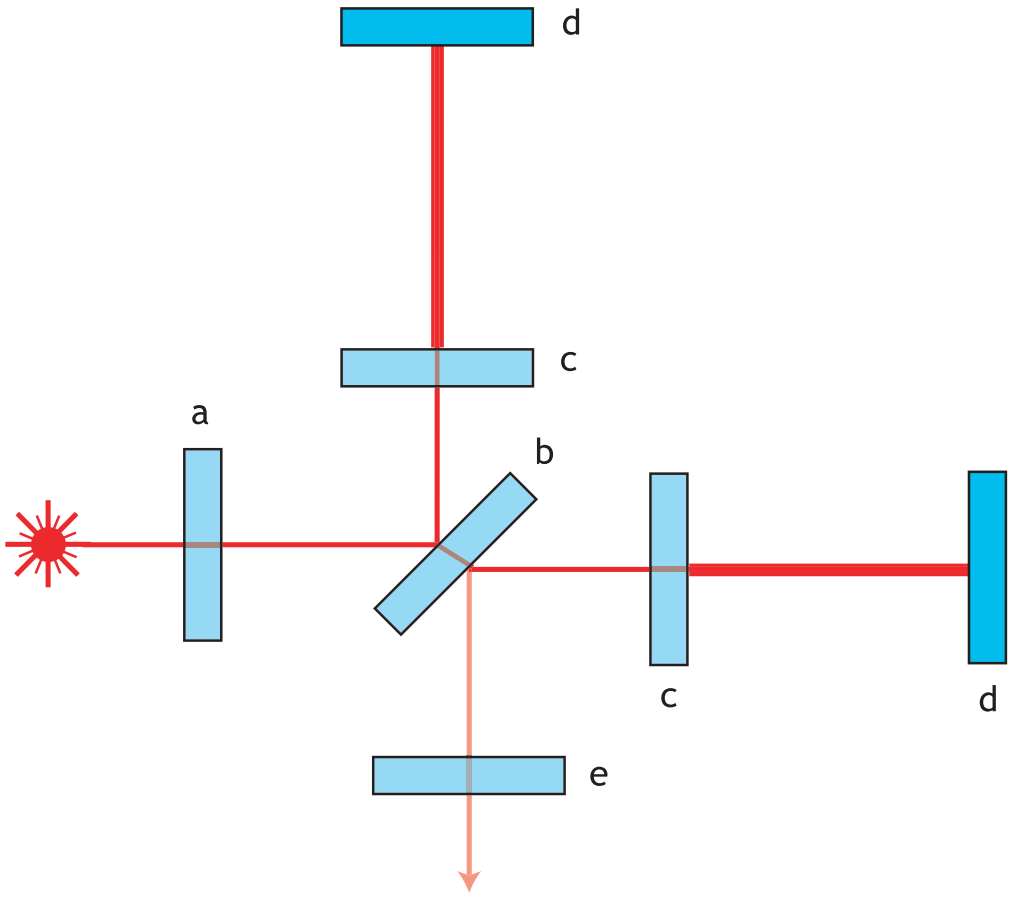

Figure 1.4: A schematic of a Michelson interferometer incorporating power recycling, signal recycling, and Fabry-Perot cavities in the arms. In order, the optical components are: a) power recycling mirror, b) beam-splitter optic, c) cavity input couplers, d) cavity end mirrors, e) signal recycling mirror.

This increases confidence in any signal detected (it is unlikely that an erroneous signal will appear in all detectors at the same time) and also provides directional information on the signal source. The LIGO detectors employ power recycling and Fabry-Perot cavities in the arms to increase sensitivity, and are currently operating at their design sensitivity of $h \sim 10^{-23}/\sqrt{\text{Hz}}$, where h is now the strain sensitivity: see *Figure 1.6*.

Virgo [20], in Cascina near Pisa, Italy, is a 3 km long system, featuring extensive seismic isolation for the main optical components to allow operation at a strain sensitivity of 10^{-22} $/\sqrt{\text{Hz}}$ down to 50 Hz. Similarly to LIGO, it uses power recycling techniques and Fabry-Perot cavities.

Figure 1.5: View northeast along Virgo's 3 km arm.

TAMA300 [21] is a 300 m arm-length interferometer based near Tokyo, Japan, and is the test-bed for a larger planned detector – the LCGT [22]. The LCGT (Large-scale Cryogenic Gravitational wave Telescope) baseline design specifies 3 km arm-length, cryogenic cooling of the major optical components and the use of sapphire substrates for the mirrors.

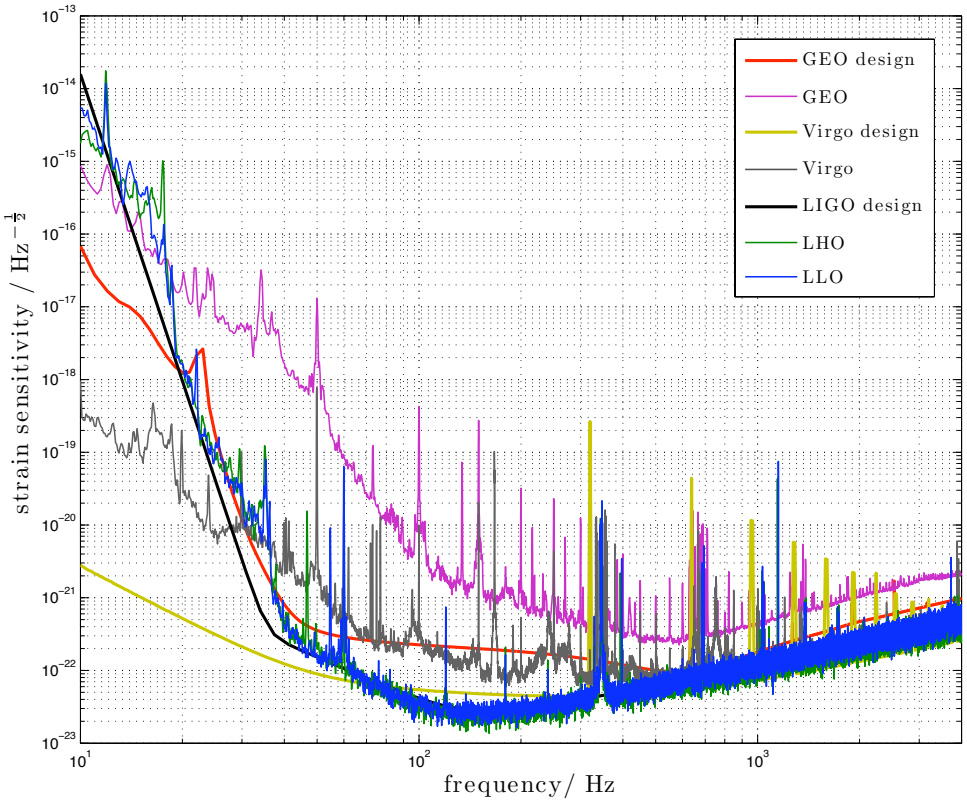

Figure 1.6: The calibrated strain sensitivities of the two 4 km LIGO instruments (green and blue), Virgo (grey) and GEO600 (magenta) as of May 2007. The design goals are shown in black, yellow and red respectively.

GEO600 (see *Figure 1.7*) is a British-German collaboration, situated near Hannover in Germany. With a 600 m baseline it is inherently less sensitive than the larger detectors, but employs several advanced techniques to allow it to compete in sensitivity.

Figure 1.7: The GEO600 interferometric gravitational wave detector, situated in Ruthe, near Hannover, Germany. The two perpendicular 600 m long arms can be seen, meeting in the central station which houses the input laser, beam-splitter, and other optics and control electronics.

GEO600 has pioneered the use of quasi-monolithic fused-silica suspension systems to suspend the interferometer optics, reducing the thermal noise effects associated with the wire suspensions used in other current generation detectors. It also features multiple-stage suspension systems, used to seismically isolate the mirrors and employs advanced optical techniques – signal and power recycling – to both boost the input power into the interferometer, and selectively amplify certain signal frequencies, thus reshaping the device's sensitivity curve. Both signal recycling and quasi-monolithic fused-silica suspension techniques are due to be incorporated in future detectors.

Although LIGO successfully meets its design criteria and sensitivity, it may not be sensitive enough to perform meaningful astronomy. Advanced LIGO, an upgrade due to be performing observations in 2014, will provide an order of magnitude increase in

sensitivity, with event rates for neutron star binary coalescences predicted to be between two a year, and three per day [23]. LIGO will also undergo an incremental upgrade – Enhanced LIGO – increasing sensitivity by a factor of ~ 2 before Advanced LIGO is fully implemented. As the time-averaged antenna response pattern for a Michelson interferometer to gravitational waves is almost spherically symmetric [5], this increase in sensitivity will increase the predicted event rate by a factor of ~ 8.

Virgo+ [24], a proposed sequential upgrade to the Virgo detector follows similar lines to Enhanced LIGO, upgrading the laser system and suspension setup to improve sensitivity. Similarly, there are plans to upgrade the Virgo detector to Advanced Virgo [24], with comparable sensitivity to Advanced LIGO.

GEO600 has plans to improve its system through upgrading its laser, optimising its sensitivity curve for high frequencies [25] (hence the name "GEO-HF"), and investigating new optical techniques; complementing the push towards low frequencies by other detectors.

Efforts are also underway to design and build a third-generation detector, called the Einstein gravitational wave Telescope [26] or "ET", utilising the knowledge gained in the research and development of current and next generation devices. This project is still in the design stage, but could provide an order of magnitude increase of sensitivity over even Advanced LIGO.

Another future detector is the space-based LISA [27] (Laser Interferometer Space Antenna). This consists of three identical spacecraft flying in a triangular formation, 20° behind the Earth, forming an interferometer with a 5×10^6 km arm-length. With seismic noise restrictions lifted, it should be able to observe gravitational waves at much lower frequencies (0.1 mHz $\rightarrow$ 0.1 Hz), and as such observe different sources such as galactic white dwarf binaries and supermassive black holes. LISA Pathfinder, a technology demonstration mission for LISA, is due to fly in 2010 [28].

1.6 Astrophysical Results

In 1993, Hulse and Taylor [29] were awarded the Nobel Prize for their work showing that the binary pulsar PSR1913+16 system was in an orbit decaying according to general relativistic predictions for gravitational wave emission. Although their results agreed with general relativity, they still relied on electromagnetic information, and as such are *indirect* evidence for gravitational emission. More recently, similar indirect evidence has been observed in the timing of quasi-periodic optical outbursts from quasar OJ 287; where the measured timing of the bursts agrees with the general relativistic predictions of energy loss due to gravitational emission in the system to within 10% [30].

The properties of the radio galaxy 3C 66B have been constrained by analysing long-term timing residuals from pulsar PSR B1855+09. The notion that the system may be a supermassive black-hole binary – as tentatively suggested from radio astronomy observations [31] – was ruled out with 95% confidence when no gravitational wave signature was detected in the pulsar timing data [32].

While a direct, confirmed detection remains elusive, several meaningful astrophysical results have already been obtained from the interferometric gravitational wave detectors currently operational. For instance: an upper limit of $h = 2.6 \times 10^{-25}$ has been placed on the gravitational emission from pulsar PSR J1603-7202, and the equatorial ellipticity of pulsar PSR J2124-3358 has been constrained to $\epsilon < 10^{-6}$ [33]. Additionally, data from LIGO's fifth science run imposed the limit that the gravitational wave power radiated from the Crab pulsar is no more than 4% of the total emitted power from its rotational spin-down [34].

Chapter 2

Noise & Control

2.1 Introduction

In order to detect gravitational wave induced strains, interferometers must be able to confidently detect displacement spectral densities of the order of 10^{-21} m/$\sqrt{\text{Hz}}$ or better, over a frequency range of a few tens of Hz to several hundred Hz. It is vitally important that all relevant noise sources that may limit the sensitivity of a gravitational wave detector are reliably characterised or estimated. As the fundamental aim is the detection of a phenomenon whose characteristics have yet to be observed, all non-signal variables must be accounted for to allow any sort of meaningful detection. If a detector operates with an unexplained noise peak above its design sensitivity, this can – philosophically at least – be equally well ascribed to a gravitational wave source as it can an arbitrary noise source, thereby reducing the probability of a reliable detection.

It is essential that the main optical components of the interferometer act as free, isolated masses, and that the interferometer be held on resonance (or "locked") for

long periods of time by e.g. controlling the position of one or more optics, or altering the frequency of the input laser.

2.2 Noise

There are several sources of noise that limit the sensitivity of gravitational wave interferometers. Broadly, these fall into three categories: seismic, thermal and photon noise.

The first fundamentally limits detection at low frequencies (typically between $\sim$ 10 Hz and 100 Hz) and – while it can be heavily attenuated – can never be completely removed from a ground-based instrument.

Thermal noise limits gravitational wave interferometers in their most sensitive regime, between the limit imposed by seismic noise and the effect of photon shot noise (which typically dominates the sensitivity above several hundred Hz) and manifests itself through several related, yet distinct, temperature dependent mechanisms, covered in more depth in *Section 2.2.2*. Theoretically these effects could be removed, or heavily reduced, through extensive cryogenic cooling.

Noise from the light source also inevitably couples in to the detected signal in various ways, discussed further in *Section 2.2.3*.

Additionally, a further limit to detection arises from gravitational gradient effects, detailed in *Section 2.2.4*.

2.2.1 Seismic Noise

An inevitable consequence of siting a detection system on the ground is that seismic motions may couple into the detection. In the case of a gravitational wave interfer-

ometer, this is especially important as the suspended masses are required to be *free* (along the axis of the propagating laser beam) in the sense that they are acted on by external (non-gravitational[1]) sources as little as possible in the detection band.

Seismic noise along any direction, at a suitably quiet site, can be roughly characterised by the following displacement noise spectrum above $\sim 1\,\mathrm{Hz}$ [35];

$$\tilde{x}_{\mathrm{seis}} = \left(\frac{10^{-7}}{f^2}\right)\ \mathrm{m}/\sqrt{\mathrm{Hz}}. \tag{2.1}$$

This implies that without any seismic isolation, the residual motion of an interferometer optic due to seismic effects will be $\sim 10^{-11}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$ at 100 Hz, some ten orders of magnitude above the $\sim 10^{-21}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$ required for gravitational wave detection.

Suspending the main interferometer optics as the lowest stage of a multiple pendulum system – in conjunction with a pre-isolation platform – provides the required level of seismic isolation. A train of n multiple pendulums will suppress the horizontal coupling of seismic noise by a factor dependent on the number of pendulums and their associated resonant frequency (or equivalently, length) as follows,

$$\tilde{x}_{\mathrm{seis}} = \left(\frac{{f_0}^2}{f^2}\right)^n \left(\frac{10^{-7}}{f^2}\right)\ \mathrm{m}/\sqrt{\mathrm{Hz}} \tag{2.2}$$

where f_0 is the resonance frequency of the pendulum, and $f_0 \ll f$. If we assume an optic is suspended as a train of three identical pendulums of 40 cm length, then the isolation supplied by the pendulums will reduce the effective seismic motion on the final optic to the required $\sim 10^{-21}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$ at $\sim 50\,\mathrm{Hz}$. Extending this further, a setup of 8 identical multiple pendulums pushes this limit down to $\sim 5\,\mathrm{Hz}$, where gravity gradient effects start to dominate – see *Section 2.2.4*. The suppression of seismic noise in the vertical direction (which over a 4 km detector will directly couple into horizontal motion with a factor of $\sim 0.1\%$ [36]) is achieved through the use of multiple stages of cantilever blade springs. These are designed to exhibit vertical resonances at lower frequencies than the suspension wires, and thus provide suitable isolation over the required range.

[1] gravity gradient noise being an unavoidable exception

Additionally, it should be noted that while pendulums may supply sufficient isolation at frequencies above their resonances, any motion near those resonances will be effectively amplified. Therefore, pendulum resonances must be kept outside of the detector's measurement band, and suitably damped.

2.2.2 Thermal Noise

Introduction

Thermal noise arises from the random thermally induced motions of the atoms comprising the interferometer's test-mass mirrors and suspension systems. The magnitude of this noise source can be reduced in two ways; lowering the temperature to directly reduce unwanted atomic motion; suitably shaping the energy spectrum of the thermal motion. The difficulties associated with cryogenically cooling large, isolated optical systems has meant that the latter method has been favoured for first and second generation gravitational wave detectors. Third generation detectors may use a combination of both approaches.

The mirrors in interferometric detectors are typically made from a high quality, low optical absorption substrate of fused-silica (SiO_2) and are then layered (see *Figure 2.1*) with alternating dielectric materials (e.g. SiO_2, Ta_2O_5, Al_2O_3, TiO_2) to form a highly reflective coating at the appropriate laser wavelength – 1064nm for an Nd:YAG laser – with very low levels of scattering and absorption. Thermal noise effects can arise from both the substrate material, and the coating, as well as from the elements used to suspend the mirrors.

Theory

In 1877, Ludwig Boltzmann connected the realm of the microscopic to the macroscopic through his formulation of the *Second Law of Thermodynamics*:

$$S = k_{\mathrm{b}} \ln \Omega \ . \tag{2.3}$$

The macroscopic quantity of entropy (S) is directly related to the number of available microscopic states (Ω) via Boltzmann's constant – k_b.

The generalised Nyquist relation (hereafter referred to as the Fluctuation Dissipation Theorem, or simply *FDT*) of Callen and Welten [37][2] similarly relates the *macroscopic dissipation* of a system to its level of *microscopic fluctuation*.

One can use this relation to express the displacement power spectral density of a substrate surface as [40],

$$\Xi(f) = \frac{k_\mathrm{b}T}{\pi^2 f^2}\Re[Y(f)] = \xi^2(f) \ , \tag{2.4}$$

where $\xi(f) \geq 0$ is then the associated amplitude spectral density in $\mathrm{m}/\sqrt{\mathrm{Hz}}$, T is temperature in K, f is frequency, and $Y(f)$ is the admittance function of a thermally driven mechanical oscillator.

The measured thermal noise contribution can then be directly estimated by defining the admittance function as

$$Y(f) = i\omega\frac{\xi(f))}{F(f)} \ , \tag{2.5}$$

where the external driving force F then comes from the force exerted on the front face of the optic by the laser beam. This then leads [41] to the following general expression for displacement power spectral density:

$$\Xi(f) = \frac{2k_\mathrm{b}T}{\pi^2 f^2}\frac{W_\mathrm{diss}}{F_0^2} \ . \tag{2.6}$$

This expression relates the average power dissipated (W_diss) from an arbitrary force (F_0) to the resulting displacement power spectral density. An expression for the displacement power spectral density due to any arbitrary effect can then be found, assuming the quantity W_diss resulting from the effect is known.

[2] further discussed by Callen and Greene in [38][39]

Mirror Substrates and Dielectric Coatings

Test-mass mirrors in interferometric detectors are typically constructed from a low loss, low optical absorption material such as fused silica, and then coated with a multi-layer stack of alternating thin dielectric films – see *Figure 2.1.*

If the substrate, and two coating materials have refractive indices n_s, n_1, n_2 (with $n_2 > n_1$), then the power reflectivity coefficient of the resultant mirror (with an even number of coating layers – n) will be [42],

$$R_n = \left(\frac{n_s \left(\frac{n_2}{n_1}\right)^n - 1}{n_s \left(\frac{n_2}{n_1}\right)^n + 1} \right)^2 . \tag{2.7}$$

The contributions towards displacement sensitivity of several thermal processes involving the coatings and substrate are outlined below, where Ξ_x denotes the resultant displacement power spectral density from a process x, and the radius of the laser beam (r_0) is taken to be at the point where the intensity of the beam is $\frac{P_{\max}}{e^2}$. Additionally, it is assumed that $r_0 \ll R_{\text{mass}}$, where R_{mass} is the radius of the test-mass optic.

Brownian Noise in Substrates and Coatings

Both the optic substrate, and the coating will have intrinsic levels of Brownian thermal motion. *Equation 2.6* can be extended [43][44] to produce,

$$\Xi_{\text{brown}}(f) = \left(\frac{2k_{\text{b}}T}{\pi^{\frac{3}{2}} f r_0 Y_{\text{s}}} \right) \phi_{\text{eff}} , \tag{2.8}$$

where r_0 is the radius of the laser beam, f is frequency, Y_{s} is the Young's modulus of the substrate material and ϕ_{eff} is the effective loss angle (in radians) of the combined substrate and coating stack.

In general, the loss angle ϕ of a material is related to the material's quality factor

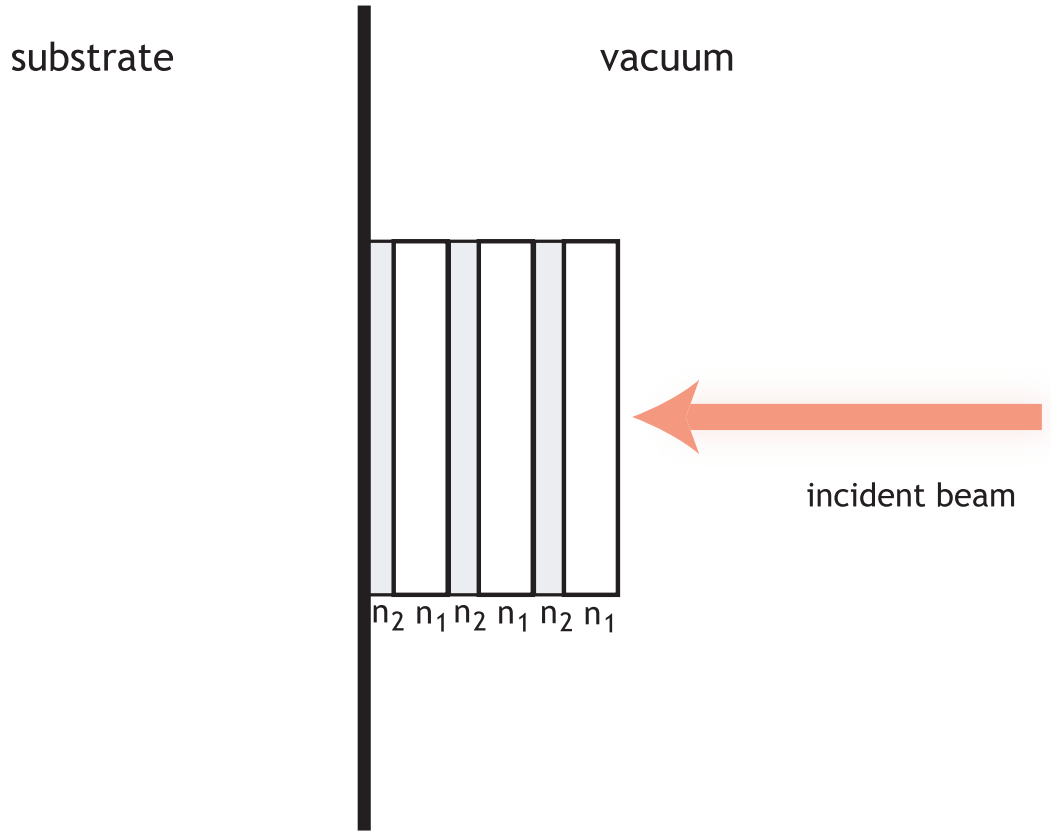

Figure 2.1: Schematic representation of mirror substrate with a multi-layered dielectric coating. The substrate (typically fused-silica) has refractive index n_s. The coating materials have refractive indices n_1 and n_2 respectively. Each layer of coating has the physical thickness $d_1 = \frac{d_{op}}{n_1}$ and $d_2 = \frac{d_{op}}{n_2}$ respectively, where the optical thickness d_{op} is typically $\frac{\lambda}{4}$, with λ being the laser wavelength.

Q as follows,

$$Q(\omega_0) = \phi(\omega_0)^{-1} = 2\pi \left(\frac{E_{\text{stored}}}{E_{\text{lost per cycle}}} \right) , \tag{2.9}$$

where ω_0 denotes a resonant frequency, and the final term represents the fractional energy lost in one cycle due to the dissipation.

The effective loss angle for a combined substrate and coating is,

$$\phi_{\text{eff}} = \phi_{\text{s}} + \frac{d}{r_0\sqrt{\pi}} \left(\frac{Y_{\parallel}\phi_{\parallel}}{Y_{\text{s}}} + \frac{Y_{\text{s}}\phi_{\perp}}{Y_{\perp}} \right) , \tag{2.10}$$

where ϕ_{s} is the mechanical loss angle of the substrate, d is the total physical thickness of the coating layer, and Y_{s} is the Young's modulus of the substrate. The quantities $Y_{\perp}$, $Y_{\parallel}$, $\phi_{\perp}$, $\phi_{\parallel}$ are defined as follows;

$$Y_{\parallel} = \frac{Y_1 d_1 + Y_2 d_2}{d} , \tag{2.11}$$

$$Y_{\perp} = \frac{d}{\frac{d_1}{Y_1} + \frac{d_2}{Y_2}} , \tag{2.12}$$

$$\phi_{\parallel} = \frac{Y_1\phi_1 d_1 + Y_2\phi_2 d_2}{dY_{\parallel}} , \tag{2.13}$$

$$\phi_{\perp} = Y_{\perp} \left(\frac{\phi_1 d_1}{Y_1} + \frac{\phi_2 d_2}{Y_2} \right) . \tag{2.14}$$

$Y_{\perp}$, $Y_{\parallel}$, $\phi_{\perp}$, $\phi_{\parallel}$ are the combined coating layer's Young's modulus and loss angles; parallel, or perpendicular, to the coating/substrate interface.

Y_1 and Y_2 are the individual coating material's Young's moduli, d_1 and d_2 are their respective physical thicknesses, and d is the total physical thickness of the coating.

Knowledge of these material properties can then provide an estimate of the Brownian thermal noise in the combined coating and optic substrate, according to *Equation 2.8*.

Thermoelastic Noise in Substrates and Coatings

Any material with a non-zero thermal expansion coefficient (α) will experience an inhomogeneous temperature distribution, associated with strain gradients in the material, when a localised force is applied. This effect causes a flow of heat energy in a test-mass, leading to dissipation. The thermal expansion coefficient is defined as,

$$\alpha = \frac{1}{l}\frac{dl}{dT} \ , \tag{2.15}$$

where dT is a change in temperature and dl the temperature induced change in length of a length l.

The displacement spectral density due to thermoelastic dissipation in the substrate of an optic is [45],

$$\Xi_{\text{sub thermo}}(\omega) \simeq \frac{8}{\sqrt{2\pi}} \frac{k_{\text{b}} T^2 \alpha_{\text{s}}^2 (1+\sigma_{\text{s}})^2 \kappa_{\text{s}}}{\rho_{\text{s}}^2 C_{\text{s}}^2 \frac{r_0^3}{\sqrt{2}} \omega^2} \ , \tag{2.16}$$

where α_{s}, σ_{s}, κ_{s}, ρ_{s} and C_{s} are the substrate's thermal expansion coefficient, Poisson ratio, thermal conductivity, density and specific heat capacity respectively.

Thermoelastic dissipation effects will also occur in the multi-layer dielectric coating stack. The power spectral density from a coating layer can be expressed as [46],

$$\Xi_{\text{coat thermo}}(f) \simeq \frac{8\sqrt{2} k_{\text{b}} T^2}{\pi\sqrt{2\pi f}} \frac{d^2}{2r_0^2} (1+\sigma_{\text{s}})^2 \frac{C_{\text{avg}}^2}{C_{\text{s}}^2} \frac{\alpha_{\text{s}}^2}{\sqrt{\kappa_{\text{s}} C_{\text{s}}}} \Delta^2 \tag{2.17}$$

where d is the total coating physical thickness and the quantity Δ^2 is,

$$\Delta^2 = \left\{ \frac{C_{\text{s}}}{2\alpha_{\text{s}} C_{\text{avg}}} \left(\frac{\alpha_{\text{avg}}}{1-\sigma_{\text{avg}}} \left[\frac{1+\sigma_{\text{avg}}}{1+\sigma_{\text{s}}} + (1-2\sigma_{\text{s}}) \frac{Y_{\text{avg}}}{Y_{\text{s}}} \right] \right) - 1 \right\}^2 \tag{2.18}$$

and a quantity X_{avg} is a weighted combination of the two material properties,

$$X_{\text{avg}} = \frac{d_1}{d} X_1 + \frac{d_2}{d} X_2 \ . \tag{2.19}$$

Thermorefractive Noise

Dissipation from thermorefractive effects are analogous to those from the thermoelastic process: the thermoelastic effect derives from a non-zero coefficient of thermal expansion, and the thermorefactive effect from a temperature dependent refractive index. Any temperature inhomogeneities in the optic will lead to localised fluctuations in the refractive index, and therefore the phase-front of any probing laser beam, causing an effective displacement noise to be induced in any measurement.

The displacement power spectral density from this effect in the mirror substrate is given by [47][48],

$$\Xi_{\text{thermoref}}(f) = \frac{\sqrt{2}\beta_0^2 k_{\text{b}} T^2 \lambda^2}{\pi r_0^2 \sqrt{2\pi\rho_{\text{s}} C_{\text{s}} \kappa_{\text{s}} f}} \tag{2.20}$$

where λ is the wavelength of the light used and the other symbols represent the quantities defined previously.

The temperature dependence of the refractive index is defined as,

$$\beta = \frac{dn}{dT} \ , \tag{2.21}$$

where n is the refractive index, and dT a change in temperature.

In *Equation 2.20* the term β_0 denotes the effective temperature dependence of the refractive index for either the substrate ($\beta_0 = \beta_{\text{s}}$) or a combined expression for the coating layer,

$$\beta_0 = \frac{n_1 n_2 (\beta_1 + \beta_2)}{4(n_1^2 - n_2^2)} \ , \tag{2.22}$$

where n_1, n_2, β_1, β_2 are the refractive indices and refractive index temperature dependencies of the two coating materials.

Suspension Thermal Noise

Another contribution towards the thermal noise spectrum comes from the mechanical dissipation of the pendulum suspension elements.

Pendulum Mode: The main pendulum mode of a suspended optic can be modelled as a harmonic oscillator, with a driving force $F(\omega)$, acting on a mass m with effective spring constant k,

$$F(\omega) = m\ddot{x} + k(1 + i\phi(\omega))x \ , \tag{2.23}$$

where ϕ is then the loss angle of the system. Applying the *FDT* gives the following expression [5] for pendulum mode displacement power spectral density:

$$\Xi_{\rm pend}(\omega) = \frac{4k_{\rm b}T}{\omega} \frac{\phi_{\rm p}(\omega)\omega_0^2}{m[w_0^4\phi_{\rm p}^2(\omega) + (\omega_0^2 - \omega^2)^2]} \ . \tag{2.24}$$

The loss angle of the pendulum ($\phi_{\rm p}$) is related to the loss angle of its constituent material ($\phi_{\rm s}$) [49],

$$\phi_{\rm p}(\omega_0) \simeq \phi_{\rm s}(\omega_0)\frac{n\sqrt{PY_{\rm s}I}}{2mgl} \ , \tag{2.25}$$

where P is the tension in one suspension element, n is the number of suspension elements, I is the moment of inertia of the suspension element ($I = \frac{\pi r^4}{4}$ for a cylindrical wire/fibre), $Y_{\rm s}$ is the material's Young's modulus, m is the mass of the optic, and l the length of the pendulum.

Violin Modes: The transverse vibrational modes (commonly termed *violin modes*) of the suspension elements will also have an impact on the thermally induced displacement noise. The violin modes form a slightly anharmonic series, with resonant frequencies defined [49] as,

$$f_n = \frac{n}{2l}\sqrt{\frac{P}{\rho_{\rm lin}}}\left[1 + \frac{2}{l}\sqrt{\frac{Y_{\rm s}I}{P}} + \frac{Y_{\rm s}I}{2P}\left(\frac{n\pi}{l}\right)^2\right] \ , \tag{2.26}$$

where f_n is the frequency of the nth mode, $\rho_{\rm lin}$ is the linear density (kg m^{-1}), and l is the length of the pendulum.

The periodicity of the violin modes implies that a suspension system designed to have the main pendulum resonance peak outside the detection band will still have violin resonances inside the detection band. The use of low loss suspension materials

such as fused-silica can help alleviate this by effecting very narrow resonances which can be identified and subsequently removed from measured data.

The contribution to the displacement noise spectral density from the nth violin mode of a cylindrical element of length l and radius r is given by [35],

$$\Xi_n^{\mathrm{viol}}(\omega) = \frac{8k_{\mathrm{b}}T\omega_0^2\phi_n}{m\omega[w_n^4\phi_n^2(\omega) + (\omega_n^2 - \omega^2)^2]} \ , \tag{2.27}$$

where the loss of the nth violin mode (ϕ_n) is,

$$\phi_n \simeq D_n^{-1}\left[\left(1 + \frac{8d_{\mathrm{s}}}{r}\right)\phi_{\mathrm{s}} + \phi_{\mathrm{nonlin}}\right] \ . \tag{2.28}$$

The term d_{s} denotes the dissipation depth – related to the typical depth of micro-cracks in the material surface, which contributes towards dissipation. For a fused-silica fibre produced in a flame pulling machine, this is $\sim 1\times10^{-4}$ m [35] [50].

The *dilution factor* D_n is given by,

$$D_n^{-1} = \frac{2}{l\sqrt{k}}\left[1 + \frac{\left(4 + \frac{(n\pi)^2}{2}\right)}{l\sqrt{k}}\right] \ , \tag{2.29}$$

with the value k being

$$k = \frac{P + \sqrt{P^2 + 4Y_{\mathrm{s}}I\rho_{\mathrm{lin}}\omega_n^2} + P}{2Y_{\mathrm{s}}I} \ . \tag{2.30}$$

Finally, the loss angle contribution from nonlinear thermoelastic effects (associated with thermal gradients from stress induced during fibre bending) is,

$$\phi_{\mathrm{nonlin}} = \frac{Y_{\mathrm{s}}T\left(\alpha_{\mathrm{s}} - \frac{P\gamma}{(\pi r^2 Y_{\mathrm{s}})}\right)}{\rho_{\mathrm{s}}C_{\mathrm{s}}}\ \frac{\omega\tau}{1 + (\omega\tau)^2} \ , \tag{2.31}$$

where

$$\tau = \frac{4r^2}{13.55\kappa_{\mathrm{s}}} \ , \tag{2.32}$$

and

$$\gamma_{\mathrm{s}} = \frac{1}{Y_{\mathrm{s}}}\frac{dY_{\mathrm{s}}}{dT} \ ; \tag{2.33}$$

where γ_s is the temperature dependancy of the material's Young's modulus.

2.2.3 Laser Noise

Intensity and Frequency Noise

Fluctuations in both intensity and frequency will couple into the output of an interferometric gravitational wave detector, adversely affecting the sensitivity. Frequency fluctuations can form from internal length changes in the laser cavity, and impose a limit on the detector's strain sensitivity,

$$h_{\text{freq}} = \frac{\Delta f}{f}\frac{\Delta L}{L} \tag{2.34}$$

where Δf is a fluctuation in frequency f, and ΔL is the offset between arms of length L.

Intensity fluctuations impose a similar limit,

$$h_{\text{P}} = \frac{\Delta P}{P}\frac{\Delta L}{L} \tag{2.35}$$

where P is the laser power, ΔP is the power fluctuation, L is the arm length, and ΔL is the length offset between the interferometer arms.

The Standard Quantum Limit

The *Standard Quantum Limit* represents the best possible measurement sensitivity as defined by *Heisenberg's Uncertainty Principle*. The limit describes the optimal operating boundary between the effects of radiation pressure noise, and photon-counting shot noise. The latter scales inversely with the square root of the input laser power (shot noise is reduced at higher powers) while the former increases with the square root of the input power – more photons implies more "back action" on the test-mass mirror. The optimum laser power for a given mass m in a simple Michelson interferometer is

given by [5],

$$P_{\text{opt}} = \pi c \lambda m f^2 \ . \tag{2.36}$$

Future detectors may use *squeezing* techniques to evade this limit. The uncertainty principle can be expressed as,

$$\Delta n \Delta \phi \geq \frac{1}{2} \ , \tag{2.37}$$

where Δn is the uncertainty in the number of photons, and $\Delta \phi$ is the uncertainty in phase. Light noise may be "squeezed" in one quadrature ($\Delta \phi$ for phase measurements), at the expense of the other, thus decreasing phase noise below that allowed by the standard quantum limit while still obeying the uncertainty principle. The limits imposed by shot noise are dealt with in more detail in *Section 2.3.1*.

2.2.4 Gravity Gradient Noise

A final limit to the sensitivity of gravitational wave detectors comes in the form of gravity gradient noise. This limit arises from the changing mass distribution in and around the interferometer altering the local gravitational field, which then couples directly into the positions of the test-mass optics. This effect is small, but may fundamentally limit ground-based detectors, defining their low-frequency measurement regime [5].

The seismic induced gravitational gradient limit for the LIGO Hanford observatory, in terms of strain sensitivity, has been estimated [51] as $\sim 6 \times 10^{-23} \ \ /\sqrt{\text{Hz}}$ at 10 Hz. While below LIGO's current sensitivity at this frequency (see *Figure 1.6*) this effect may affect the potential sensitivities of future instruments.

2.3 Control

In order to operate an interferometer, it is essential that the optical cavities within the system are controlled, and held to the correct operating point. The Fabry-Perot cavities in interferometers are typically held on resonance through use of the *Pound-Drever-Hall* [52] scheme (also known as *RF Reflection locking*) or extensions thereof, and controlled using negative feedback systems. Both of these concepts are outlined below.

2.3.1 Fabry-Perot Cavities

Background

Fabry-Perot cavities are essential components of gravitational wave interferometers, and consist of two partially transmitting (but typically highly reflective) mirrors. Here, ρ_1 and ρ_2 are the respective mirror's amplitude reflection coefficients, and τ_1, τ_2 the amplitude transmission coefficients, defined such that,

$$\rho_n^2 + \tau_n^2 = 1 \ . \tag{2.38}$$

Figure 2.2: A schematic of a Fabry-Perot cavity. Mirrors 1 and 2 (with amplitude reflection and transmission coefficients $\rho_1, \rho_2, \tau_1, \tau_2$) are separated by a length L, with their reflective surfaces facing each other. The light reflected from the system is separated from the input light with a polarising beam-splitter and quarter wave-plate setup.

If the input field (setting the phase at the input coupler to zero) is given by,

$$E_{\text{in}} = E_0 e^{i\omega t} \ , \tag{2.39}$$

where ω is the angular frequency of the laser and t is time, then the light field transmitted through the cavity can be found by considering the superimposed fields;

$$E_{\text{trans}} = E_{\text{in}} \left(\tau_1 \tau_2 e^{-ikL} + \tau_1 \rho_2 \rho_1 \tau_2 e^{-3ikL} + \tau_1 \rho_2 \rho_1 \rho_2 \rho_1 \tau_2 e^{-5ikL} + \ldots \right) . \tag{2.40}$$

The first term in parentheses represents the field transmitted straight through the system, the second term the field that has reflected off the second mirror, then returned from the first, and so on. This can be simplified by summing the terms, using the identity $\sum_{n=0}^{\infty} a^n = \frac{1}{1-a}$ (where $a < 1$),

$$E_{\text{trans}} = E_{\text{in}} \left[\tau_1 \tau_2 e^{-ikL} \sum_{n=0}^{\infty} \left(\rho_1 \rho_2 e^{-i2kL} \right)^n \right] = E_{\text{in}} \left(\frac{\tau_1 \tau_2 e^{-ikL}}{1 - \rho_1 \rho_2 e^{-2ikL}} \right) . \tag{2.41}$$

A similar expression can be found for the reflected field:

$$E_{\text{ref}} = E_{\text{in}} \left[-\rho_1 + \tau_1^2 \rho_2 \sum_{n=0}^{\infty} \left(\rho_1 \rho_2 e^{-i2kL} \right)^n \right] = E_{\text{in}} \left(-\rho_1 + \frac{\tau_1^2 \rho_2}{1 - \rho_1 \rho_2 e^{-i2kL}} \right) . \tag{2.42}$$

We can now plot the transmitted field as a function of one of the mirror's offset from resonance[3] – see *Figure 2.3*. Here, as the second mirror is shifted through 40°, the amplitude of the transmitted field is seen rising, reaching a maximum as the phase offset becomes zero at the resonance. This 40° range is equivalent to $\sim 1.2 \times 10^{-7}$ m offset in cavity length.

It is now useful to define some parameters commonly used when dealing with optical cavities.

The *finesse* of a cavity is given by,

$$\mathcal{F} = \frac{\pi \sqrt{\rho_1 \rho_2}}{1 - \rho_1 \rho_2} = \frac{FSR}{FWHM} , \tag{2.43}$$

where FSR is the *free spectral range* of the cavity (the separation between distinct

[3] the optical simulation software FINESSE [53] is used for these simulations

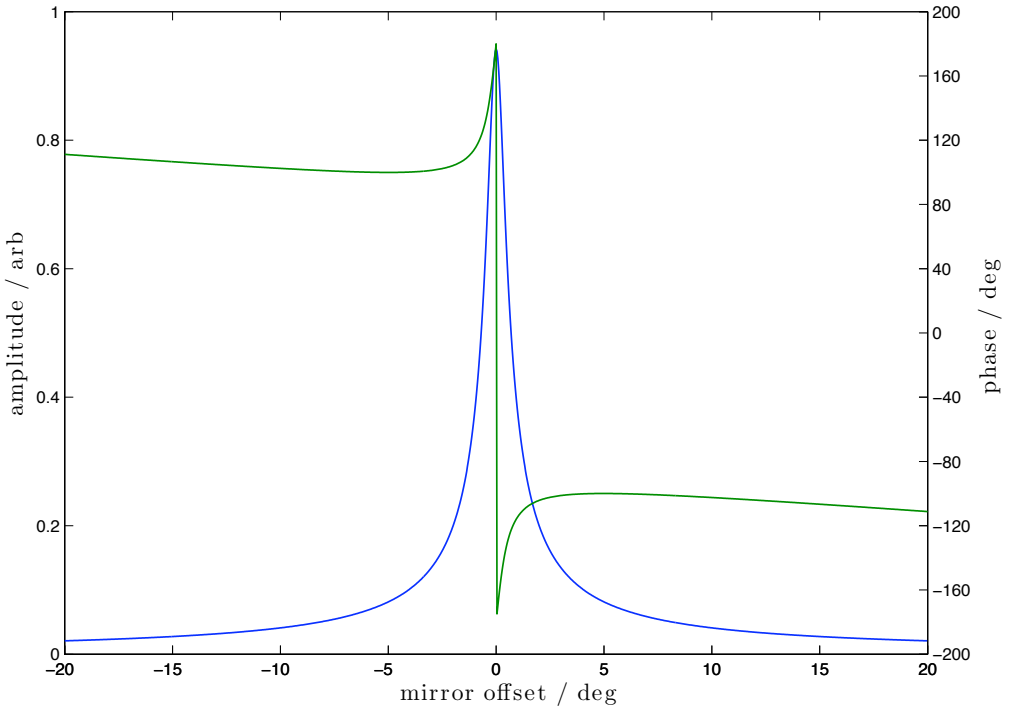

Figure 2.3: The amplitude and phase evolution of the throughput light of a Fabry-Perot cavity over a single resonance. The blue trace shows the amplitude of the field, the green trace depicts the phase. The cavity in this example consists of two mirrors, the first with an amplitude reflectivity $\rho_1 = 0.9899$, the second with $\rho_2 = 0.995$, separated by $L = 10\,\mathrm{m}$.

resonances in Hz) and *FWHM* is the "full-width half-maximum" measure of a single resonance feature, also measured in Hz – see *Figure 2.4*. The *FWHM* value is also termed the *linewidth* of the cavity, and is twice the value of the cavity *bandwidth*.

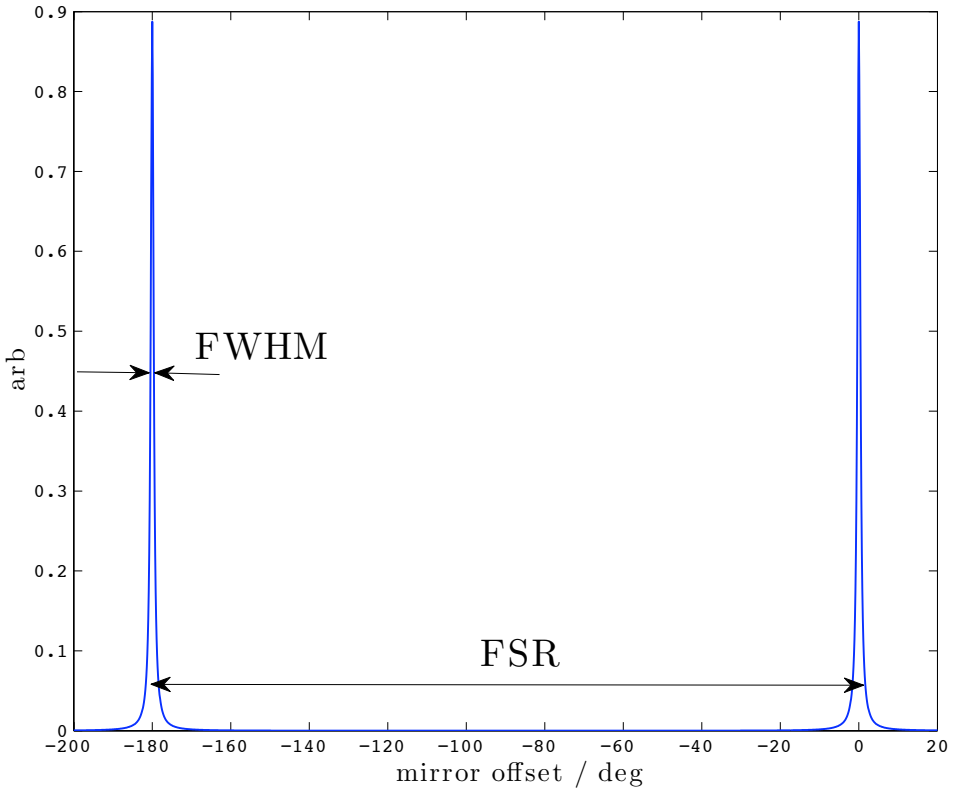

Figure 2.4: Depiction of the free spectral range (FSR) and linewidth (FWHM) of the same cavity as in *Figure 2.3*. Transmitted power is shown in blue.

The FSR is given by,

$$FSR = \frac{c}{2L} \ , \tag{2.44}$$

where L is the length of the cavity, and the free spectral range is then the inverse of the time for a photon to traverse the cavity and return.

Modulation and Locking

Fabry-Perot cavities can be held on resonance through use of the technique known as *Pound-Drever-Hall* locking. This involves imposing a set of *sidebands* on to the main

carrier light, by either amplitude or phase modulating the carrier beam.

As a cavity will (ideally) only allow light within its own linewidth to resonate, the sideband fields – which are offset from the main carrier in frequency – will reflect from the front face of the cavity when on resonance. This provides a static phase reference, so that light returning from the cavity (which may have additional phase imparted to it from a gravitational wave, or other effect) will then interfere with the sideband field, causing a beat at the sideband modulation frequency. This effect is then detectable using a standard photodiode and mixer arrangement, where the photodiode acts as a square-law detector of the light field and the mixer is used to extract the component oscillating at the modulation frequency. The demodulated signal provides a bi-polar error signal which can be used to feed back to either the length of the cavity, or the frequency of the laser, to hold the system on resonance.

Optical modulators (in the form of $LiNbO_3$ crystals) are used to create the phase modulation sidebands for this technique, typically at RF frequencies $\gtrsim 10$ MHz.

A sinusoidal phase modulation imposed on the input beam can be expressed as,

$$E_{\mathrm{mod}} = E_0 e^{i(\omega t + m\cos(\omega_{\mathrm{m}} t))} \ , \tag{2.45}$$

where E_0 is the input laser field, ω is the laser's angular frequency, t is time, β is the *modulation index*, ω_{m} is the angular frequency of the modulation and E_{mod} is the resultant modulated field.

Using the *Jacobi-Anger* identity,

$$e^{iz\cos\theta} = \sum_{n=-\infty}^{\infty} i^n J_n(z) e^{in\theta} \ , \tag{2.46}$$

this can then be rewritten as,

$$E_{\mathrm{mod}} = E_0 e^{i\omega t}\left[J_0(m) + J_1(m)e^{i\omega_{\mathrm{m}} t} - J_1(m)e^{-i\omega_{\mathrm{m}} t}\right] \ , \tag{2.47}$$

where higher order sidebands have been omitted (i.e. assuming $m \ll 1$) and the $J_n(m)$ terms are Bessel functions of the first kind, defined as,

$$J_n(m) = \sum_{x=0}^{\infty} \frac{(-1)^x}{x!(x+n)!} \left(\frac{m}{2}\right)^{n+2x} \quad . \tag{2.48}$$

A modulation index of $m = 0.2$ then, implies that a factor of $\sim \left(\frac{m}{2}\right)^2 = 1\%$ of the original carrier power will be directed into each sideband.

The superposition of phase-shifted carrier and phase-reference sideband field is then detected on a photodiode and demodulated using a mixer and a local oscillator of the same frequency as the phase modulation. When the phase of the local oscillator is set correctly, the output from the mixer will be an appropriate bi-polar error signal: see *Figure 2.5*.

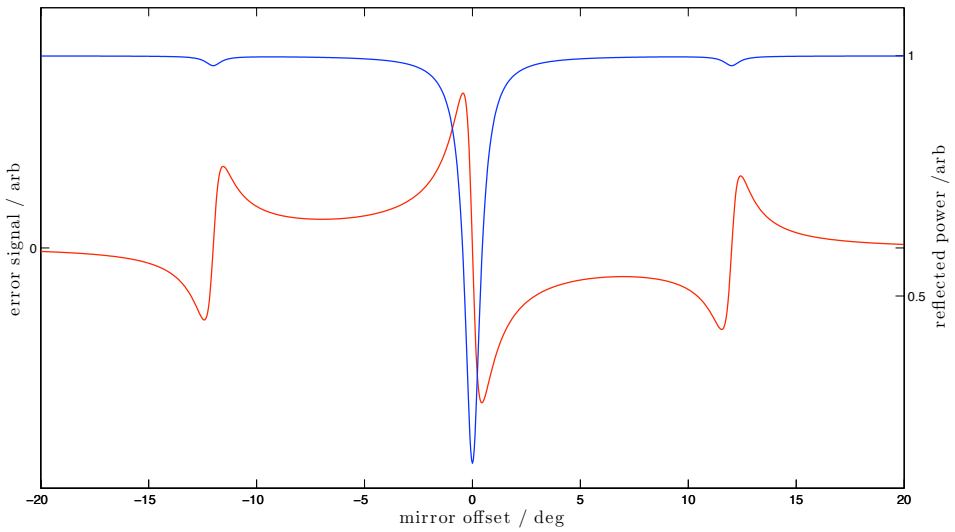

Figure 2.5: A typical *Pound-Drever-Hall* error signal for a Fabry-Perot cavity. Here, the cavity is identical to previous examples, but with a 1 MHz phase modulation added to the input field, with modulation index $m = 0.2$. The blue trace is the reflected power from the cavity, the red trace is the demodulated output from an in-phase mixer, connected to a photodiode detecting the return light from the cavity.

When the phase of the local oscillator is correctly set, the slope of the error signal for the carrier light resonance is maximised. When the local oscillator is 90° out of phase, the error signal slopes for the sidebands are then emphasised: see *Figure 2.6*.

It is possible then, to use a combination of phase modulators, mixers, and photo-

diodes to produce a suitable error signal for either feeding back to the light source – thus stabilising the frequency of the laser to the length of the cavity – or feeding back to the length of the cavity, effectively measuring the displacement noise of the cavity with reference to the stability of the laser frequency: see *Figure 2.7*.

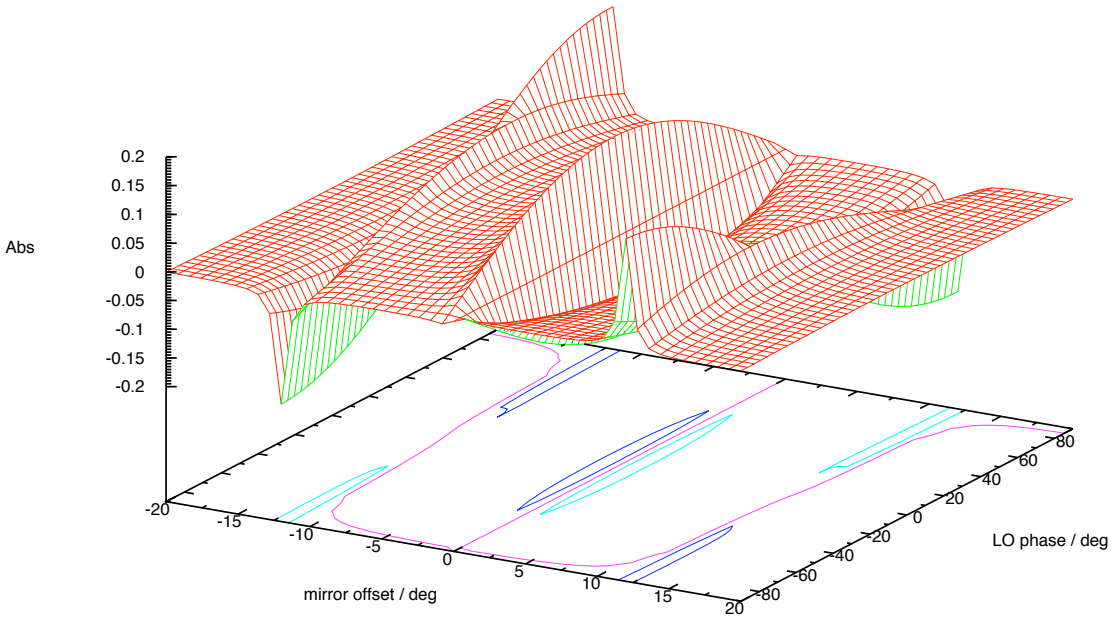

Figure 2.6: Variance of the *Pound-Drever-Hall* error signal with local oscillator phase. The cavity properties here are identical to those in *Figure 2.5*, which corresponds to the error signal when the local oscillator phase is 0°.

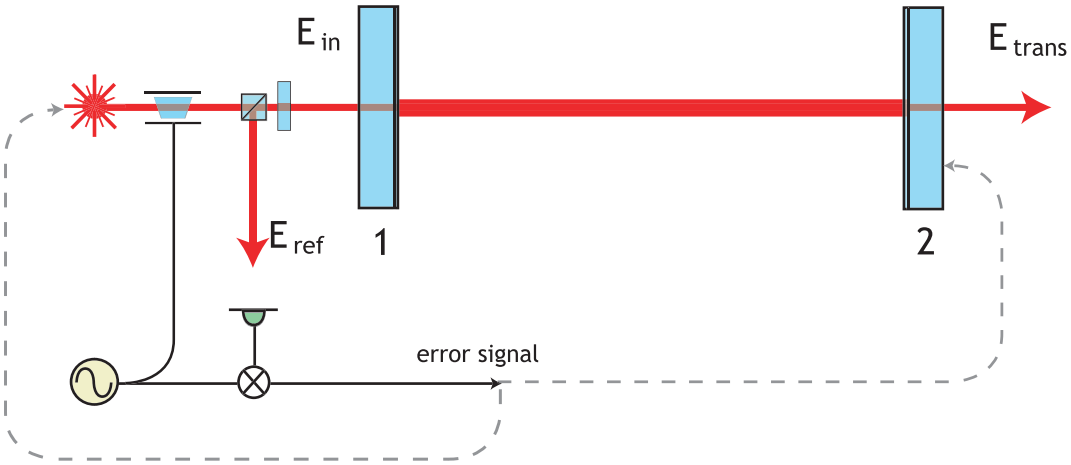

Figure 2.7: Schematic of a Fabry-Perot cavity with optical phase modulator, local oscillator source, photodiode, demodulating mixer and possible feedback paths.

Measurement Limits

The quantum nature of light fundamentally limits the resolution to which laser frequency noise or cavity displacement noise can be measured.

The light used to make these measurements is the reflected light from a Fabry-Perot cavity. For a simplified, symmetric cavity ($\rho_1 = \rho_2 = \rho \simeq 1$) with high reflectivity mirrors ($\tau_1 \simeq 0$) *Equation 2.42* can be simplified to,

$$E_{\rm ref} = E_{\rm in}\left(\frac{\rho(e^{i\phi}-1)}{1-\rho^2 e^{i\phi}}\right) \quad = \quad E_{\rm in}\Gamma(\phi)\ , \tag{2.49}$$

where $\phi = 2kL$, and $\Gamma(\phi)$ is then a *complex reflection coefficient*.

The reflected light power from the cavity (including phase modulation) can be given by the expression,

$$P_{\rm ref} = E_{\rm mod}E^*_{\rm mod} = \underbrace{|C|^2 + |S_+|^2 + |S_-|^2}_{DC\ terms} + (CS^*_+ + C^*S_+) + (CS^*_- + C^*S_-) + \underbrace{\ldots}_{2\phi\ terms} \tag{2.50}$$

where the terms representing the carrier, upper sideband and lower sideband from *Equation 2.47* have been rewritten as,

$$C = J_0(m)E_0e^{i\omega t} = ae^{i\phi}\Gamma(\phi)\ , \tag{2.51}$$

$$S_+ = J_1(m)E_0e^{i\omega t}e^{i\omega_{\rm m}t} = be^{i\Phi}\Gamma(\Phi)\ , \tag{2.52}$$

and

$$S_- = J_1(m)E_0e^{i\omega t}e^{-i\omega_{\rm m}t} = be^{-i\Phi}\Gamma(-\Phi)\ , \tag{2.53}$$

for simplicity. Also, $\pm\Phi = \phi \pm \phi_{\rm m}$ where $\phi_{\rm m} = \omega_{\rm m}t$ represents the differing phase evolution of the sidebands with respect to the carrier.

The demodulation performed by the mixer provides an error signal proportional to the modulation frequency terms. So, discarding the DC and $2\omega_{\rm m}$ terms gives the error signal ϵ as,

$$\begin{aligned}\epsilon &= 2ab\ \{\Re\left[\Gamma(\phi)\Gamma^*(\Phi) - \Gamma^*(\phi)\Gamma(-\Phi)\right]\cos(\phi_{\rm m}) \\ &\quad + \Im[\Gamma(\phi)\Gamma^*(\Phi) - \Gamma^*(\phi)\Gamma(-\Phi)]\sin(\phi_{\rm m})\}\end{aligned} \tag{2.54}$$

$$\Rightarrow d\epsilon = 4\sqrt{P_{\rm c}P_{\rm s}}\ \Im[\Gamma(d\phi)]\ , \tag{2.55}$$

where $d\epsilon$ is the error signal around resonance (taking the terms proportional to $\sin(\phi_{\rm m})$ for an optimally phased local oscillator), P_c is the power of the carrier light and P_s is the power in one sideband. $\Gamma(\Phi) = -1$ as the sidebands should be entirely reflected, and $\Gamma^*(d\phi) = -\Gamma(d\phi)$ around the resonance, where $\phi = d\phi$.

Near the resonance, the error signal from the mixer is linear, and the change in phase detected is,

$$d\phi = 2k\ dL = \frac{4\pi}{\lambda}dL\ , \tag{2.56}$$

implying that the carrier reflection coefficient around resonance $\Gamma(d\phi)$ can be written as,

$$\Gamma(dL) \simeq \left(\frac{\rho}{1-\rho^2}\right)\left(\frac{i4\pi}{\lambda}dL\right) \tag{2.57}$$

$$\simeq i\frac{4\mathcal{F}}{\lambda}dL\ . \tag{2.58}$$

Substituting this into *Equation 2.55* gives,

$$d\epsilon = 16\sqrt{P_{\rm c}P_{\rm s}}\mathcal{F}\frac{dL}{\lambda}\ , \tag{2.59}$$

and we now define the term μ as,

$$\mu = \frac{d\epsilon}{dL}\ . \tag{2.60}$$

The measurement quantity involved is typically the displacement amplitude spectral

density of a cavity, (or the frequency amplitude spectral density of the light source) which is related to μ by,

$$\xi_x = \frac{\xi_\epsilon}{\mu} \tag{2.61}$$

where ξ_x is the displacement amplitude spectral density in $\mathrm{m}/\sqrt{\mathrm{Hz}}$, ξ_ϵ is the amplitude spectral density of the error signal in $\mathrm{V}/\sqrt{\mathrm{Hz}}$ and μ is then in Vm^{-1} if the detector used has a linear voltage response to incident light power.

The shot noise in the error signal ϵ has the following form [5] [54] within the cavity bandwidth,

$$\xi_\epsilon = \sqrt{\frac{4hcP_\mathrm{s}}{\lambda}}\ , \tag{2.62}$$

which then leads to the following expression for the shot noise limited displacement amplitude spectral density of a cavity:

$$\xi_x = \frac{1}{8\mathcal{F}}\sqrt{\frac{hc\lambda}{P_\mathrm{c}}}\ . \tag{2.63}$$

Equivalently, this can be expressed as the shot noise limited frequency measurement on the laser,

$$\xi_f \quad = \quad \frac{f}{L}\xi_x \quad = \quad \frac{1}{8\mathcal{F}L}\sqrt{\frac{hc^3}{\lambda P_\mathrm{c}}}\ . \tag{2.64}$$

Equation 2.63 and *Equation 2.64* represent the limits to which the cavity displacement noise and laser frequency noise can be measured, in absence of greater noise sources.

2.3.2 Feedback & Control Systems

It has been shown that an optical cavity (or indeed, a system of optical cavities) can be used to measure the frequency fluctuations of a laser, and that a complementary setup can be used to measure the displacement noise of a cavity. However, for such

measurements to be performed practically, the system must be held on resonance. This is achieved by first attaining an appropriate error signal (as outlined in *Section 2.3.1*). This error signal may then require filtering, before being applied to either the laser or the cavity mirror in question, forming a *negative feedback control loop*.

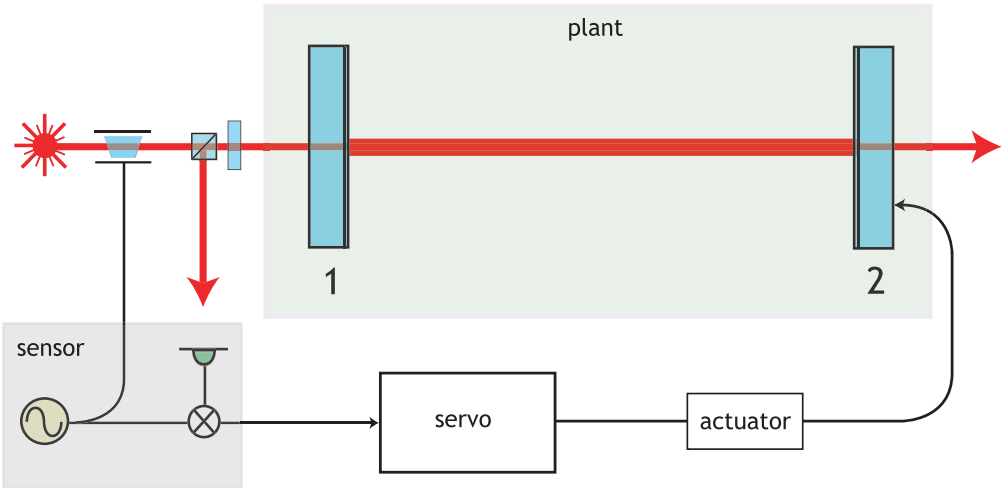

Figure 2.8: Schematic of a control loop applied to the length feedback of a Fabry-Perot cavity. The optical system is described as the *plant*, with the photodiode arrangement then the *sensor*. A servo and actuator then process the error point signal and provide appropriate feedback to hold the cavity on resonance. Equally, feedback could be applied to the laser.

The system is then characterised by a complex transfer function, which is the product of the individual transfer functions of the different subsystems,

$$G_{\text{tot}} = G_{\text{servo}}[G_{\text{sensor}}G_{\text{actuator}}G_{\text{plant}}] = GH \ , \tag{2.65}$$

where $G_{\text{servo}} = H$ is the complex transfer function of the filter stage (typically known), and G is the combined complex transfer function of the other subsystems.

The quantity G_{tot} is termed the *open-loop gain* of the system, and as such is the complex gain of all the appropriate feedback elements, when no feedback is being actively applied.

When the loop is then *closed* and negative feedback is applied, the sensor supplies a control signal, which is then filtered by the servo and sent to the actuator, which

applies the feedback. The feedback signal supplied by the actuator then alters the plant, bringing it towards equilibrium (i.e the operating point). If noise (ξ_{n}) is present (either specifically injected or fundamental) at the point of actuation ξ_{a}, and the output of the servo system is ξ_{servo} then,

$$\xi_{\mathrm{a}} = \xi_{\mathrm{servo}} + \xi_{\mathrm{n}} \ , \tag{2.66}$$

$$\xi_{\mathrm{servo}} = GH\xi_{\mathrm{a}} \ . \tag{2.67}$$

Rearranging, this gives;

$$\xi_{\mathrm{a}} = \frac{\xi_{\mathrm{n}}}{1 - GH} \ , \tag{2.68}$$

$$\xi_{\mathrm{servo}} = \left(\frac{GH}{1 - GH}\right) \xi_{\mathrm{n}} \ . \tag{2.69}$$

So, the noise in term ξ_{a} is suppressed by the factor $\frac{1}{1-GH}$, where GH is the open-loop gain, and suppression occurs when $|GH| > 1$. The factor $\frac{1}{1-GH}$ is also termed the *closed-loop gain*. Note that as the servo acts to suppress changes in the plant, $GH < 0$ will typically be true within the desired control bandwidth.

While there is no theoretical limit to the amount of noise suppression shown here, in reality any measurement will be limited by e.g. shot noise effects, as outlined previously in *Section 2.3.1*.

Stability

For a control loop to be stable, the phase of the complex open-loop gain must be within $\pm 180°$ (where the phase at DC is taken as $0°$) when a unity gain frequency is reached – i.e. when the amplitude of the open-loop gain is equal to one. Failing to meet this condition results in positive feedback. The difference between the phase of the open-loop gain at the unity gain frequency, and the region of instability is termed the *phase margin*. Practically it is deemed prudent to build systems with a phase margin

of at least 40°, to take account of potential unknown phase delays in the subsystems. Similarly, the term *gain margin* is defined as the difference between the maximum and minimum safe possible gain – i.e. those that do not move the unity gain frequency into a region of instability. A large gain margin allows adjustments in the gain to be made safely without causing the system to lose stability. Additionally, if multiple feedback paths (where actuators only act over a certain frequency range for instance) are used then at the *crossover points*, where one path becomes dominant over the other, the phase difference between the paths must be less than 180° and the frequency response of the paths must differ by ≤ 2 powers of frequency.

Transfer Functions

Characterising the open and closed loop gains of a system (where not all the transfer functions are necessarily known) can be achieved by injecting known signals into one point of the loop.

If a signal ξ_{inj} is injected into the system directly before the servo filtering stage then,

$$\xi_0 = \xi_{\text{ep}} + \xi_{\text{inj}} \ , \tag{2.70}$$

$$\xi_{\text{ep}} = \xi_0 GH \ , \tag{2.71}$$

where ξ_0 is the input to the servo stage, and ξ_{ep} is the output of the plant. Rearranging gives,

$$\frac{\xi_0}{\xi_{\text{inj}}} = \frac{1}{1-GH} \ , \tag{2.72}$$

$$\frac{\xi_{\text{ep}}}{\xi_{\text{inj}}} = \frac{GH}{1-GH} \ , \tag{2.73}$$

which are the open-loop and closed-loop transfer functions respectively. Injecting a signal into a closed loop system, and monitoring the system's response, yields the appropriate transfer functions.

Chapter 3

A Direct Thermal Noise Measurement Experiment: *Design*

3.1 Introduction & Motivation

Thermal noise effects are the dominant limit to sensitivity in a gravitational wave interferometer's most sensitive frequency range. It is important to be able to characterise this phenomenon, so that different mirror coatings, substrates and test-mass compositions can be evaluated reliably and accurately.

Typically, the thermal noise of an optic or substrate is estimated by measuring the object's mechanical quality factor or Q, and extrapolating the thermal noise contribution in the gravitational wave detection band using the *Fluctuation-Dissipation Theorem* as follows [55] :

The object in question is suspended by e.g. a single loop of silk thread in a vacuum system, where its resonant modes are excited by an electrostatic driving plate. A Michelson interferometer is used to sense the mechanical ringdown of the object's resonant modes. The optic acts as the end mirror of one interferometer arm, with the other arm being formed by two mirrors – one mounted on a loudspeaker, the other on a piezo-electric transducer (PZT). Feedback to the loudspeaker (to compensate for slow drifts of the optic position) and the PZT (above the audio range of the speaker) allows the system to stay locked and facilitates a measurement of the mechanical ringdown. The exponential envelope of the ringdown can then be related to the Q, and the associated loss angle at the resonant frequency ϕ_{ω_0} (where $Q_{\omega_0} = \frac{1}{\phi_{\omega_0}}$) through the following expressions:

$$A(t) = A_0 e^{-\frac{\omega_0 t}{2Q}} = A_0 e^{-\frac{(\omega_0 t)\phi_{\omega_0}}{2}} \ , \tag{3.1}$$

$$\ln[A(t)] = -\frac{(\omega_0 t)\phi_{\omega_0}}{2} + \ln[A_0] \ . \tag{3.2}$$

Here, the excited resonance frequency is ω_0, and the initial amplitude of the ringdown signal is A_0. The loss angle can be extracted by plotting the natural log of the decaying amplitude $A(t)$ against time.

Typically, mechanical Qs in the region of 10^8 will be measured for resonant modes in the $50 \rightarrow 100\,\mathrm{kHz}$ range, for a fused silica mirror substrate.

This figure must then be translated into a relevant level of thermal noise in terms of displacement noise spectral density. From the *FDT*, the thermally induced displacement noise power spectral density from such an harmonic oscillator is given by,

$$\Xi(\omega) = \left(\frac{4k_\mathrm{b}T}{\omega}\right) \frac{\phi(\omega)\omega^2}{m[{\omega_0}^4\phi^2(\omega) + ({\omega_0}^2 - \omega^2)^2]} \ . \tag{3.3}$$

Knowledge of the Q, and some assumptions about the frequency dependence of the loss angle $\phi(\omega)$ then allow the thermal noise in the gravitational wave band to be estimated. For example, a mode at 60 kHz exhibiting a Q of 10^8 implies that the

thermal displacement noise at 500 Hz will be $\sim 6\times10^{-23}$ m/$\sqrt{\text{Hz}}$, if $\phi(\omega)$ is frequency independent.

While useful, this method has limitations: it infers the level of thermal noise as opposed to sensing it explicitly, and it is unavoidably sensitive to other forms of loss (in the suspension or clamping structure) adversely affecting the Q. It can therefore be difficult to separate the intrinsic properties from those associated with the measurement process.

As future generation detectors require increasingly stringent limits on noise sources, the importance of characterising the thermal properties of critical optics becomes paramount. New materials, coatings, dopants, and bonding techniques may affect the thermal noise contribution in unknown ways that may not be readily quantifiable from a typical Q measurement.

An interferometric measurement senses the thermal noise contribution directly, in the measurement band of interest – where it will limit advanced gravitational wave interferometers – and should eliminate the inherent uncertainties of mechanical resonance measurements.

3.2 Experimental Outline

3.2.1 Introduction

To directly measure the thermal noise contribution, one can imagine constructing an experiment where a single thermally limited Fabry-Perot cavity is locked to a monochromatic light source. The Fabry-Perot cavity effectively amplifies the phase-front alterations caused by the fluctuating mirror surfaces, and the feedback signal required to lock the system is then related to the displacement noise of the optics. While feasible, this concept needs considerable refinement before it can ultimately be employed.

An *over-coupled* cavity has the input mirror more transmissive than the sum of all other losses in the cavity[1], and is desirable for use in phase or displacement mea-

[1]For a suitably low loss system it can equivalently be said that an over-coupled cavity has the output optic more reflective than the input coupler.

surements as the light reflected from such a system exhibits a larger variation of phase shift around a resonance than the *under-coupled* case [56]. Additionally, any extra losses accrued on the optic surfaces with time will have a lesser adverse effect on the signal in an over-coupled system. As thermal noise due to the coating increases with the number of dielectric layers (and therefore also with reflectivity) it is the end, more reflective, optic that will exhibit the larger intrinsic thermal fluctuation.

Recall from *Section 2.2.2* that a small measurement beam leads to a correspondingly larger thermal noise signal. As such it may seem prudent to arrange a small beam waist on the end mirror, and a larger spot on the input optic – see *Figure 3.1*. This would cause the thermal noise signal from the end optic to dominate the measurement further.

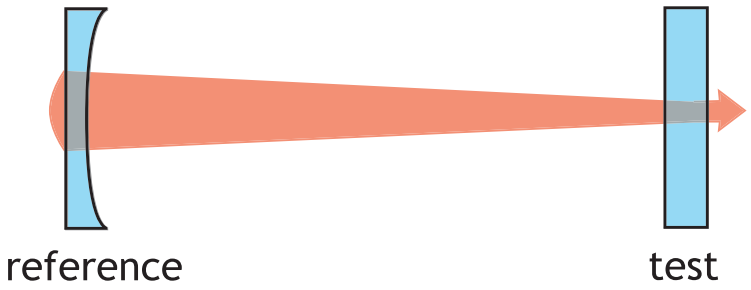

Figure 3.1: A naive schematic of the layout of a thermal noise measurement cavity: the final optic is more reflective than the input coupler, is probed at a beam waist and used as the dominating test optic of the system.

However, this requires that the end optic should always be more reflective than the input optic – to maintain the over-coupled state – which imposes restrictions on the optic substrate and coating when performing new measurements. Indeed, it requires that the test optic should always have more dielectric layers (and thus a higher level of intrinsic thermal noise) than the reference optic – so we can *never* effectively test a configuration where the test optic has relatively low levels of thermal fluctuation. The only appropriate topology is to define the input optic of the cavity to be flat, with the end mirror concave and of higher reflectivity – see *Figure 3.2*. Although the end mirror has more coating layers and therefore higher intrinsic thermal noise, the input optic may still dominate the measurement through having a suitably small beam waist on its surface, relative to that on the end mirror.

Figure 3.2: A schematic layout of a thermal noise measurement cavity: the input coupler is flat, with its reflecting surface coinciding with a beam waist. The end mirror is concave and more reflective than the input optic. The thermal noise contribution from the input mirror dominates the measurement.

To successfully sense the thermal variations of the optic surfaces, the laser light used to probe the surface must be sufficiently frequency stable; otherwise the frequency variations in the measurement beam will dominate the thermal noise signal from the test cavity. The frequency noise of a typical free-running Nd:YAG laser is $\frac{\sim 10^4}{f}$ Hz/$\sqrt{\text{Hz}}$ [57], while the achievable level of sensitivity can be expressed as,

$$\frac{\Delta f}{f} = \frac{\Delta L}{L}, \tag{3.4}$$

where Δf is the frequency noise of a laser of frequency f in Hz/$\sqrt{\text{Hz}}$, and ΔL is the displacement noise of a cavity with length L in m/$\sqrt{\text{Hz}}$. Using an unmodified free-running laser, the sensitivity to displacement noise in a 1 m optical cavity is $\sim 3{\times}10^{-14}$ m/$\sqrt{\text{Hz}}$ at 1 kHz. A thermally induced displacement noise of this level is, in principle at least, detectable – if an appropriately sized beam is used to make the measurement. An order of magnitude estimate can be obtained using *Equation 2.8* to estimate the required beam radius to sense the Brownian thermal noise of a fused silica substrate[2]. This implies that a suitable beam radius is $r_0 \simeq 1{\times}10^{-15}$ m. Clearly, this is nonsensical and another solution must be found. The level of laser stabilisation must be improved.

As outlined in *Section 2.3.1* it is possible to both lock an optical cavity to the frequency of a light source, and lock the frequency of a laser to the length of an optical

[2]Material properties can be found in *Appendix A*

cavity.

If the laser is stabilised to a cavity of length L_1 and the test cavity has a length L_2 then the measurable displacement noise of the test cavity is (recalling *Equation 2.64* – the shot noise limited frequency noise of a laser locked to a cavity),

$$\Delta L_2 = L_2 \left(\frac{\Delta f}{f} \right) = \left(\frac{L_2}{L_1} \right) \frac{1}{8\mathcal{F}_1 f} \sqrt{\frac{hc^3}{\lambda P}}, \tag{3.5}$$

where ΔL_2 is the shot noise limited displacement noise of the test cavity, Δf is the frequency noise of a laser of frequency f, L_1 and L_2 are the respective cavity lengths, $\mathcal{F}_1$ is the finesse of the cavity, λ is the laser wavelength and P is the light power of the carrier light in the laser stabilisation cavity.

By minimising the ratio $\frac{L_2}{L_1}$ (through having the measurement cavity considerably shorter than the cavity used for the laser stabilisation) it is possible to maximise the measurement's sensitivity to ΔL_2.

In summary: to achieve the goal of direct thermal noise measurement we need to construct a two cavity system where the term $\Delta L_2 < \xi_{\text{thermal}}$; i.e. where the displacement noise that we *can* measure is less than the thermally induced displacement noise we expect, and where all other noise sources are also minimised or suppressed. This can be effected through the use of a short high finesse test cavity – with respect to a longer laser stabilisation cavity – and a small beam waist on the designated test optic, to ensure that it dominates the measurement.

3.2.2 Experimental Layout

The JIF Facility

To this end, an experiment was constructed in the IGR JIF facility, which houses a 10 m gravitational wave interferometer prototype in a Class 1000 clean area. The facility was designed to test systems and optical configurations for advanced full-scale gravitational wave detectors and houses several main optics suspensions similar to those found in

GEO600.

The lab consists of nine vacuum tanks, forming a folded L shape interferometer – see *Figure 3.3*. One arm houses a Fabry-Perot cavity of length 9.78 m, made up of two mirrors – an inboard test mass (*ITM*) and an end test mass (*ETM*). Two additional steering mirrors are used to direct the beam into the cavity. Each tank contains a support structure formed from "Bosch" aluminium extrusions, large enough to support two triple pendulum systems culminating in 2.7 kg fused silica optics, with an upper stage supported by rubber damping elements. The input laser is a single-mode, single-frequency, continuous-wave *Innolight Mephisto* Nd:YAG non-planar ring oscillator (NPRO), operating at 2 W.

The *ITM* and *ETM* masses are ultra-low loss silica optics with multi-layered dielectric reflective coatings, suspended as the final stages in a triple pendulum, and can be acted on via *reaction pendulums* suspended behind them. The lowest stage of the reaction pendulum is an aluminium mass, with identical dimensions to the silica optics. This mass has attached coils, wired to feedthroughs in the surrounding tank. The main optics have small magnets bonded to their non-reflective sides; inducing a current in the reaction mass coils allows sufficiently delicate control of the position of the main optics.

The low frequency pendulum modes, and alignment of the test-mass suspensions, are controlled via a combined system of coils, magnets, infra-red LEDs and shadow sensing techniques, governed by a LabVIEW$^{\text{TM}}$ control system. The uppermost stage of the triple pendulum has combined magnet/flag attachments on its extremities, which pass through LED/coil/sensor elements attached to the frame superstructure. Changes in the light levels for each degree of freedom are then interpreted as pendulum motion, and a corrective signal sent through the coils, effectively damping the motion.

The optic labelled *PRM* is (nominally) the power recycling mirror for the system, and is a standard 1" diameter optic housed in an aluminium mass, suspended as a double pendulum. For these measurements it was replaced with a blank, anti-reflection coated optic, with the same radius of curvature as the reflective power recycling mirror

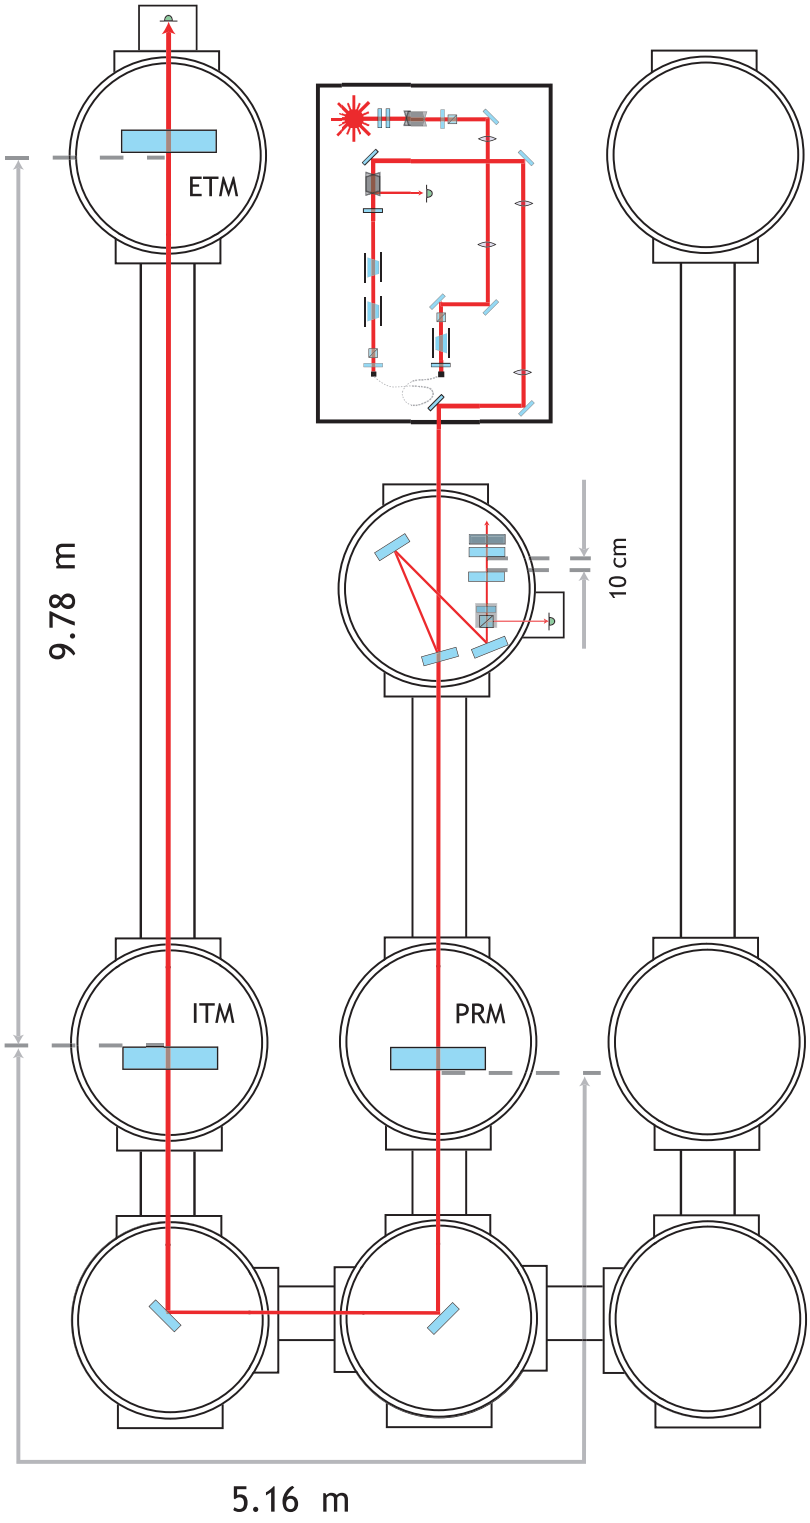

Figure 3.3: JIF lab layout: the system consists of nine interconnected vacuum tanks, with one main arm cavity, an optional power recycling cavity, and a short test cavity for thermal noise measurements. The laser bench holds the laser, phase modulators, single mode optical fibre, Faraday isolators, detection photodiode and mode-matching lenses, as well as other associated optics.

($R_c = -15$ m), to keep the mode-matching condition into the main arm cavity valid for different experimental configurations.

The optical cavity in the main arm (hereafter termed the *arm cavity* or simply AC) serves as the frequency reference for the laser. A short, 10 cm thermal noise measurement cavity is housed in the first input tank.

The secondary arm in the system was not used in any of these experiments.

The thermal noise experiment (TNE) tank contains seven suspensions: two fused silica suspensions comprising the test cavity, four small optic suspensions – required to steer, mode-match and separate the outgoing light – and one reaction suspension, to provide feedback signals to the test cavity. The test optics used in the thermal noise measurement cavity are smaller in size (3" diameter) and mass (~ 360 g) than the main optics used in the AC. In more detail, they are:

- The inboard test mass suspension, or ITM_{TNE}. Its final stage is the thermal noise test optic. This mirror is suspended from the intermediate stage by silica fibres, and the upper stage by a single wire-loop from cantilever springs.
- The end test mass suspension, or ETM_{TNE} comprises the silica end mirror of the test cavity, and a silica intermediate stage. Aside from the optical properties of the final stage mirror, this suspension is identical to ITM_{TNE}.
- RM is the reaction mass suspension. This is an all metal construction whose final stage is an aluminium mass (of the same dimensions as the silica test optics), with four coils attached to the front side. These are then used to actuate on four corresponding magnets bonded to the rear face of the ETM_{TNE}. It is also suspended from cantilever blades to aid vertical isolation.
- Immediately preceding the ITM_{TNE} is the *mounting unit*, or MU. It is a com-

bined polarising beam-splitter and quarter wave-plate setup. This effectively splits the light returning from the cavity from the ingoing light, and ejects it through a viewport in the tank, to be detected by the photodiode on a small optical bench.

- A flat steering mirror ($S1$) is used to direct light into the cavity.
- A suspension with a mode-matching (and steering) optic – MM. This is used to both correctly set the mode-matching condition into the system, and steer the light. The radius of curvature of the mode-matching optic is $R_{\mathrm{c}} = 0.15$ m.
- The beam-splitter suspension (BS) is used to pick light off from the main system and inject it into the thermal noise measurement cavity. It has a power reflectivity of 33%.

A rendered mock-up of the thermal noise measurement tank can be seen in *Figure 3.4*, and a schematic showing the relative positions of the optics in *Figure 4.19*.

The optical properties of the critical optics from both cavities are summarised in the table below.

$\mathbf{PRM}_{\mathrm{blank}}$	**ITM**	**ETM**	$\mathbf{ITM}_{\mathrm{TNE}}$	$\mathbf{ETM}_{\mathrm{TNE}}$
$\rho_{\mathrm{PRM}} = 0$	$\rho_{\mathrm{ITM}} = 0.99497$	$\rho_{\mathrm{ETM}} = 0.99998$	$T = 600_{\mathrm{ppm}}$	$T = 10_{\mathrm{ppm}}$
$R_{\mathrm{c}} = -15\,\mathrm{m}$	$R_{\mathrm{c}} = \infty$	$R_{\mathrm{c}} = 15\,\mathrm{m}$	$R_{\mathrm{c}} = \infty$	$R_{\mathrm{c}} = 0.15\,\mathrm{m}$

Table 3.1: Mirror properties: the values for the arm cavity optics are given as amplitude reflectivity coefficients, and as transmitted power (in *parts per million*) for the thermal noise measurement optics. Note that the transmission values for the thermal measurement optics are design specifications, and were not measured. Also, all transmissive optics were AR coated on their non-reflective faces.

The theoretical finesse values for the respective cavities can then be calculated (using *Equation 2.43*) as $\mathcal{F}_{AC} = 621$ and $\mathcal{F}_{TNE} = 10^4$.

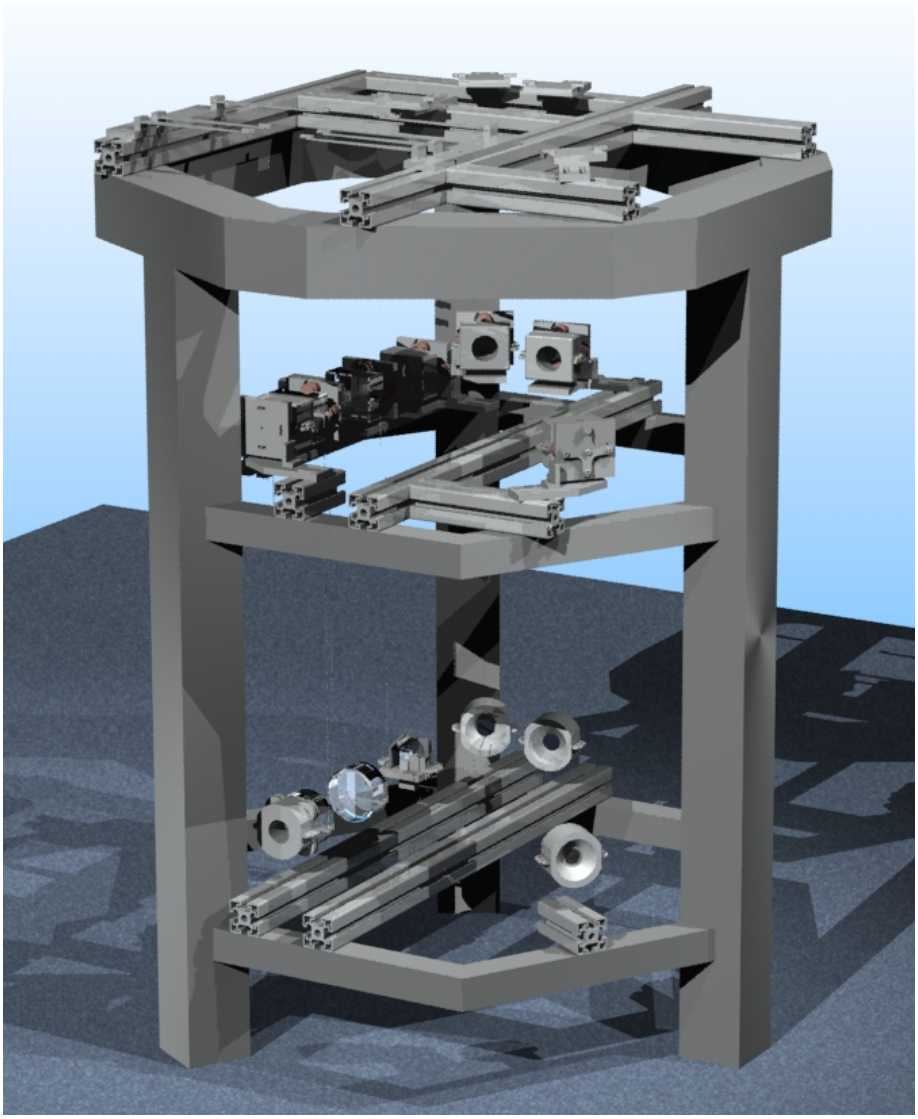

Figure 3.4: SolidWorks® mock-up of the thermal noise measurement tank [58].

Measurement Scheme

Using phase modulation sidebands at 46.526 MHz, the laser is stabilised to the length of the *AC* using a *Pound-Drever-Hall* method as detailed in *Section 2.3.1*, via feedback to a (*PZT*) bonded to the laser crystal, and (via a high-voltage amplifier) an electro-optic modulator in the beam path. The servo for the laser stabilisation feedback is designed to provide stabilisation over a large range, from DC to $\sim$ 40 kHz. Feedback to the position of the *ETM* is used at low frequencies ($\leq$ 1 Hz). The thermal noise test cavity is then locked via feedback to the ETM_{TNE} position using a separate 10 MHz locking scheme – see *Figure 3.5* – with the servo having a unity gain frequency of $\sim$ 150 Hz. The error-point of the feedback above the unity gain frequency can then be calibrated to give the desired displacement noise measurement.

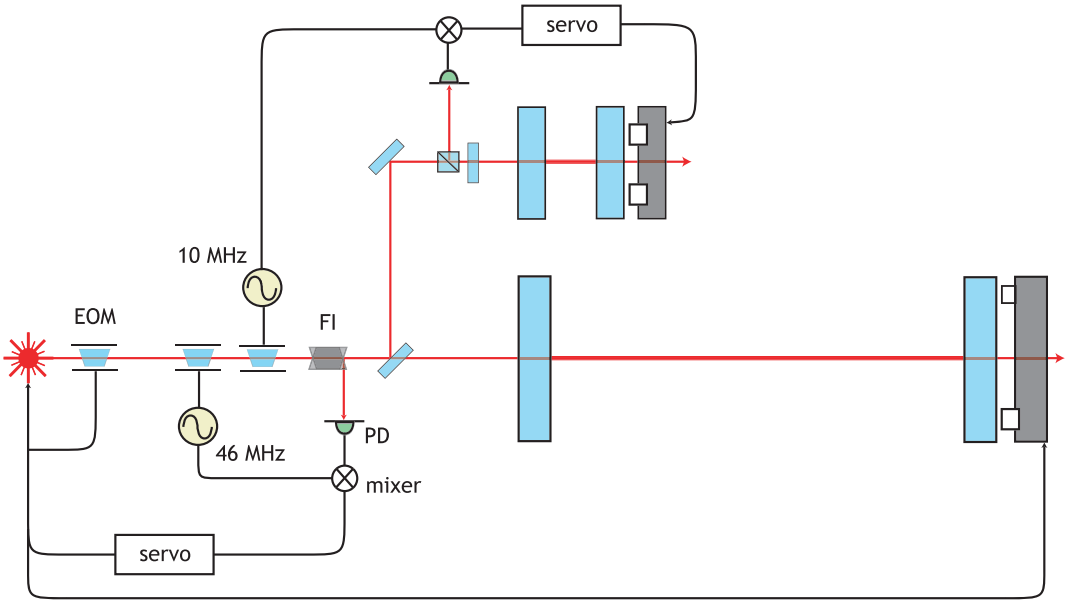

Figure 3.5: Control scheme layout for the direct thermal noise measurement experiment: Phase modulation sidebands at 10 MHz and 46.526 MHz are imparted on the light by two electro-optic modulators. The control signal derived from the 46.526 MHz signal is filtered through a servo and used to feed back to the frequency of the laser via a *PZT* and electro-optic modulator – and additionally to feed back to the cavity length at low frequencies. Similarly, the 10 MHz sideband signal is used to control the length of the measurement cavity.

3.2.3 Sensitivity & Design

As detailed in *Section 2*, there are myriad noise sources that must be quantified and structured such that the thermally induced displacement noise of the test optic is the dominant noise source in the measurement.

Shot Noise Limit

The shot noise limited frequency noise of the signal readout required to lock a laser to a cavity "x" (or equivalently to lock a cavity length to a laser) can be expressed as [57],

$$\Delta f_x(f) = \Lambda_x \sqrt{\left(\frac{hc(1-J_0^2(m_x)M_xV_x)}{16\lambda\eta_x P_x}\right)\left[1+\left(\frac{2f}{\Lambda_x}\right)^2\right]} \times \left(\frac{1}{M_1 J_0(m_x)J_1(m_x)}\right)\left(\frac{1}{1-\sqrt{1-V_x}}\right) \text{ Hz}/\sqrt{\text{Hz}}\,, \tag{3.6}$$

where Λ_x is the linewidth of the cavity, $J_0(m_x)$ is the amplitude of the main carrier light, $J_1(m_x)$ is the amplitude of a phase modulation sideband, m_x is the modulation index of the sidebands, M_x is the fraction of input power P_x matched into the cavity, η_x is the quantum efficiency of the photodetector, and V_x is the cavity *visibility* defined as $V = \frac{P_{\max}-P_{\min}}{P_{\max}}$, where $P_{\min}$ and $P_{\max}$ are the minimum and maximum intensities measured when the system passes through a resonance – where the cavity here is assumed to be over-coupled.

As there are two distinct measurement points in this system – one detection point controls the laser frequency, and another the length of the short cavity – there will be two uncorrelated contributions: Δf_1 and Δf_2.

Recall from *Equation 3.4* the relation,

$$\frac{\Delta f}{f} = \frac{\Delta L_1}{L_1} = \frac{\Delta L_2}{L_2}\,. \tag{3.7}$$

By rearranging, it can be seen that,

$$\Delta L_2 \quad = \quad \frac{L_2}{f}\Delta f \quad = \quad \frac{L_2}{f}\sqrt{{\Delta f_1}^2 + {\Delta f_2}^2}\ . \tag{3.8}$$

For the measurement to succeed we require the shot noise limited frequency noise contribution from the laser to be below the equivalent contribution from the thermal noise fluctuations of the test optic, i.e. $\Delta f_1 < \Delta f_2$. Neglecting sideband size and photodetector efficiency – and recalling that light is injected into the *TNE* via a 33% beam-splitter – this condition will be met by a factor of [59],

$$\frac{\mathcal{F}_1 L_1}{\mathcal{F}_2 L_2}\sqrt{\frac{P_1}{P_2}} \quad = \quad \sqrt{2}\left(\frac{\mathcal{F}_1 L_1}{\mathcal{F}_2 L_2}\right) \quad \simeq \quad 10\ . \tag{3.9}$$

Radiation Pressure & Intensity Noise

Fluctuations in laser power can give rise to fluctuations in test-mass position, through radiation pressure effects. As the test masses in the *TNE* are relatively light, this should be evaluated. The contribution from this noise source can be estimated as [5],

$$\xi_{\mathrm{rad}}(f) = \frac{\mathcal{F}}{c\pi^3 f^2 m}\Delta P\ , \tag{3.10}$$

where

$$\Delta P = \sqrt{\frac{Phc}{\lambda}}\ , \tag{3.11}$$

for a shot noise limited case, and

$$\Delta P = P \times RIN \tag{3.12}$$

for the case where the power fluctuations are dominated by the intensity noise of the laser. Here RIN is the relative intensity noise in units of $\mathrm{Hz}^{-\frac{1}{2}}$. Assuming an

input power of 333 mW, a finesse of 10^4, a mirror mass of 360 g and an intensity noise dominated RIN level of $3 \times 10^{-8}/\sqrt{\text{Hz}}$ [60], this yields a level of 3×10^{-20} m$/\sqrt{\text{Hz}}$ at 1 kHz for the thermal noise test cavity. Given that the optics in the AC are more massive, the finesse lower and the cavity longer – the final noise contribution from this effect in that cavity will be attenuated by a factor of $\sim 10^4$.

Seismic Noise

To attenuate seismic noise (see *Section 2.2.1*), all the *TNE* optics are suspended as double stage pendulums, with the optics comprising the test cavity being suspended as quasi-monolithic assemblies, with fused silica fibres forming the lower suspension elements. The length of the upper stage of the suspensions is 38.6 cm and the lower stage 56.2 cm, pushing the main pendulum mode resonance down to ≤ 700 mHz.

The relatively long suspension stages (and correspondingly lower pendulum resonances) provide greater horizontal isolation at frequencies above the pendulum resonance than typical, shorter suspensions, thus widening the range where thermal effects from the optical substrates and coatings will limit the system. See *Figure 4.1*, *Figure 4.10*, and *Figure 4.15* for images of the *TNE* suspensions.

All of the suspensions in the *TNE* tank are aligned through the use of 10 mm diameter magnets bonded to the rear of the upper mass. These magnets interact with coils wound round copper formers, attached to the frame structure in the tank, allowing alignment forces to be exacted through passing a current through these coils. The main pendulum resonance is damped via the eddy-current damping effect that arises from the magnets from the alignment system interacting with the copper formers of the alignment coils. This technique has been shown [61] to reduce resonant pendulum mode quality factors to $Q \simeq 5$. See *Figure 4.16* for a picture of the eddy-current damping and alignment systems.

Vertical seismic motion can couple into horizontal mirror motion, due to both pendulum imperfections and – for full-scale detectors – the curvature of the Earth. The magnitude of this coupling effect can be considered to be $\sim 0.1\%$ [62]. To ameliorate

this effect, the test cavity (and reaction) suspensions are suspended by cantilever blades attached to the frame. These provide a soft vertical spring (in comparison to the vertical resonance from the suspension wires) with a low resonance frequency. The topmost section of the frame (from which all the optics are suspended), sits on damped rubber spring elements (formed from GE RTV615 silicone rubber with added graphite) which help isolate the system [63] further – both horizontally and vertically. Additionally, as all the optics in the *TNE* tank are suspended from effectively the same point, any differential vertical motion should be minimised.

The seismic noise contribution to the measurement from the *AC* optics is considerably smaller compared to that from the *TNE* optics. This is due to two factors: firstly, the motion of the mirrors contributes less to the readout signal by the factor $\frac{L_1}{L_2}$; secondly, the *ITM* and *ETM* optics are suspended more rigorously. As they are designed to be very similar to GEO600 main suspensions, they have three pendulum stages for horizontal isolation and two stages of cantilever blades to aid vertical isolation. Although the main pendulum resonances occur at slightly higher frequencies (due to shorter wire lengths), the horizontal isolation increases as $\sim f^{-6}$ above a few Hz, due to the triple pendulum. As such the seismic noise contribution from the *TNE* test cavity is the dominating seismic effect.

Suspension Thermal Noise

As outlined in *Section 2.2.2*, there can be a significant thermal noise contribution from the suspension elements in any suspended optical system, from both the main pendulum mode, and transverse violin modes. Silica fibres are used to suspend the optics in the thermal noise test cavity, as the loss of standard steel wire loops – or wire and clamp arrangements – may be large enough to overcome the thermal noise contribution from the optics themselves. The low loss of the silica suspension elements effectively channels the thermal energy into the main pendulum and transverse violin modes, reducing suspension induced thermal noise contributions in the measurement band.

The loss angle of the pendulum defines its thermal noise contribution, and is proportional to $\frac{r^2}{l}$ (see *Equation 2.25*) where r is the radius of a cylindrical suspension element, and l is the length. To reduce the effect of pendulum thermal noise, long, thin silica fibres must be used. Fibres of length 57 cm and diameter 100 μm were chosen, as these provide sufficient pendulum thermal noise attenuation without imposing undue practical limitations.

However, constructing a low loss, high Q, system also implies that the transverse violin modes (see *Figure 3.6*) will have similarly high Q values[3].

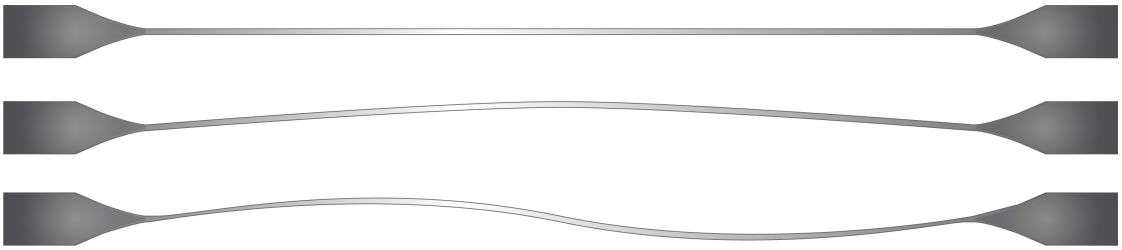

Figure 3.6: Violin modes: schematic representation of a fibre with no excitation, then first and second order violin excitation respectively. Note that the motion depicted here is highly exaggerated.

From *Equation 2.26* it can be seen that these modes will form a slightly anharmonic series beginning at $\sim$ 200 Hz for the *TNE* parameters. Their thermal noise contribution can be calculated from *Equation 2.27*.

Coating & Substrate Thermal Noise

Using the formulae from *Section 2.2.2*, it is possible to estimate the level of thermally induced displacement noise expected from the test cavity. Assuming no optical losses in the mirrors, that they meet their design specification and that they consist of alternating $\frac{\lambda}{4}$ layers of SiO_2/ Ta_2O_5; and estimating a beam radius of 155μ at the front reflective face of the test optic, the various thermal noise contributions can be estimated, as shown in *Figure 3.7*. The dashed lines indicate contributions from the end mirror (with solid lines denoting effects from the primary test optic), and the thick

[3]More specifically, $Q_{\text{violin}} = \frac{Q_{\text{pendulum}}}{2}$

black line is then the total noise. As can be seen, the thermal noise from the test optic dominates over that of the end mirror.

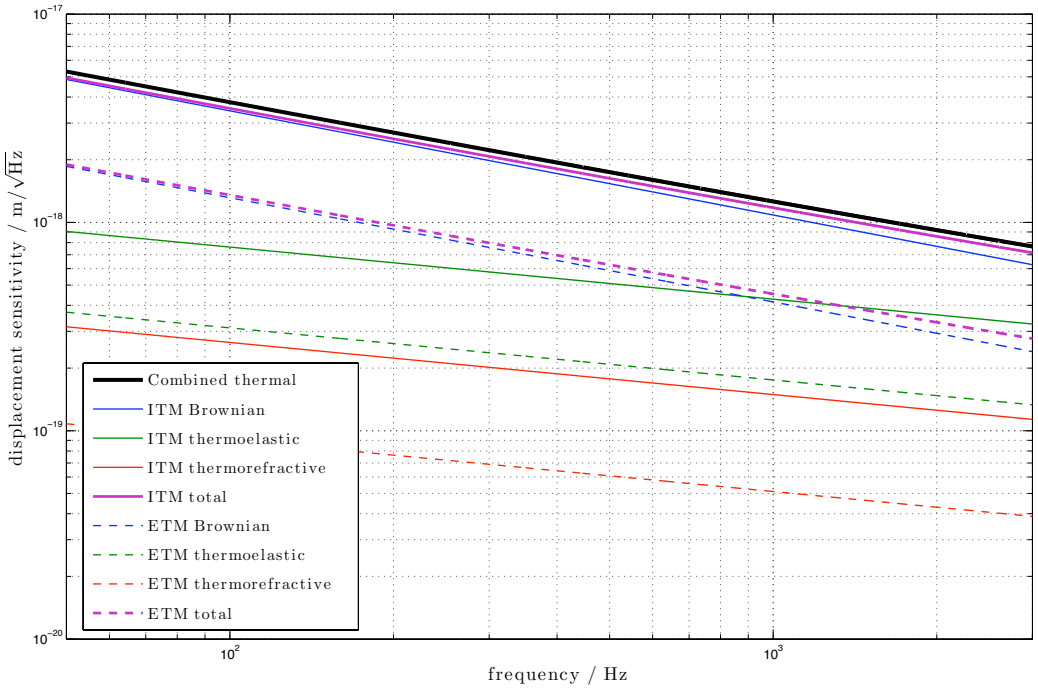

Figure 3.7: Thermal noise contributions from various sources, over the two test cavity optics. Brownian noise, thermoelastic and thermorefractive effects are shown. The overall sensitivity is dominated by the primary test optic.

Although the *AC* optics also contribute coating and substrate thermal noise to the measurement, the contribution is reduced by both the length ratio of the cavities ($\sim 10^2$) and the ratio between the respective measurement beam sizes – the beam radius on the *ITM* is 1.5 mm compared to $1.5{\times}10^{-4}$ m on the thermal noise test optic.

Finally, summing all the appropriate noise sources gives a sensitivity curve for the experiment – see *Figure 3.8*

The measurement is dominated by seismic noise at low frequencies, the thermal noise of the test cavity above ~ 100 Hz, and lastly by shot noise at frequencies above the cavity bandwidths.

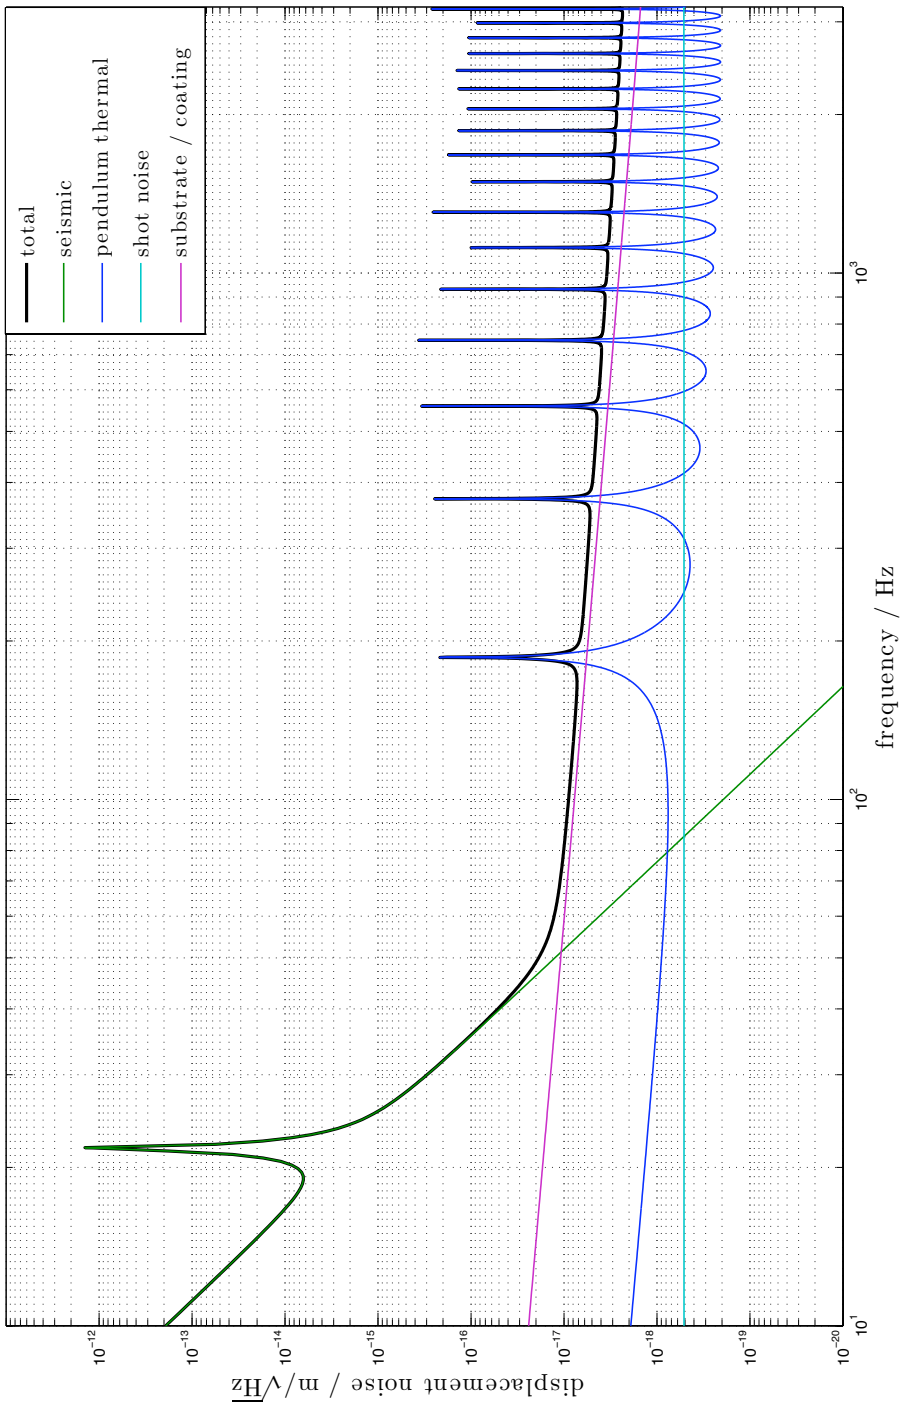

Figure 3.8: Noise budget for the *TNE* experiment. Thermal noise from the optics (dominated by the contribution from the test optic – see *Figure 3.7*) is shown as the dominant contribution in purple, the shot noise limit is shown in light blue, and the seismic noise contribution in green. The peak at $\sim 20\,\mathrm{Hz}$ is due to the vertical resonance of the silica fibre suspension elements. Pendulum thermal contributions are depicted in blue, with the numerous peaks being caused by the violin resonances of the silica fibres – no damping of these resonances is assumed here. Finally, the total sensitivity curve is shown in black.

Chapter 4

A Direct Thermal Noise Measurement Experiment: *Implementation*

4.1 Introduction

The previous chapter detailed the design and measurement ideas behind a direct interferometric thermal noise measurement experiment, which may help to inform the choice of optical substrates and coatings for future gravitational wave interferometers. This chapter expands these ideas and focuses on the implementation of the experiment. The production of suitable silica fibres, optical suspension systems, the measurement scheme and electronics will be discussed, and final results presented.

4.2 Silica Fibre Production

As detailed in *Section 3.2.3*, it is necessary to suspend the critical measurement optics in the experiment from suspension elements made from a high Q, low loss, material (such as fused silica), to reduce the contribution of the suspension element thermal noise to the measured thermal noise spectrum. Long, thin fused silica fibres are required, and must be affixed in such a way that no extra loss is induced – ensuring that the thermal noise of the coating and substrate is the dominant thermal effect.

A fibre pulling machine was constructed to produce fibres of length 57 cm, and diameter $\sim 100\,\mu$m. These were then to be welded to fused silica ears, which were silicate bonded to the masses. This bonding process introduces an aqueous hydroxide bonding solution to the two pieces of silica to be bonded – here the ears and the silica masses. The hydrophilic silica is then "etched", forming a layer of silicate gel between the pieces which – over time – becomes a solid bond layer.[1] Ideally the bond layer is ~ 100 nm thick; this can be accomplished if the bonding surfaces are polished to a flatness of $\lambda/10$, where λ is 633 nm – the wavelength of the He:Ne laser typically used to make the appropriate measurement.

As both the upper and lower masses were fashioned from fused silica, this combination of fused silica suspension elements, attachment ears and silicate bonding created a quasi-monolithic structure, where the only non-contiguous part was the thin bond layer between the ears and the test-masses. Although the bond layer can exhibit relatively significant loss, the small ratio of the bond volume to the test optic volume (and the small beam size localising the measurement away from the bond area) effectively reduces the bond's contribution to the thermal noise spectral density to insignificant levels [64] [66].

These silicate bonding and fibre welding techniques were successfully pioneered on the suspensions used in the GEO600 detector [35].

In addition to a quasi-monolithic lower stage, the silica suspensions had a wire-loop element suspending the upper mass from maraging steel cantilever blades for vertical

[1] More details on this technique can be found in other documents: [55] [64] [65]

seismic isolation: a schematic can be seen in *Figure 4.1*. Various dimensions of the suspension systems can also be found in *Table 4.1*.

4.2.1 Fibre Pulling Machine

To produce suitable silica suspension elements, a rod of Suprasil 300 grade fused silica was melted by a hydrogen/oxygen flame and rapidly stretched in a controlled fashion, creating a long cylindrical fibre as the molten silica cooled. Similar methods have been used to successfully produce both cylindrical fibres for GEO600, and rectangular cross-section fused silica "ribbons" [49].

The flame pulling machine constructed to make the fibres for this experiment can be seen in detail in *Figure 4.2*, and a close-up of the burning stage in *Figure 4.3*. The operating procedure was as follows:

Two pieces of silica stock of 5 mm diameter were held in the vertical position by two screw clamps, with each clamp being attached to rigid arms that could be moved vertically – in opposing directions – by a motor-driven chain system controlled by LabVIEW. A piece of thinner, 1.5 mm silica rod was then welded (using an electrolytic torch[2]) between the two 5 mm pieces, effectively connecting the two movable arms. Once the silica was in place, hydrogen gas was fed through the system at a rate of 10 L/min, exiting via five nozzles in a circular ring-burner (see *Figure 4.3*).

This was ignited using a standard spark gun. Oxygen was then fed in, up to a rate of 6 L/min, until the flame became a sharp blue. Gas flow rates were controlled with a combination of fine needle valves (calibrated for each gas) and coarser on/off controls. Additionally, a small fan was mounted above the ring burner to cool the upper aluminium clamp. This helped minimise excess expansion of the aluminium under the heat, and consequent slipping of the silica rod.

[2] A self-contained hydrogen and oxygen generator and burner. The flame produced is clean and easily controlled, providing ~ 100 W of heat, at high temperature, over a small area – adjustable via interchangeable burner nozzles.

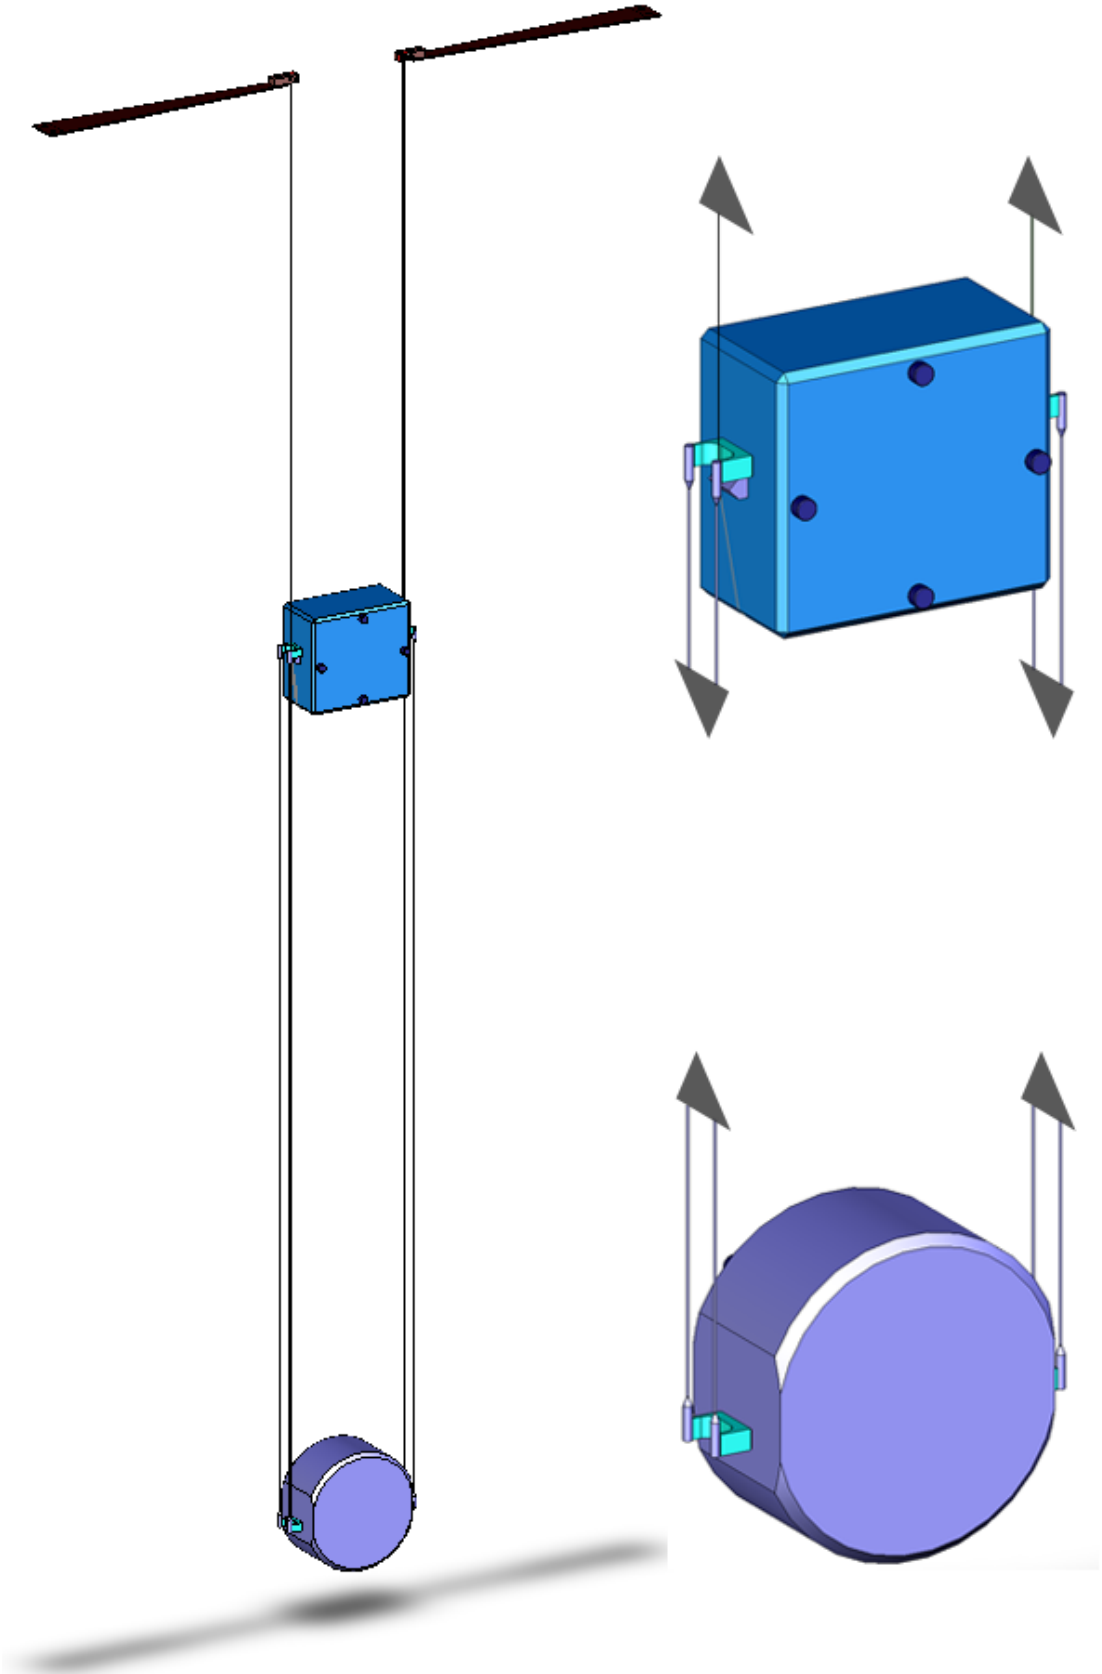

Figure 4.1: SolidWorks render [58] of a *TNE* quasi-monolithic fused silica suspension, used to suspend the measurement optics. The upper, dark blue detail figure shows the upper mass, with magnets bonded to one face, and fused silica ears and break-off prisms (to provide a defined bending point for the upper wire loop) silicate bonded to the two side faces. A single wire loop passes under the mass, connecting to the blade/clamp assemblies above. The light blue figure depicts the lower silica mass, with identical ears bonded to the flat sides – providing attachment points for the fused silica fibres that connect it to the upper mass.

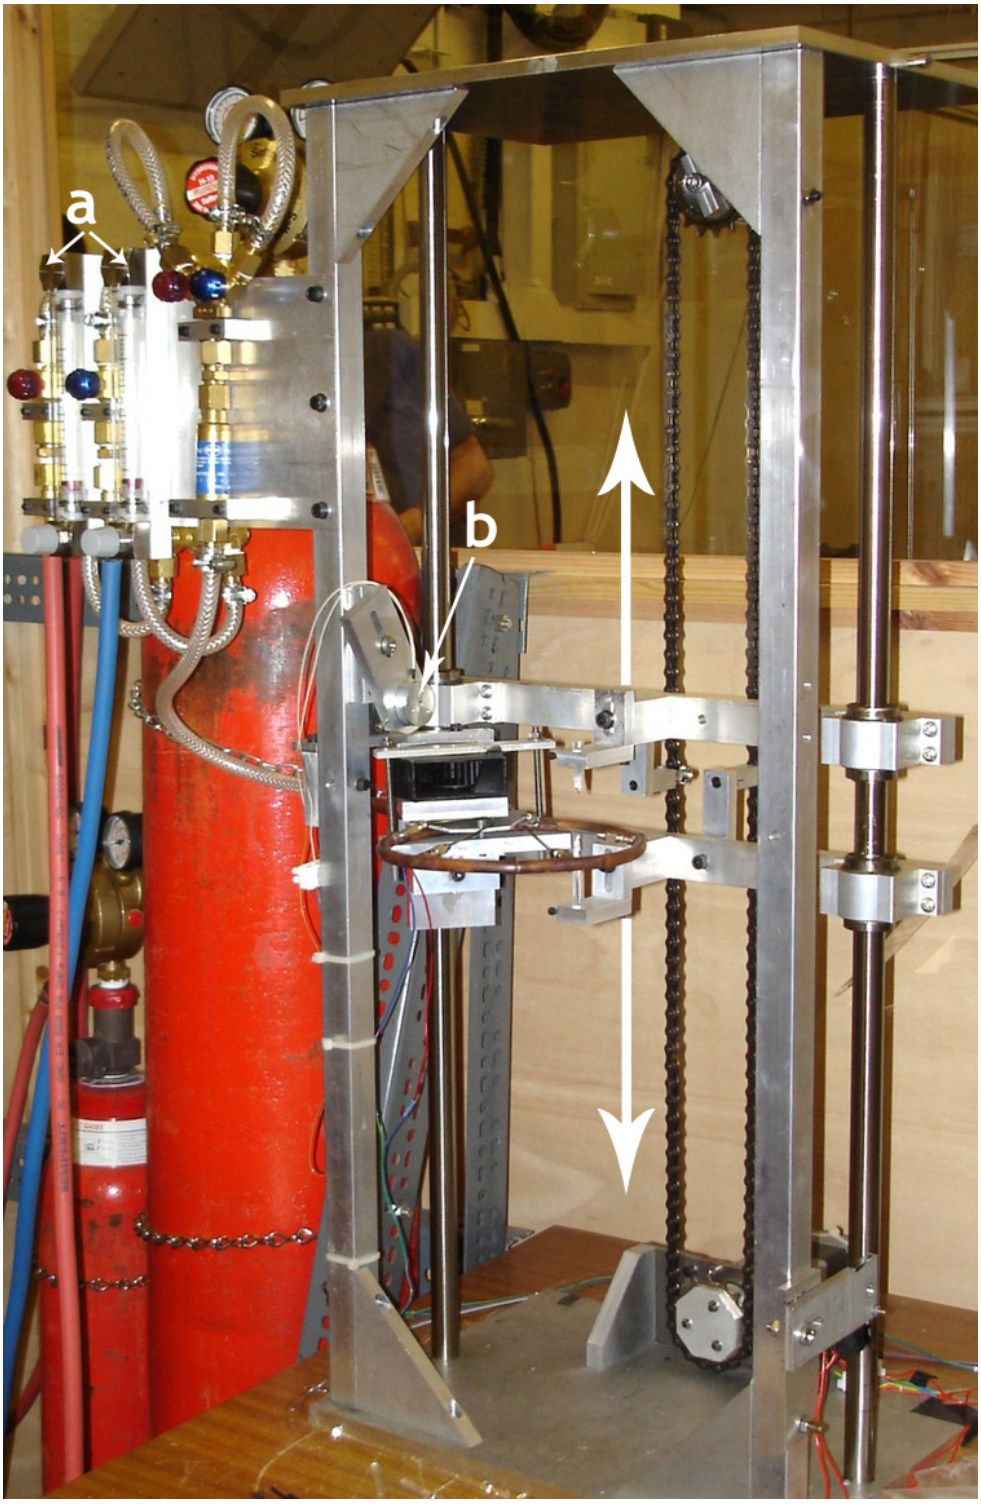

Figure 4.2: The flame pulling machine, used to manufacture the silica fibres used in the *TNE* experiment. **a:** the gas inlet valves for hydrogen (red) and oxygen (blue). **b:** the electromagnet used to hold the burner in place during the burn period. The arrows depict the motion of the arms during a pull. The burner and fan assembly can be seen in more detail in *Figure 4.3*.

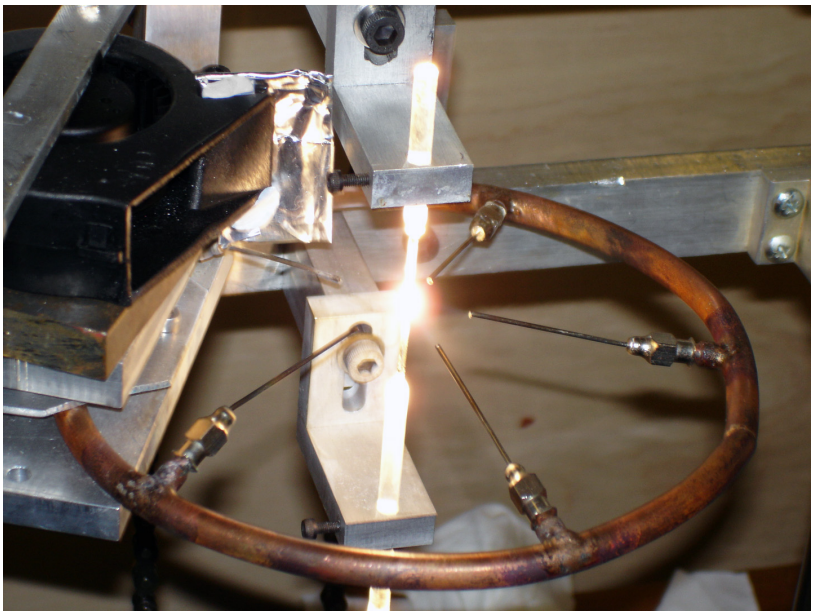

Figure 4.3: Detail of the pulling machine ring-burner. Five nozzles direct the flow of hydrogen towards a central point. After this is ignited, oxygen is fed in and the ring burner swung into place – with the centre coincident with the silica rod. A small black fan can be seen on the left, cooling the upper clamp to prevent slippage of the silica.

Once ignited, the ring burner was swung into position, with the confluence of the flames coinciding with the thin silica rod. This motion tripped a microswitch, signalling to the *LabVIEW* control system that the burn period had begun. This caused an electromagnet to engage, holding the ring burner in position for a set time (typically 2 seconds) before releasing, whereupon the burner swung away from the silica and the motor was driven to separate the clamping arms, thus stretching the molten silica. The arms were then stopped at the required length by the *LabVIEW* system, and additionally a fail-safe switch cut the power to the motor before the lower arm hit the end of its rail. The gas to the burner, and the fan, were then turned off manually. Finally, fibres were carefully removed using a diamond file and tweezers, before being placed in a perspex shielded box to protect them from dust and accidental damage.

4.2.2 Fibre Characterisation

Using this method, fibre lengths were found to vary by $\sim 5\,\mathrm{mm}$, so care had to be taken to select fibres of consistent length. The storage box was designed to be the correct height with respect to a reference fibre, and fibres that proved too long or too short

were discarded.

Additionally, fibres were found to be asymmetric in diameter, often (but unpredictably) being thinner at the lower end. The most homogenous fibres were chosen by carefully bending the fibre into an Ω shape. The least symmetrical fibres exhibited an obvious asymmetry and were discarded.

The origin of this asymmetry was investigated, and found to be independent of both burn time and pulling speed. Another possible factor was the asymmetric cooling imposed by the fan.

As the fan was turned off manually after each pulling cycle, it continued to operate during the pulling stage. Although the burner (and hence the fan) swung away from the fibre at this point, it is possible that the residual airflow caused an asymmetric cooling. Several tests were performed without the fan, but this frequently led to the upper silica piece slipping and as such it was not possible to produce reliable fibres for comparison. Similarly, the burner was also turned off manually after the pulling cycle, providing an asymmetric source of heat. Due to the rapidity of the pulling stage, it was not possible to manually shut off both the power to the fan and the gas to the burner in order to evaluate this effect properly.

Another possible source of the asymmetry was the variable angle and tension of the silica rod. As the 1.5 mm diameter stock was too thin to clamp sufficiently tightly without breaking, it was welded between two thicker pieces of silica, which were held in the clamps permanently. Each time a fibre was pulled, a short piece of 1.5 mm rod was welded to the upper thick silica piece. This was then heated with the electrolytic torch until the newly welded piece loosened, and became vertical due to gravity. The two arms of the pulling machine were then brought together, with the lower end of the thin rod meeting the thicker silica. If not correctly aligned at this point, the upper weld would have to be redone. Even when optimally aligned, a slight angle and tension would be introduced to the rod on releasing the lower arm after the lower weld – see *Figure 4.4*. Certainly, it was noted that particularly bad welds – with a visible angle imparted to the thin silica stock – would lead to either heavily asymmetric fibres or

even cause the fibre to break mid-pull. It seems that this method of clamping and welding was the likely source of the fibre asymmetry.

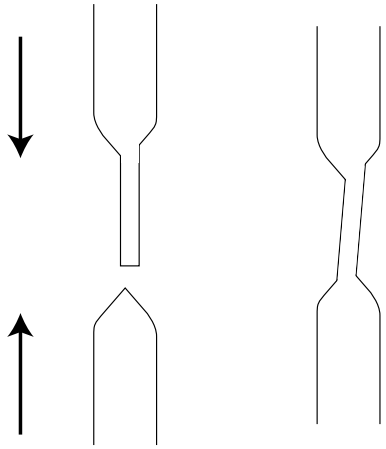

Figure 4.4: Representation of a "bad weld". The thinness of the silica stock used to create the fibres caused a systematic uncertainty in the pulling process, leading to non-uniformity and non-repeatability of the fibres.

The diameter of the fibres was determined by measuring their vertical bounce frequencies and length, and relating this to the fibre thickness through the relation below[3],

$$f_{\text{vert}} = \frac{1}{2\pi}\sqrt{\frac{Y\pi r^2}{lm}}, \tag{4.1}$$

where Y is the Young's modulus of the fibre material, f_{vert} is the vertical bounce frequency, l is the fibre length, m the mass suspended, and r the fibre radius.

The apparatus used for these measurements can be seen in *Figure 4.5*. Typically, the fibres were found to have a vertical bounce mode of $\sim 22\,\text{Hz}$, corresponding to a diameter of $\sim 120\,\mu\text{m}$.

In summary: using the flame pulling machine – and bounce test rig – fused silica suspension elements satisfying the design requirements were produced. Although some fibres exhibited asymmetry in diameter and variability in length, these were easily discarded in favour of more consistent examples.

[3]this work was carried out by Mr Nicolas Crouzet.

Figure 4.5: Apparatus for "bounce testing" the silica fibres, to determine their average diameter. Fibres were clamped at the upper stage (**a**), and attached to a free mass (representing the design payload) at the lower end – **b**. A magnet attached to this lower mass was then excited by a coil (**c**), driven by the voltage output of a signal generator. The frequency of maximum vertical response of the system could then be measured by observing the trace on an oscilloscope of a flag moving past a split-photodiode (**d**).

4.2.3 Violin Modes and Fibre Coating

When using silica fibres as a suspension element, one must take account of the high mechanical Q of their transverse violin modes. These modes form a slightly anharmonic series: recall *Equation 2.26*

$$f_n = \frac{n}{2l}\sqrt{\frac{T}{\rho_{\text{lin}}}}\left[1 + \frac{2}{l}\sqrt{\frac{EI}{T}} + \frac{EI}{2T}\left(\frac{n\pi}{l}\right)^2\right], \tag{4.2}$$

where n is the mode number, T is the load on the fibre, ρ_{lin} is the linear density, l is the length of the fibre, E is the Young's modulus and I is the moment of inertia. As relatively long fibres are required (compared to, for example, the ~ 30 cm GEO600 silica suspension elements) the violin modes will therefore occur at relatively low frequencies for a given fibre diameter – see *Figure 4.6* below.

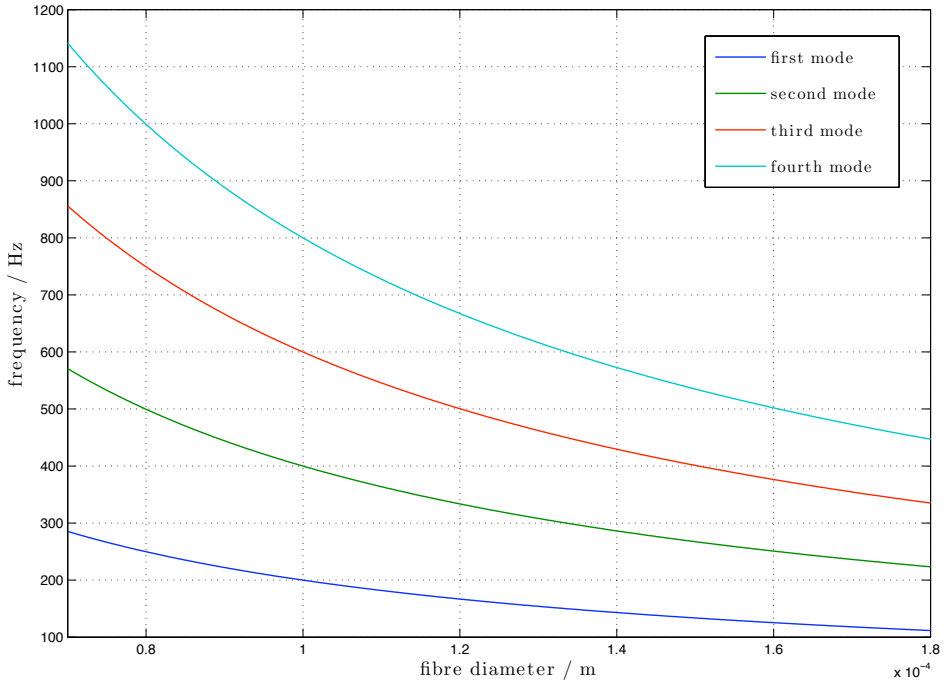

Figure 4.6: Relation between fibre thickness and violin resonances for a 57 cm fused silica fibre.

Given that fibre diameters of $\sim 100\,\mu$m were used, violin resonances from ~ 200 Hz upwards could be expected. This is within the measurement band where we wish to characterise the thermal noise due to the optic substrate and coating. The violin modes

can also be expected to have mechanical quality factors of the order of $\frac{Q_{\text{pend}}}{2}$ where Q_{pend} is the Q of the fundamental pendulum mode. As these transverse resonances are both low loss *and* in the desired measurement range, they must be either damped or controlled.

A technique used in the GEO600 suspension systems [35], is to selectively coat the fibre surface at specific points – halfway and a quarter length along the fibre – with a lossy material. This acts to damp the anti-nodes of the violin mode resonances, while preserving the mechanical Q of the fundamental pendulum mode – required for low pendulum thermal noise. The damped fibres used in the GEO600 detector show a reduction in violin mode Q of a factor of ~ 100 [67].

To achieve this, a coating jig was constructed where fibres could be loaded vertically and clamped securely, before being laid horizontally and coated with a nitrocellulose lacquer[4] (pigment free commercial nail varnish) in the correct locations over $\sim 2\,\text{cm}$ lengths. A sliding magnifying glass and fixed metal rule were attached to the jig to ensure repeatability. The fibre coating jig can be seen in *Figure 4.7*, and a coated test fibre in *Figure 4.8*.

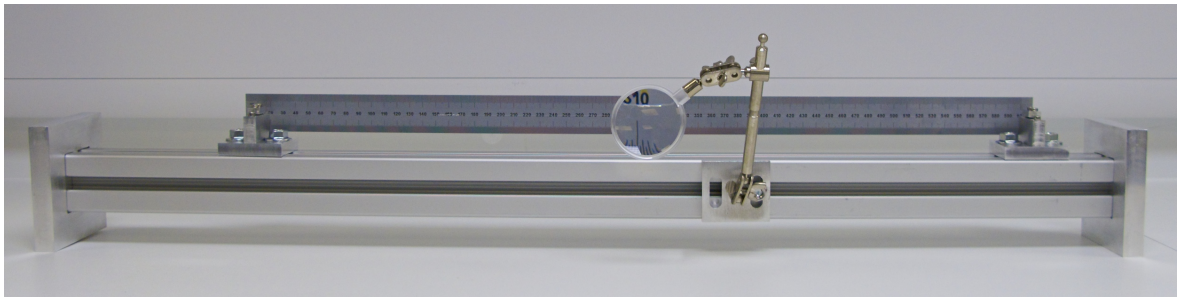

Figure 4.7: The fibre coating jig. Fibres were loaded with the jig in a vertical orientation and clamped in place. The jig could then be placed horizontally allowing the fibres to be coated. Slots in the fibre clamps allowed a metal rule to be screwed in place, and a sliding magnifying lens assembly allowed detailed viewing of the coated areas during and after coating.

The coating was applied using a piece of silicone tubing, tapered at the application end. The hollow tubing was partially filled with the lacquer, and could be gently

[4] A commercial epoxy (EPO-TEK 301[68]) was also tested but found to lack the appropriate viscosity to reliably adhere to the fibres for the required curing time.

squeezed to steadily release it onto the fibre surface. Care had to be taken to apply successive applications relatively quickly (under one minute per area), as the coating would start to set, leading to thin filaments breaking away when new layers were applied. These trailing strands could potentially contact the fibre surface, causing microcracks and a consequent degradation of the fibre's strength.

4.3 Silica Suspension Installation

Once suitable fibres had been produced, the silica suspensions could be assembled and installed in the measurement tank. The suspensions were first assembled in a jig, for the purposes of welding the fibres in place and correctly balancing the system. Once welded and finally balanced, the jig was then carefully lowered into the measurement tank for installation, and then removed once the suspensions were secured. The assembly process is summarised below:

- The silica masses were placed in their respective positions (see *Figure 4.10*, positions **c** and **e**) and secured on their front faces by small *Teflon*® pieces that were rotated into place, and on their flat sides by aluminium plates which held the masses in the lateral direction.

- A pre-assembled long wire and clamp assembly (constructed on a separate jig) was then attached; with the two clamps screwing into the topmost part of the

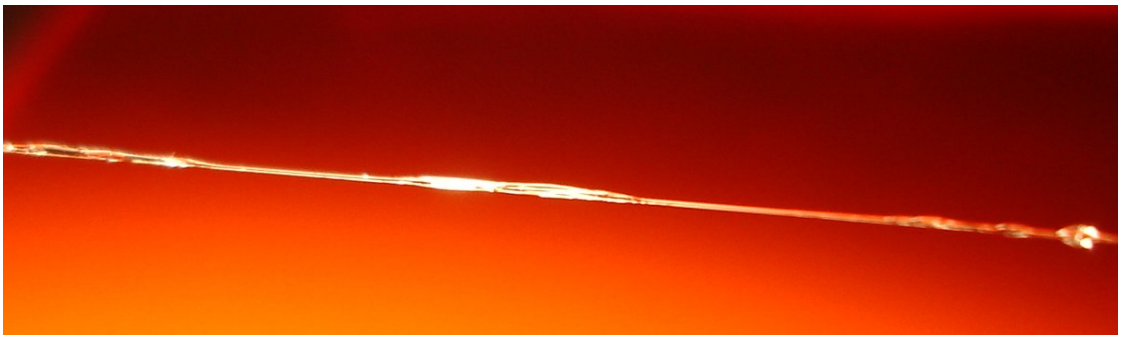

Figure 4.8: Detail of a test silica fibre, coated with damping material over three areas. Care had to be taken to ensure that the coating was as homogenous as possible and did not form unwanted globules – see the rightmost area.

installation jig (*Figure 4.10***a**), and the wire itself looping underneath the upper silica mass. As the wire was slack in this condition, it was typically held with a black wire-wrap (*Figure 4.10***b**), to ensure that it didn't interfere with either the silica ears or – later in the installation process – the silica fibres. See *Figure 4.1* for a schematic reference of the suspension layout.

- It was then possible to suspend the upper mass by lowering the platform it was resting on – (*Figure 4.10***d**). The suspension point of the jig could then be moved forward, allowing the mass to exhibit any tilt imbalances without being constrained by the jig. If any imbalance was observed at this point, it could be coarsely corrected by slightly shifting the break-off point of the wire on the silica break-off prism.

- With the installation jig in its secure position (that is, with neither mass free) the silica fibres could then be welded. As the wire-loop for the upper suspension stage passed through – and close to – the silica fibres (the distance between the weld points on the ears was 13 mm), it was necessary to hold the wire away from the weld area during welding. The procedure for welding the fibres was as follows:

  Fibre welding

  – One upper clamp was removed from the upper suspension point and attached to a post (screwed into the same optical breadboard that the jig itself was attached to) such that the wire was then horizontal with respect to the suspension.

  – A fibre was selected, carefully removed from its storage box and placed in a vertical jig consisting of two long metal tweezers rigidly attached to an aluminium bar. This held the fibre securely while welding.

  – The fibre was welded quickly – merely enough to keep it in place – at the upper ear position. This was then repeated at the lower ear.

  - When the fibre was securely attached to both ears, the welds could be improved. This was done by heating the ear, and gently pushing the butt of the fibre inwards with a silica tool piece, then thoroughly heating both pieces and repeating the process. See *Figure 4.11*.
  - Once the upper weld was satisfactory, the lower weld could be attended to in a similar fashion.
  - Typically at this point it was necessary to adjust the tension in the fibre. As such the welding at the upper and lower ears was something of an iterative process: completely finalising the weld meant that the butt of the silica (i.e. the end of the fibre) became subsumed into the ear, and any further adjustments would be difficult.
  - This process was then repeated with the other fibre, after which the wire could be replaced at its upper suspension point, and the welding process repeated for the two fibres on the other side.

- Once all four fibres were welded satisfactorily to the ears, the lower mass could be suspended from the upper mass point (*Figure 4.10**d***) to observe any imbalances. If present these imbalances could be corrected by adjusting the tension in the fibres as described previously. A suspended optic can be seen in *Figure 4.9*.

- Now the system could be fully suspended, with both masses free. Again, at this point any imbalances could be observed and corrected. It was seen that the lower mass typically exhibited some tilt – this was largely corrected through bonding steel washers to the upper mass. No extra masses could of course be bonded to the lower optic, and at this point it was difficult to adjust the fibres further without risking the loss of a fibre (and potentially its neighbouring fibre). The ETM_{TNE} suspension in particular exhibited a significant tilt, as the strong curvature of the mirror surface caused a mass deficit on the front half of the optic.

- Once finally balanced, the suspension was secured (with the cantilever blades

now in place) and transferred to the measurement tank where it was carefully lowered into place. The roots of the cantilever blades could then be attached to the upper frame section in the tank and the suspension released. In reality this was a time consuming and risky process, as the blades could only be fully secured when the upper mass was suspended. As such, an iterative routine of tightening the blade's screws, and lowering the upper mass, was performed. Additionally, it was difficult to release the lower mass due to its depth in the tank, and the risk of accidentally touching the silica fibres.

- Immediately after suspension in the tank, aluminium catchers were slid into place underneath the lower masses, and the eddy-current damping and alignment units placed behind the upper masses. These also served as rudimentary catchers for the upper masses, while the eddy-current damping effect significantly reduced the residual motions of the suspensions.

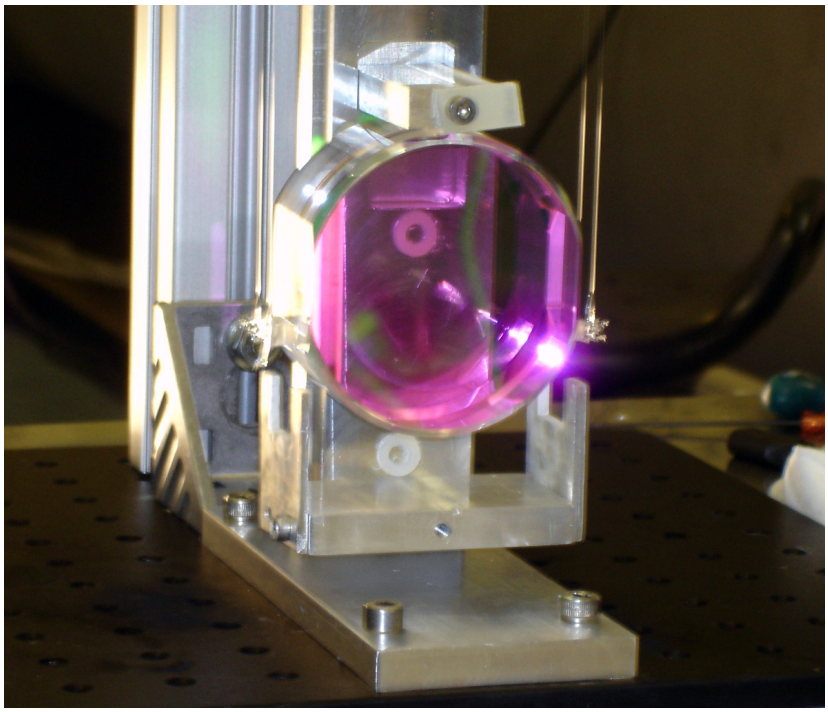

Figure 4.9: Detail of a suspended thermal noise cavity optic in the installation jig. Here the lower platform of the jig has been lowered, suspending the test optic from the upper mass via the welded silica fibres.

The ETM_{TNE} suspension was installed and aligned such that the return beam returned to the input laser bench. The ITM_{TNE} suspension followed, and the MU suspension placed such that the return beam from the inboard optic exited through the measurement tank viewport – recall *Figure 3.4* for the internal tank layout. Note that all the aluminium suspensions were installed prior to the silica assemblies to lessen the risk of accidental damage. The MU and RM suspensions were moved to the end of their respective ranges beforehand, and shifted back carefully afterwards.

4.3.1 Installation Challenges

Although the suspensions were successfully installed according to the processes outlined in *Section 4.3*, there were a number of setbacks, which may inform the design of similar experiments.

The first installation attempt of the ETM_{TNE} failed due to an excess roll present on the upper mass that was not visible when suspended in the jig. On attempting to correct for this, the upper mass violently snapped back into the kinks it had made in its suspension wire. This could have been averted if cylindrical upper stages had been used. To alleviate this as best as possible, the upper wire loop was subsequently kept taught during welding and the damping/alignment units for the silica masses were re-machined to allow for small roll offsets.

The second ETM_{TNE} installation failed due to an excess of tilt in the system. Initially, the installation jig functionality allowed the assembly to be suspended from the uppermost point and the upper mass stage, but did not allow the suspension to move forward. This effectively masked the level of tilt that was actually present in the suspension, as the surface of the jig partially constrained the suspended optics. On installation the system tilted dangerously and induced a slow bouncing motion which quickly exacerbated the tilt further, causing a fibre to fail. To correct for this the jig was altered to provide the facility to shift the suspension point forwards, allowing the masses to tilt freely. It became apparent that stainless steel counterweights had to be

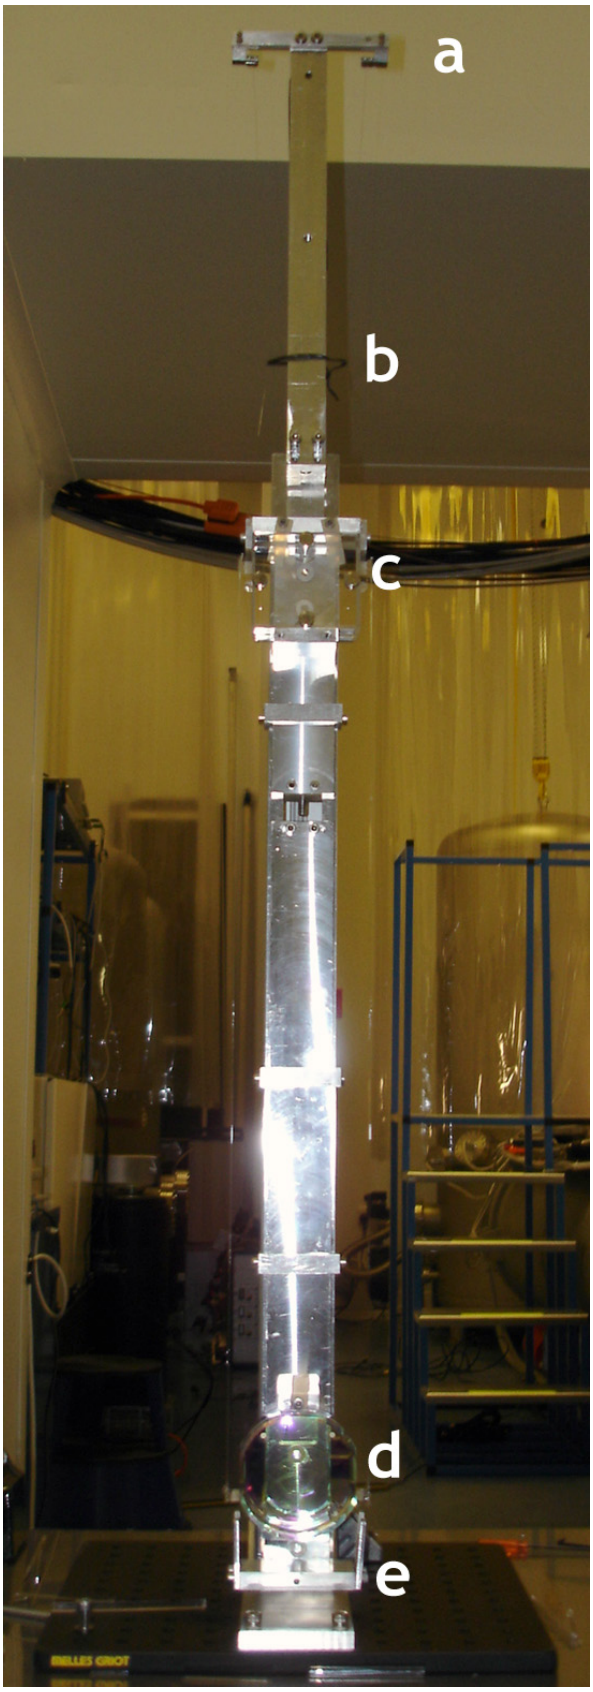

Figure 4.10: A partially suspended and fully welded silica suspension, assembled in its installation jig. The jig allowed both the upper and lower masses to be suspended so that any imbalances in the suspension could be corrected. In this figure the lower stage only is suspended – the wire suspending the upper stage (which is slack when the upper stage is not suspended) is constrained by a wire-wrap (**b**) to stop it interfering with the silica ears or fibres.

Figure 4.11: Test welding a silica fibre to an example ear. The fibre is first held in a tweezer jig, and manipulated using silica pieces as tools. The inset shows an example ear and the different stages of fibre welding: the area circled is considered monolithic, the area immediately below this is unfinished, and the near fibre is only attached via a placeholder weld.

bonded to the front side of the upper masses to counteract the mass of the magnets bonded to their reverse sides.

The silicate bonding of the ears to the upper masses was also shown to be suspect. On the third ETM_{TNE} installation attempt, a bond failed while a fibre was being welded to the associated ear. Notably, this weld location had been thermally cycled several times due to the previous aborted installation attempts. The side of the mass was then tested on a Fizeau interferometer to quantify the flatness, but it was not possible to measure reliably due to the excess residue localised around the ear – see *Figure 4.12*.

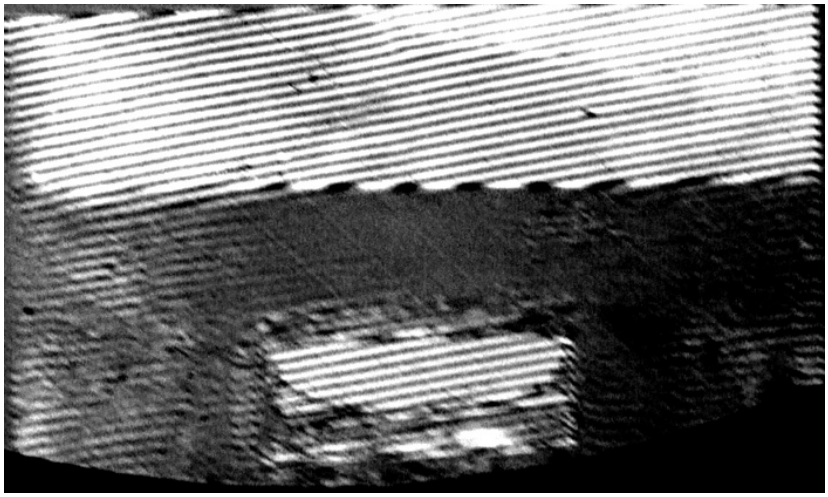

Figure 4.12: Attempted flatness measurement of the area around a failed silicate bond. The ear area is clearly visible, as is the contact region the bond had with the mass.

An identical spare mass (with no ears or break-off prisms bonded to it) was measured however, and shown to have one side polished flat to the required $\lambda/10$ level (where $\lambda = 633$ nm), with the other side only $\sim \lambda/3$ (see *Figure 4.13*), indicating that the vendor had neglected to superpolish both sides. Additionally, on close inspection another fully assembled upper mass exhibited a visible gap between the ear and the mass surface: this ear failed immediately on welding – see *Figure 4.14*.

Although the flatness requirements for silicate bonding are well known, this experience emphasises the need for stringent testing and characterisation of surfaces before committing to a bond.

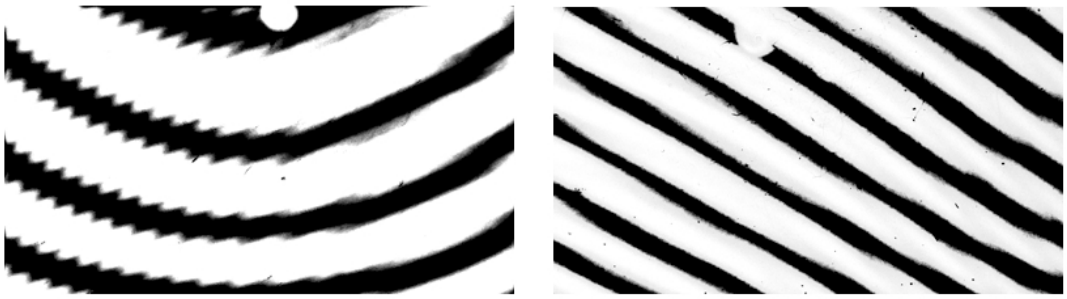

Figure 4.13: Interferometry measurement of the relative flatness of two sides of an upper stage mass. One side (the figure to the right) is seen to be polished to the requisite $\lambda/10$, where the other side is only $\sim \lambda/3$.

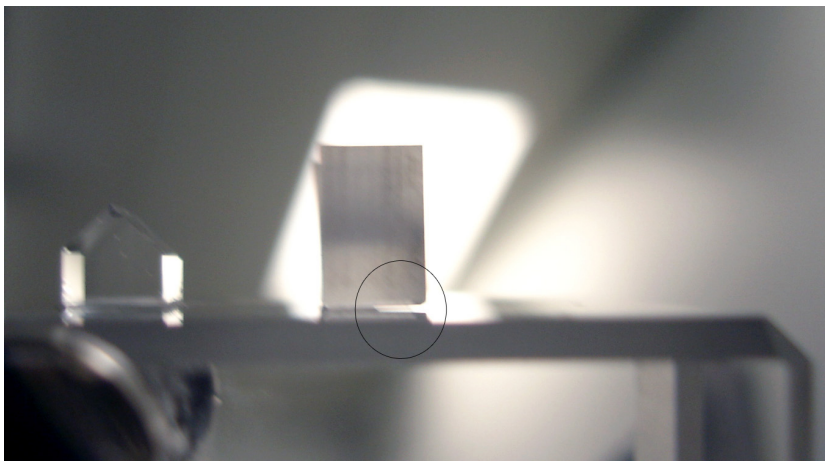

Figure 4.14: A failed silicate bond. The circled area shows a clear air gap between the ear and the surface of the silica mass. This ear immediately failed on attempting to weld a fibre.

4.3.2 Metal Suspensions

Recall from *Section 3.2.2* that there were five non-silica suspensions in the measurement tank: the reaction mass, the mounting unit, the beam-splitter, a mode-matching optic and a steering optic. The aluminium suspensions resembled the silica ones in most regards: the differences being that they were all-metal, had central holes cut in the lower masses for mounting standard mirror optics, they had wire and clamp assemblies instead of silica elements and they were suspended from simple plates bolted to the top of the frame in the tank, rather than from cantilever blades. The reaction mass suspension was essentially a hybrid of the two designs, being all aluminium, but suspended from cantilever blades, with the lower mass machined with recesses for coils (to provide the feedback to the ETM_{TNE}). These suspensions were trivially installed. A schematic can be seen in *Figure 4.15*.

	silica	**metal**
upper stage mass	360 g	348 g
lower stage mass	370 g / 350 g	382 g
upper stage length	38.5 cm	38.5 cm
lower stage length	58.5 cm	56 cm
upper stage dimensions	$69.3 \times 37 \times 62$ mm	$69.3 \times 37 \times 62$ mm
lower stage dimensions	$69.3(76.2) \times 38.1$ mm	$69.3(76.2) \times 38.1$ mm

Table 4.1: Properties of the thermal noise experiment suspensions. Note that the ETM_{TNE} lower mass was ~ 20 g lighter due to the mirror curvature, and that the masses quoted are with additions such as ears or clamps. The lower stage length for the silica suspension given was the distance from the top of the upper ear, to the bottom edge of the lower ear. The values in brackets represent the diameter of the cylindrical masses, as opposed to the distance between the flat sides.

The upper stages of all the suspensions had their fundamental pendulum mode resonance damped via a combined damping and alignment unit, mounted behind the respective masses[5]. The units consisted of an L-shaped bracket (which also served as a rudimentary catcher for the upper stages) holding a vertical plate, upon which four copper coil-formers were attached. These copper pieces were hollow and designed to allow the 10 mm wide magnets bonded to the masses to pass inside with 1 mm total clearance. They also served as formers for enameled wire coils which served as

[5]Note that the damping/alignment units were mounted in *front* of the ETM_{TNE} and ITM_{TNE} due to the close proximity of the other suspensions.

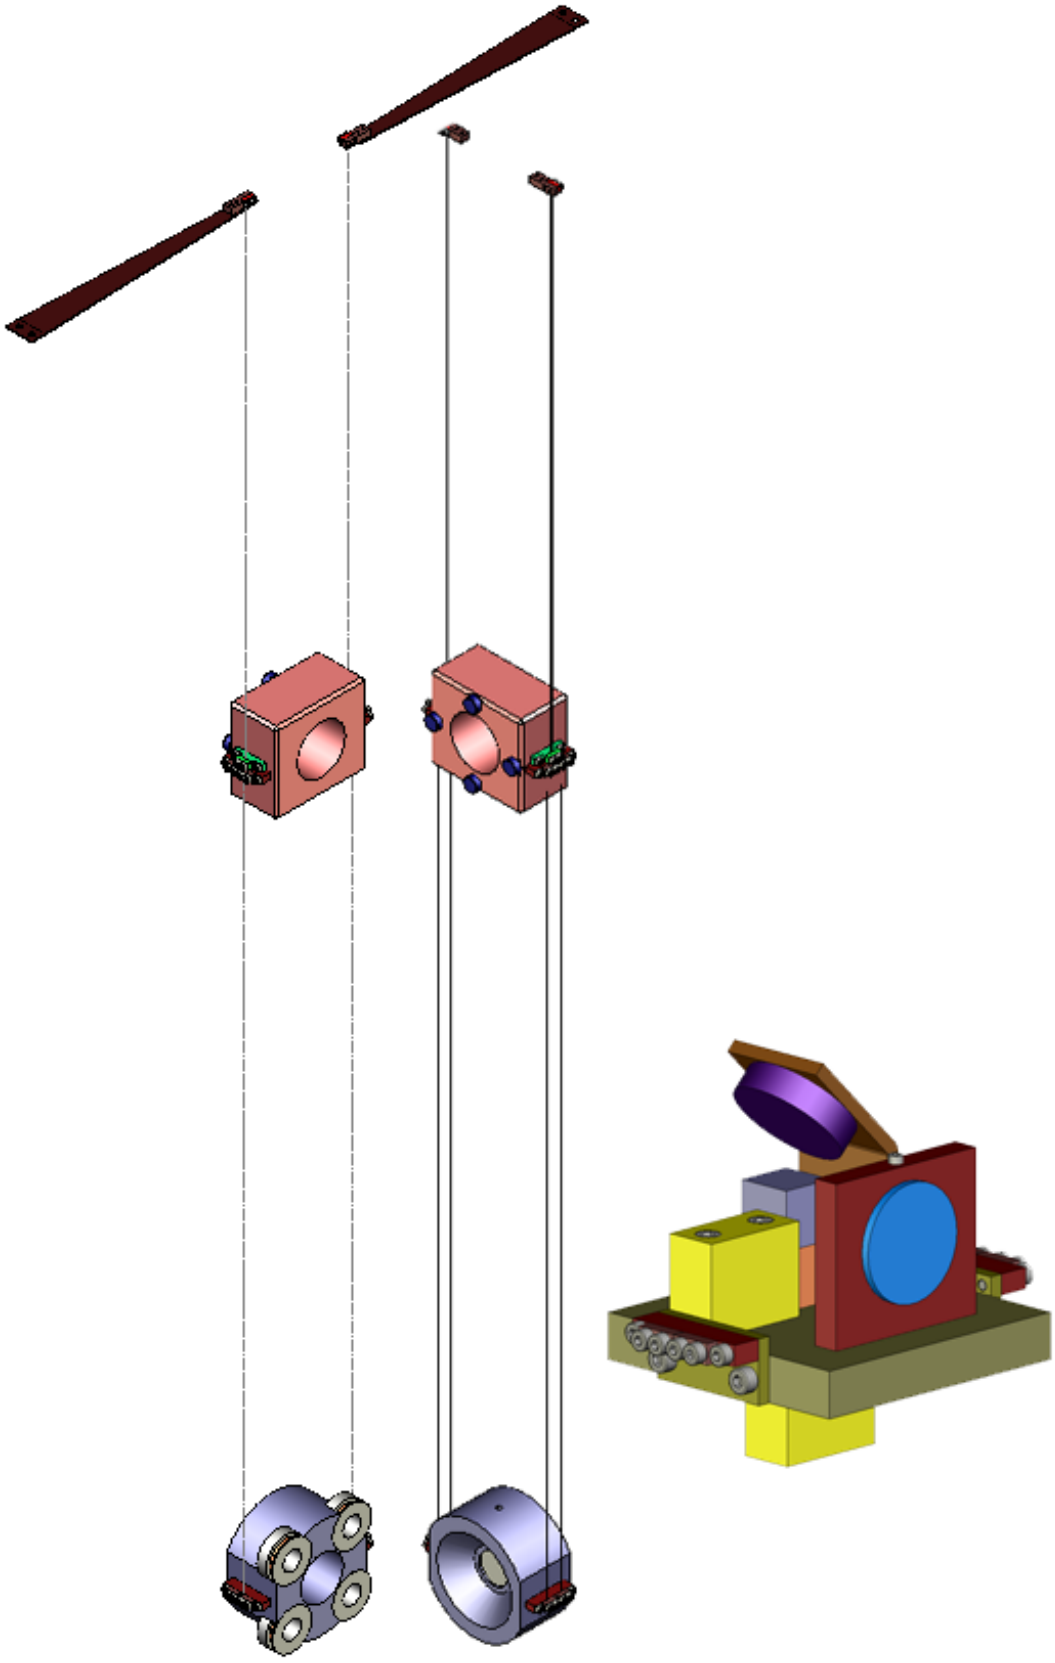

Figure 4.15: SolidWorks renders [58] of the TNE aluminium suspensions. The reaction mass (RM) suspension is shown to the left, suspended from cantilever blades and sporting coils attached to the lower mass. The BS, $S1$ and MM suspensions followed the design of the central diagram. The MU suspension had a markedly different lower stage – detailed here at the right – comprising a polarising beam-splitter (grey), quarter wave-plate (blue) and mirror (purple) to act as an optical circulator.

non-invasive alignment controls for the suspensions: see *Figure 4.16*.

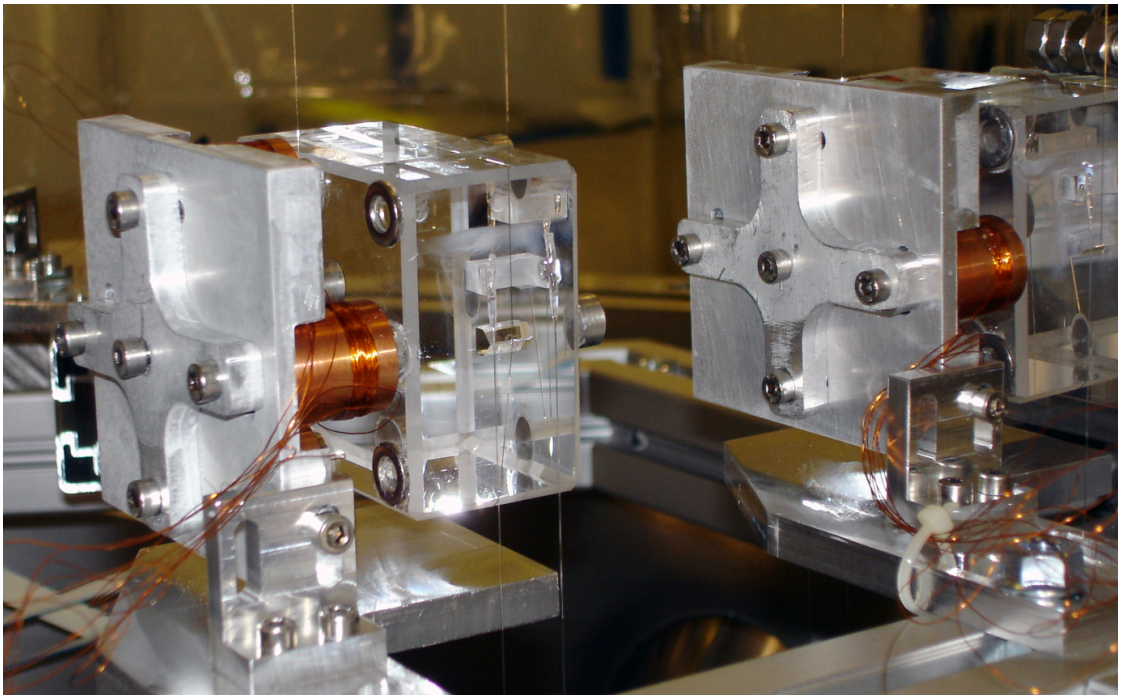

Figure 4.16: The silica upper stages of the thermal noise test cavity. The eddy-current damping and alignment units can be seen mounted behind the upper silica stages. These provided alignment control and damping of the fundamental suspension modes.

The alignment coils were soldered to 10-pin connectors, which attached to ribbon cables that exited the tank through one of two feedthroughs. Each coil could be addressed individually, but to simplify the alignment control they were wired in pairs with opposing polarity: each suspension could then be adjusted in both tilt and rotational degrees of freedom. This was achieved with simple potentiometer and op-amp circuits, applying adjustable voltages across the coils – see *Appendix B.1*.

The nominal clearance of 1 mm between the magnets and the copper caused some difficulty in the installation of the units. Although the plate holding the copper formers was free to be adjusted in tilt, and the main bracket could be easily rotated, it was difficult to judge visually whether the magnets were contacting the inner sides of the copper formers.

Additionally, on aligning the individual suspensions, the clearance of the magnets with respect to the copper formers would change, occasionally leading to contact, and a subsequent manual rotation of the suspension and damper was required to correct

the effect. It was necessary to adjust the position of the dampers several times on all suspensions before proper alignment and non-contact was achieved.

These installation difficulites were exacerbated with the silica systems, for two reasons. Firstly, it was not desirable to move the silica systems manually once installed, and secondly, it became apparent that the cantilevered systems exhibited slight offsets in roll – making installation of the regular dampers impossible without causing magnet contact. It was then necessary to retrofit the damper and alignment units with a roll adjustment feature, visible as the cross-shaped structure on the rear of the devices as shown in *Figure 4.16*.

4.4 Laser Frequency Stabilisation

As discussed in *Section 3.2.1*, it was necessary to stabilise the frequency of the measurement laser. This enabled the displacement noise of the thermal noise test cavity – and specifically the single test optic – to dominate the measurement.

4.4.1 Requirements

The laser frequency stabilisation was achieved through locking the frequency of the 2 W Nd:YAG laser to the length of a long (with respect to the thermal noise measurement cavity) optical cavity system. We wish to measure thermal noise of the order of $\sim 10^{-17}$ m/$\sqrt{\text{Hz}}$ in a 10 cm test cavity, with the laser stabilised to the length of a 9.78 m reference cavity. It was therefore decided to aim for a level of frequency stability of, or better than, $\sim 1.4 \times 10^{-4}$ Hz/$\sqrt{\text{Hz}}$. This level of frequency stabilisation has an equivalent displacement noise in the test cavity of 5×10^{-20} m/$\sqrt{\text{Hz}}$: some way below the predicted thermally induced displacement noise signal, as depicted in *Figure 3.7*.

4.4.2 Feedback Scheme

As outlined in *Section 3.2.2*, the feedback scheme for the stabilisation of the laser frequency was achieved via a standard *Pound-Drever-Hall* scheme, where phase modu-

lation sidebands are imposed on the laser light by an EOM before entering the system. These sidebands, displaced from the main carrier light by 46.526 MHz (see *Section 5.3.2* for discussion on the choice of modulation frequencies) then interfere with the phase shifted carrier light returning from the cavity, creating a beat signal that is then detected on a photodiode (similar to the one used for the TNE measurement, detailed in *Appendix B.3*) and demodulated by a mixer, thus generating a suitable locking signal – see *Figure 3.5*.

In addition to the photodiode/mixer combination, the sensing scheme also contained a signal generator and power amplifier to provide the signal to the EOM, and a low-pass filter and amplifier to send the demodulated output signal to a control room above the interferometer lab, where the signal generators and frequency stabilisation servo were housed.

4.4.3 Servo Design & Performance

The frequency stabilisation servo was required to provide the necessary level of frequency stability over at least the measurement range of the thermal noise experiment and operate stably, keeping the reference cavity locked (and thus the frequency of the laser suitably stable) for indefinite periods of time, to enable easy commissioning of the thermal noise cavity locking scheme.

The servo acted over three separate paths, governing different feedback frequency regimes. Low frequency feedback ($< 1\,\text{Hz}$) was governed by coil/magnet feedback to the position of the ETM mirror – the end mirror of the reference arm cavity. At high frequencies ($> 9\,\text{kHz}$) the feedback was governed by an EOM on the laser bench. Between these two regions, a PZT mounted on the laser crystal was the dominant feedback element. The feedback servo (see *Appendix B.6* for further details) was designed to provide both a high open-loop gain (and thus high suppression of frequency noise) over a wide band – with the high frequency feedback element (the EOM) having a unity gain frequency at $\geq 100\,\text{kHz}$ – and stability according to the criteria outlined in *Section 2.3.2*, while taking into account the frequency responses of the arm cavity and

the feedback elements. A diagram of the laser frequency stabilisation circuit can be seen in *Figure B.9*, and swept-sine responses of the different feedback paths (compared to modelled outputs) in *Figure 4.17*.

Assuming that the free-running frequency noise from the laser is approximately $\sim 10^4$ Hz/$\sqrt{\text{Hz}}$ at 1 Hz [57] and decreases proportionally with frequency, it is possible to estimate the closed-loop frequency noise of the laser when the arm cavity is held locked by the frequency stabilisation servo – see *Figure 4.18*. The modelled close-loop frequency noise is seen to meet the requirement from *Section 4.4.1* above.

4.5 Thermal Noise Measurement

Recalling *Section 3.2.2* and *Figure 3.5*, the full measurement scheme of the thermal noise experiment consists of the laser frequency stabilisation path discussed above, and another feedback loop. This feeds back to the position of the ETM_{TNE}, thus locking the thermal noise test cavity using the light stabilised from the long reference cavity, where the error-point signal for the position feedback can then be interpreted as the displacement spectrum of the cavity; and therefore the thermal noise associated with the ITM_{TNE}, as discussed throughout *Chapter 3*.

Through locking the cavity with a low unity gain frequency (ideally ~ 150 Hz) and measuring the error-point, the effect of noise sources due to the feedback process can be ignored, and the noise due to the sensing then dominates. This of course imposes the restriction that the unity gain frequency be low, which can complicate the design and operation of the feedback servo, especially when violin resonances are expected to lie outside the nominal control bandwidth.

4.5.1 Measurement Scheme

The specific layout of the TNE cavity locking and measurement scheme can be seen in *Figure 4.19*. The error signal required to lock the cavity is produced in a similar way to that for the arm cavity: phase modulation sidebands at 10 MHz (an arbitrary

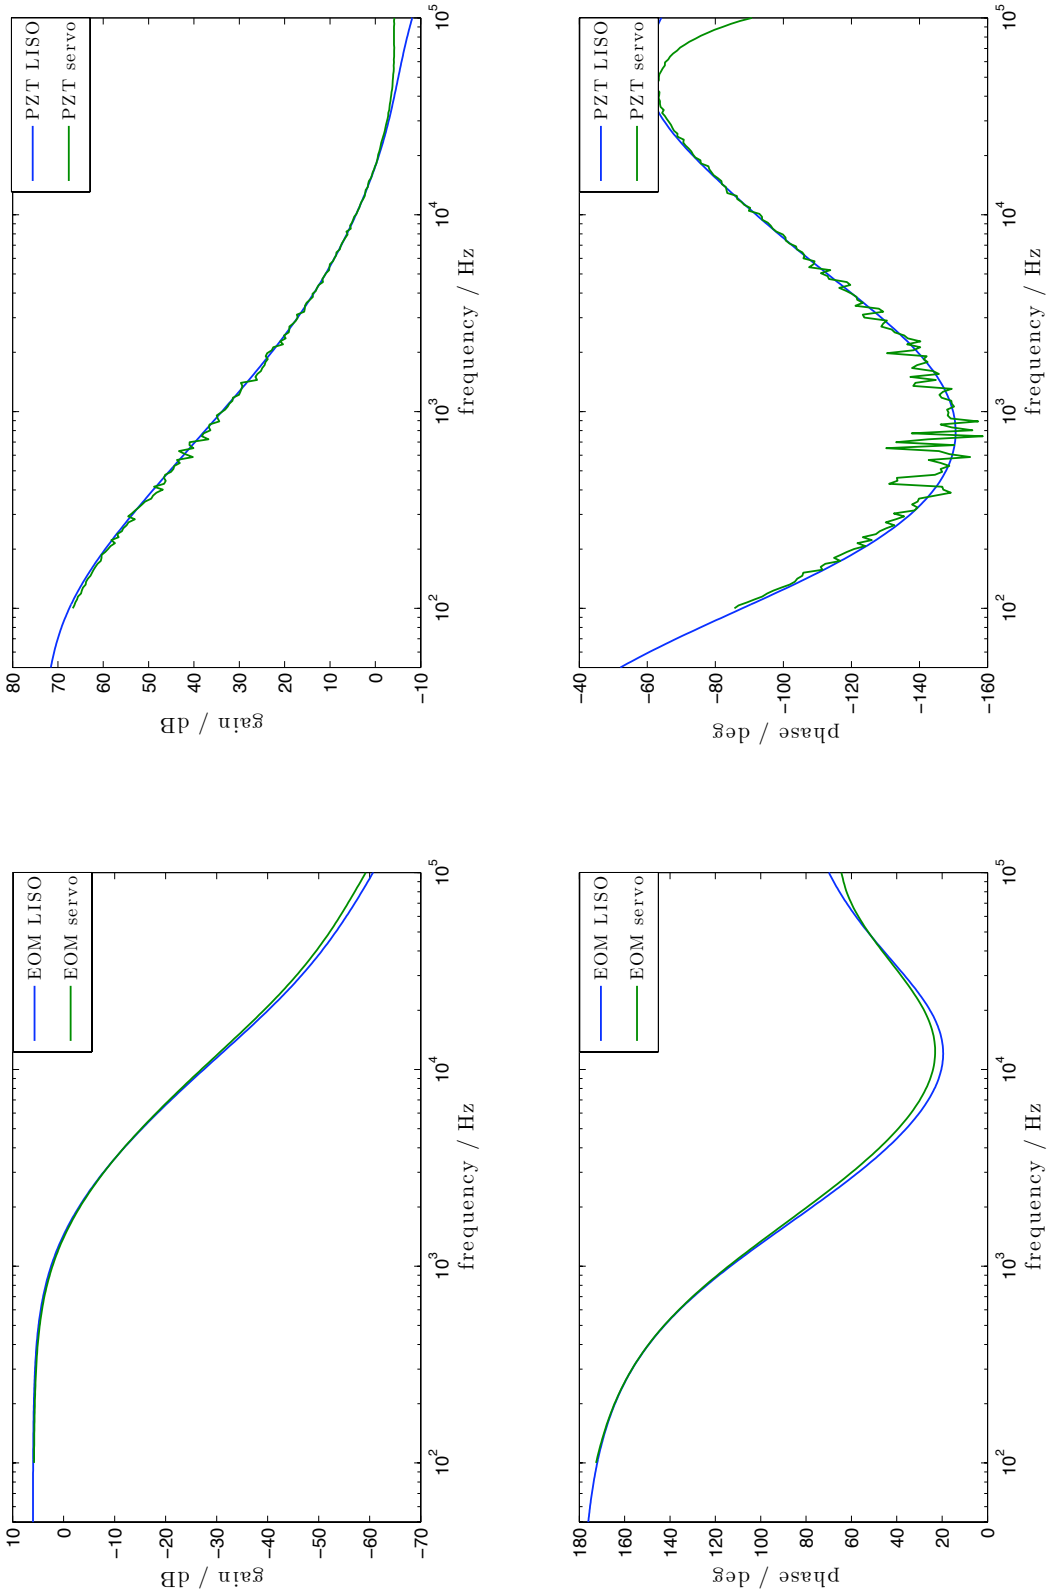

Figure 4.17: Gain and phase responses of the frequency stabilisation servo for both the EOM and PZT paths. The measured responses are compared to those from the circuit modelling software LISO [69].

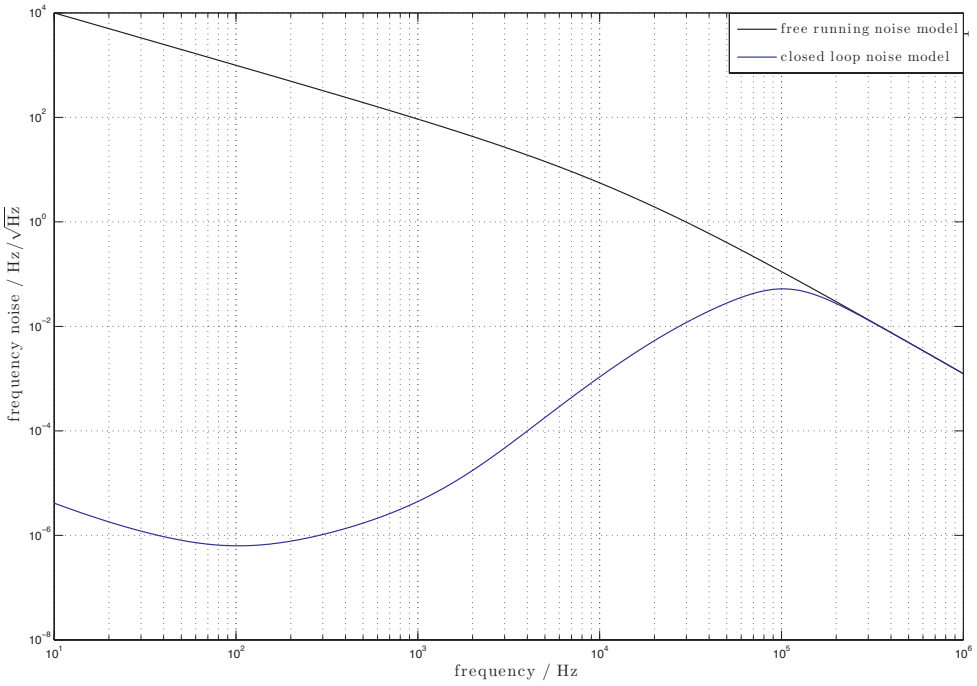

Figure 4.18: Modelled closed-loop frequency noise of the laser, when the arm cavity is locked with the frequency stabilisation servo. Note that although this estimate includes the frequency response of the cavity (an additional f^{-2} factor above $\sim 12.35\,\mathrm{kHz}$), it neglects any gain accrued from the cavity and/or the sensing scheme.

frequency chosen simply to lie outside the cavity bandwidth) were imposed by an EOM on the optical bench. The return light from the thermal noise cavity was ejected through a viewport in the vacuum tank, and a demodulated error signal produced by a TFM-3H mixer, local oscillator from an Agilent 33120A signal generator, and a resonant photodiode (see *Appendix B.3*) setup on a small optical bench, attached to the tank. This signal (along with a DC output from the photodiode) was then passed to a differential send/receive chain (to minimise electronic pick-up – see *Appendix B.4*), and into a dSPACE controller box, situated several metres away.

The dSPACE system comprised an analogue-to-digital conversion, a programmable digital filter algorithm, and a digital-to-analogue conversion on exiting the box. The flexible digital filtering of the dSPACE system was chosen to allow different servo models to be prototyped and tested quickly and easily. As the presence of several strong violin modes in the measurement band was anticipated, a relatively complex servo design could be expected and as such, the use of a flexible and easily adaptable feedback system was advantageous[6]. Sequences of filters could be defined in a simple MATLAB® script, before being compiled and uploaded. It also featured a graphical control environment, where the input and output signals could be viewed remotely. Additionally, it was straightforward to arrange triggered and timed events, e.g. the servo could be set to activate when the DC light had dropped below a certain threshold, or extra integrators at low frequency could be set to activate after several seconds of lock, to increase the gain at low frequencies.

As the dSPACE system operates digitally, it requires conversion to and from the digital domain. As such, *whitening* (and corresponding *de-whitening*) filters were used to combat the effects of quantisation noise – see *Appendix B.5*.

On exiting the dSPACE box, the feedback signal was *de-whitened* and then passed through another differential sending chain to a coil-driver box (see *Appendix B.2*) which then provided the feedback signal to the cavity by inducing currents in the coils mounted on the reaction mass, behind the ETM_{TNE}.

[6]the dSPACE controller was written by Dr Borja Sorazu.

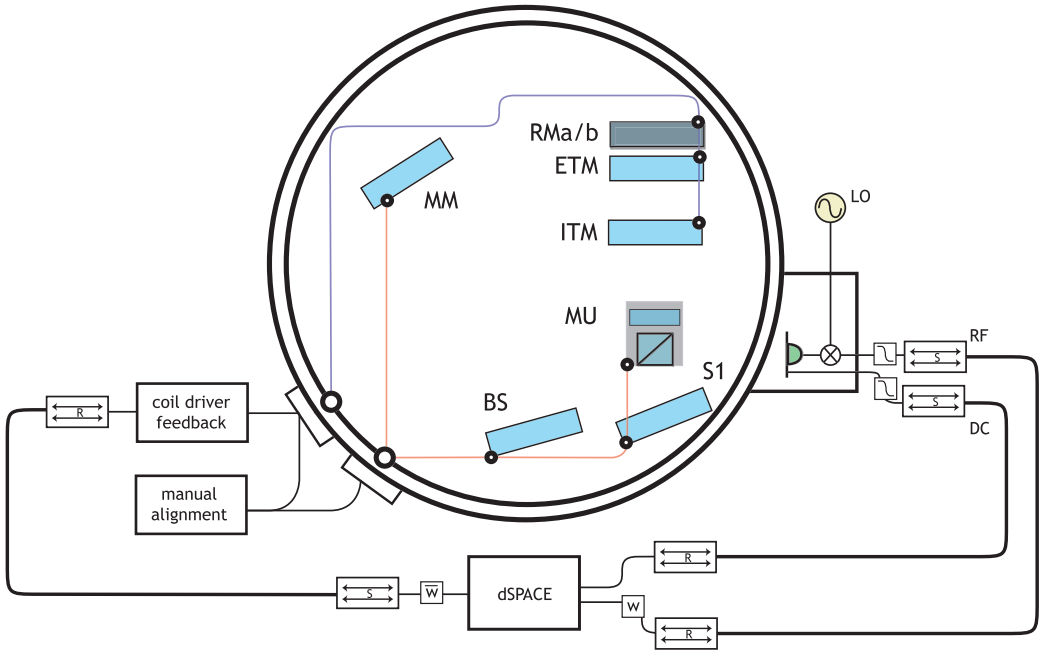

Figure 4.19: Control schematic for the thermal noise measurement tank. A signal is detected on a resonant photodiode before being demodulated and sent through a differential sending chain to a digital filter. On exiting this filter the feedback signal is sent to a coil-driver unit, which drives a current through the feedback coils mounted on the *RM*.

Alignment of the optics in the tank was achieved using manually tunable potentiometers to control the current in the alignment coils, mounted behind the upper mass stages – see *Figure 4.16*. A diagram for this circuit can be see in *Figure B.1*.

4.6 Commissioning & Results

4.6.1 Miscellaneous Measurements

Using the laser frequency stabilisation scheme as discussed in *Section 4.4*, and evacuating the interferometer system to a pressure of $\sim 10^{-5}$ torr, it was possible to lock the main arm cavity – for several hours at a time – with a fringe visibility (defined in *Section 3.2.3*) of ~ 0.13.

Optimised mode-matching was achieved through the adjustment of two converging lenses (mounted on translational stages) on the optical bench, and the measurement of the beam size using a commercial beam profiler. For the arm cavity, the fraction of the

incident power matched into the cavity was determined to be $M = 0.94$. Similarly, the mode-matching into the thermal noise cavity was performed by adjusting the position of the MM suspension optic and measuring the beam size and position, and was determined to be $M_{TNE} = 0.95$. Note that this measurement was performed before the silica optics were in place, and the beam waist position was optimised to where the reflective face of the ITM_{TNE} would be once installed.

Additionally, the modulation indices of the sidebands were determined by examining the output of a scanning Fabry-Perot cavity on the optical bench – see *Figure 4.20*. Comparing the relative power in the sidebands to that in the carrier allowed the modulation indices to be measured as $m_{46.526} = 0.73$ and $m_{10} = 0.67$ respectively for the arm cavity and thermal noise cavity control sidebands.

The power levels incident on the photodiodes, and the main cavity were measured using a commercial laser power meter, and found to be $P_{TNE} \simeq 70\,\mathrm{mW}$, $P_{AC} \simeq 50\,\mathrm{mW}$ and $P_{\mathrm{in}} \simeq 800\,\mathrm{mW}$ for the power incident on the *TNE* photodiode, the *AC* photodiode and the input power to the system respectively. All of these parameters are required to estimate the sensitivity of the measurement to the desired thermal noise fluctuation, as discussed in *Chapter 3*.

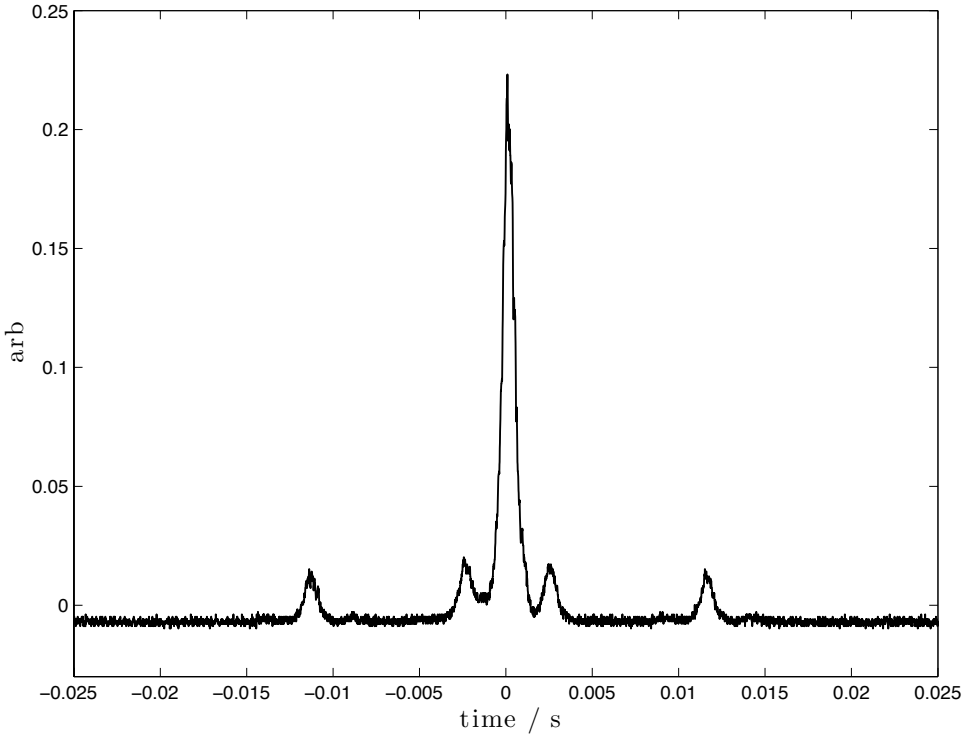

Figure 4.20: A scanning cavity measurement of the frequency spectrum of the light. The carrier, 10 MHz and 46.526 MHz sideband fields are clearly visible. The respective modulation indices are $m_{46.526} = 0.73$ and $m_{10} = 0.67$.

4.6.2 Test Cavity Operation

Initial locking of the thermal noise test cavity was achieved using a simple analogue servo – a differentiator and integrator setup giving a linear frequency response between 10 Hz and 1 kHz. This gain profile, when combined with the frequency response of the suspended optic, gives a $\frac{1}{f}$ slope over a wide frequency range. This allowed the cavity to lock for $\sim$ 10 s. The response of this prototype servo was then easily duplicated digitally using the dSPACE system.

Using the locking scheme and dSPACE system described, it was possible to lock the thermal noise cavity for tens of minutes with a fringe visibility of ~ 0.1.

With both cavities locking reliably, it was then possible to perform a displacement noise measurement.

The open and closed-loop transfer functions of the thermal noise cavity were first characterised by injecting a swept-sine signal into the alternative input of the first send/receive stage, and measuring at the alternative output of the second stage. These results (see *Figure 4.21* and *Figure 4.22*) show that the system had a unity gain frequency of $\sim$ 400 Hz, but that it exhibited significant noise below 100 Hz: specifically at 30 Hz and 100 Hz.

Similarly, by injecting a swept-sine source into the error-point input of the frequency stabilisation servo, a transfer function of the reference cavity was also measured – see *Figure 4.23*. Although this followed the approximate shape of the modelled response, the overall gain was attenuated – reaching the unity gain frequency at $\sim$ 8 kHz.

A calibrated measurement of the displacement sensitvity of the thermal noise cavity was then performed.

This was achieved by injecting a swept-sine signal at the point **a** in *Figure 4.24*.

By taking the ratio of the signal spectra obtained at points **c** and **b** it was possible to construct a calibration curve (Γ); essentially characterising the response of the thermal noise error-point to signals injected onto the PZT at each frequency. The calibrated

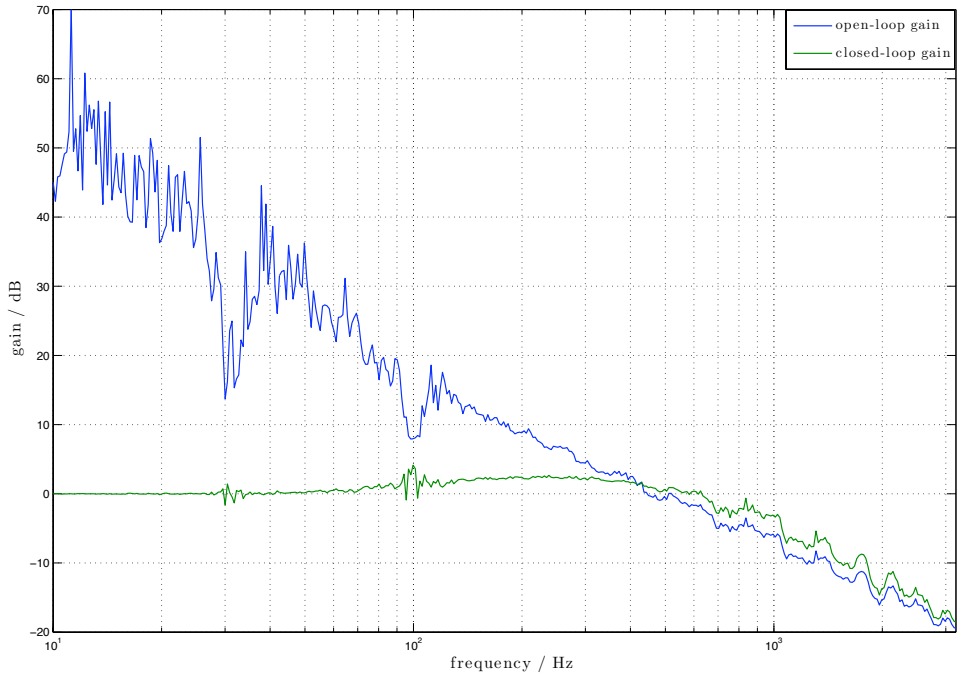

Figure 4.21: Open and closed-loop gain transfer functions of the locked thermal noise test cavity.

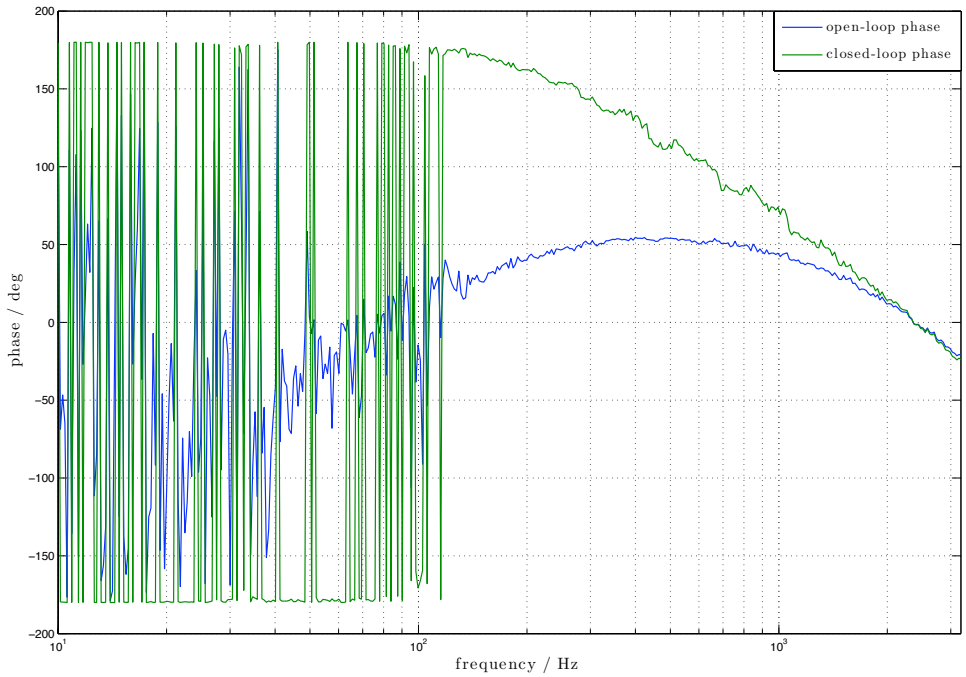

Figure 4.22: Open and closed-loop phase transfer functions of the locked thermal noise test cavity.

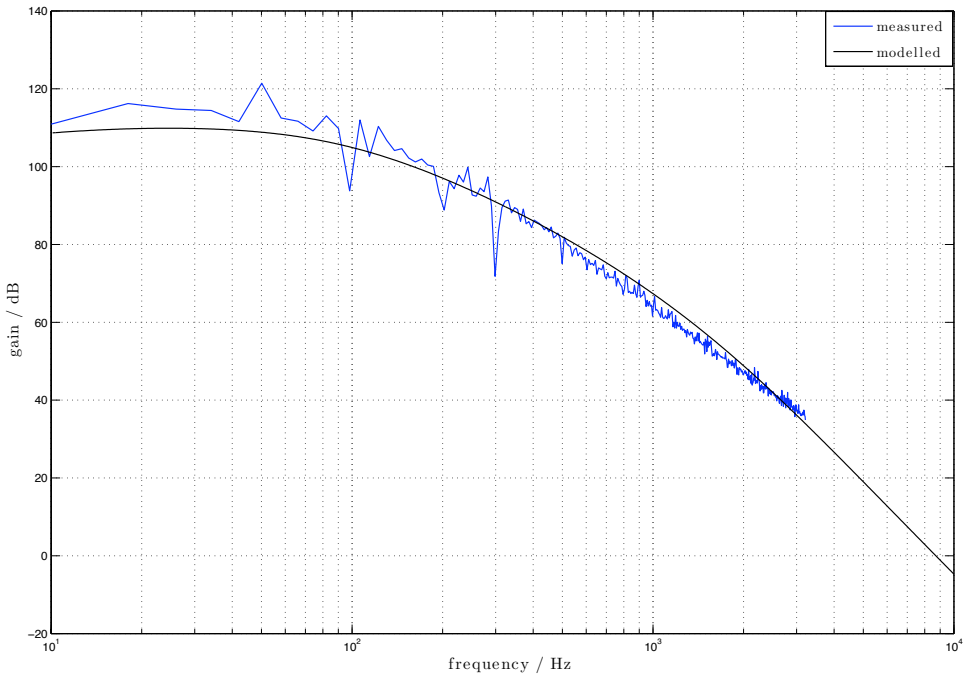

Figure 4.23: Measured and modelled open-loop response of the *AC*. The gain of the modelled response has been scaled appropriately.

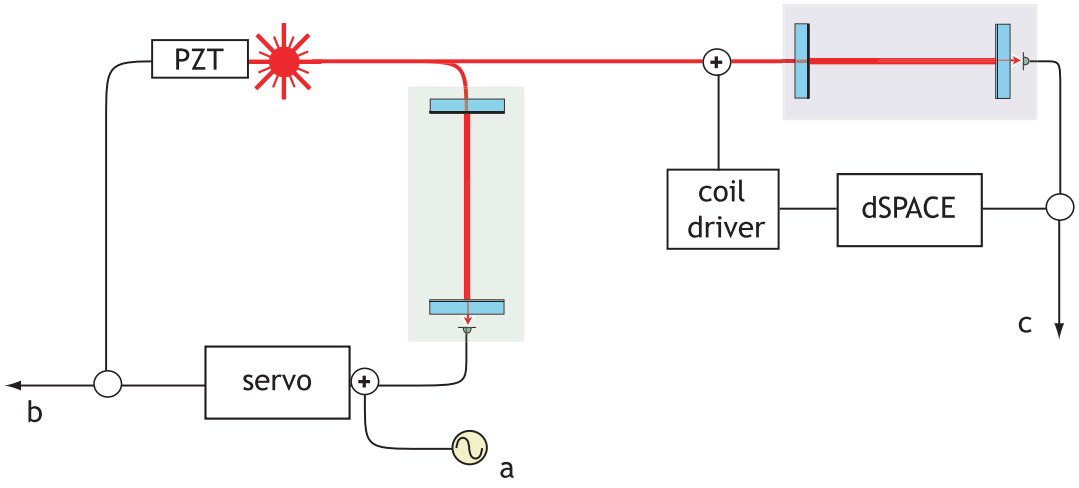

Figure 4.24: Schematic of the calibration scheme for the thermal noise cavity displacement spectra measurement. Signals could be injected into the frequency stabilisation servo for the main arm cavity, and the servo output and thermal noise error-point compared to give a calibrated displacement sensitivity for the thermal noise experiment.

displacement spectral density for any measurement was then,

$$\Delta L = \frac{L_2}{f}\Gamma\gamma\Pi\xi_{\mathrm{c}}, \tag{4.3}$$

where γ is the ratio between two corresponding calibration peaks (providing a scaling factor for the calibration curve Γ for each measurement), Π is the PZT response in $\mathrm{V}/\sqrt{\mathrm{Hz}}$, L_2 is the length of the test cavity and f the frequency of the laser and ξ_{c} is the measured error-point of the *TNE* system.

Using this method, a calibrated displacement spectra was obtained for the thermal noise cavity, along with an equivalent noise spectrum due to the excess laser frequency noise – or more accurately, the noise in the PZT path measured directly at point **b**. These are plotted in *Figure 4.25* along with measurements of the analyser noise, and the "dark noise" present when the electronics were active, but with no operational feedback applied. Although the measurements were seen to be above both analyser noise and electronic dark noise, the spurious peaks in the dark noise spectrum indicate the presence of electronic interference. The low unity gain frequency of the laser stabilisation scheme also implies that the goal for the stabilisation of the laser frequency may not have been achieved at the time of the measurement. Additionally, it was subsequently found that the power supply for the laser was faulty, causing excess noise at 100 Hz and multiples thereof, although this is not thought to have limited these measurements.

4.7 Conclusions

A suitable facility for the measurement of test-mass mirror and coating thermal noise was constructed in the IGR JIF facility. All the necessary systems required to appropriately isolate the measurement optics were successfully installed, included two quasi-monolithic fused silica suspensions – see *Figure 4.26*. Initial commissioning was performed, and while both optical cavities were capable of locking for extended periods, the performance of the system was seen to be limited by unknown electronic noise

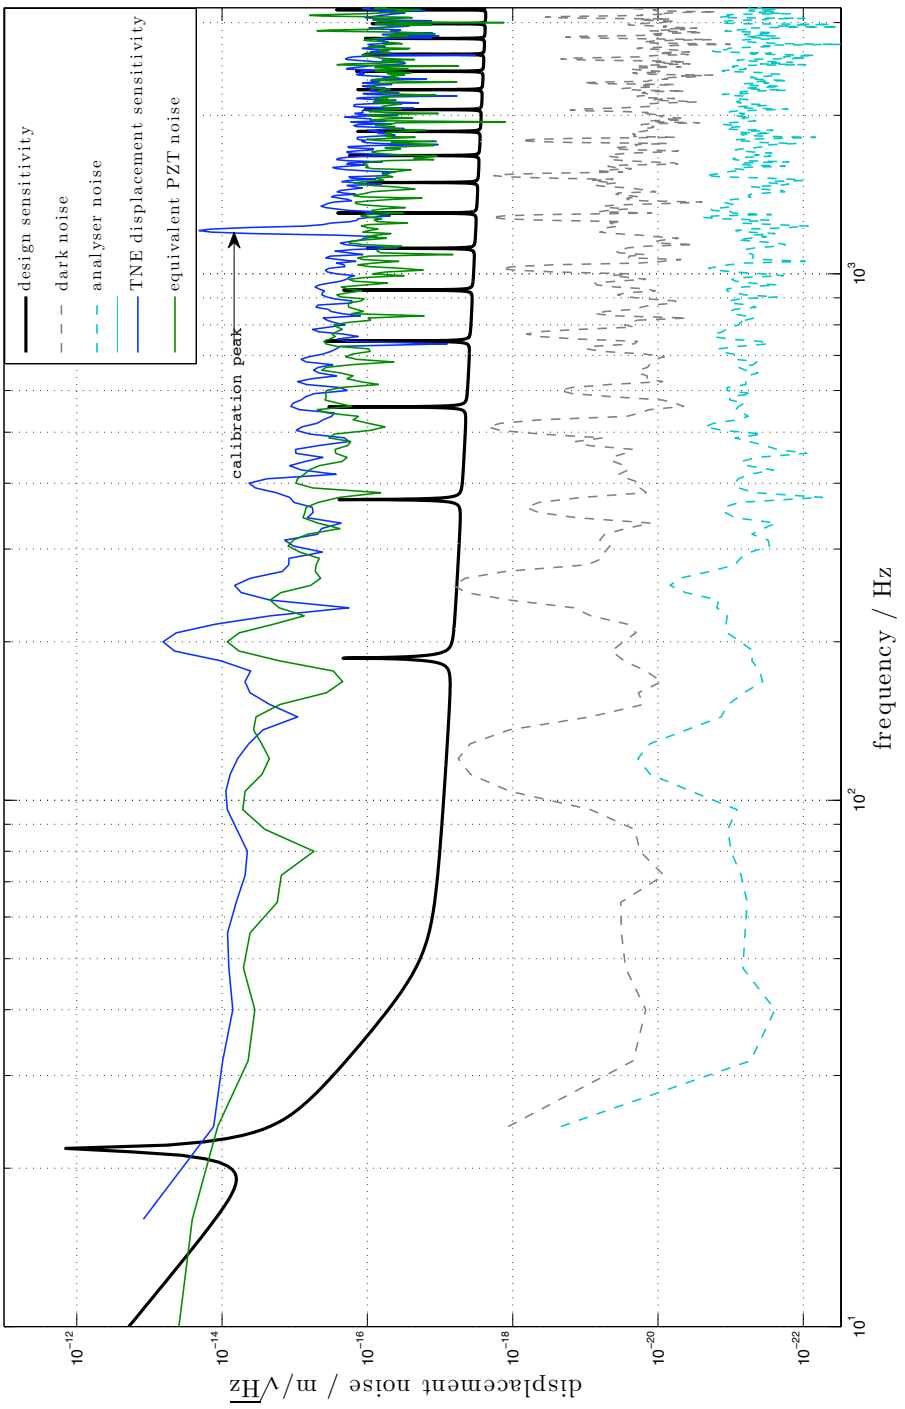

Figure 4.25: Measured calibrated displacement spectrum of the thermal noise cavity, along with measurements of the PZT noise, analyser noise, dark noise, and the design sensitivity of the system, obtained from measured parameters.

coupling into the laser frequency feedback path.

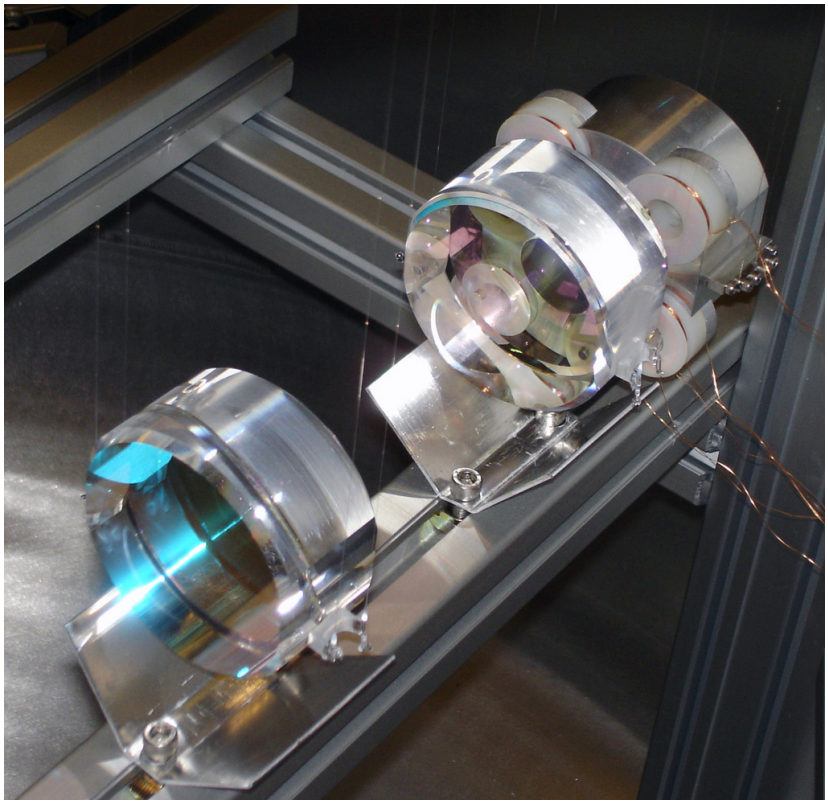

Figure 4.26: The lower stages of the thermal noise measurement cavity, suspended in the tank. The reaction mass can be seen behind the ITM_{TNE}.

However, this should not present an insurmountable difficulty – an improvement in the laser frequency stabilisation, eradication of all spurious electronic noise and proper tailoring of the feedback loop via the dSPACE system should provide a commensurate improvement in the sensitivity of the thermal noise experiment, allowing coating and substrate thermal noise effects to be directly measured.

One future experiment already planned is to evaluate the effect of silicate bonding on the thermal noise spectrum of an optic – see *Figure 4.27*. As future detectors may employ test-mass mirrors constructed as a composite of smaller masses, silicate bonded together, this as yet unmeasured potential noise source will require investigation.

Figure 4.27: A replacement thermal noise experiment test optic. Here the optic is shown in two pieces, prior to silicate bonding. This will create a composite mass with the lossy bond layer close to the measurement beam, allowing the effect of the bonding process on the thermal noise spectrum to be directly quantified.

Chapter 5

Control of a Three-Mirror Coupled Cavity

5.1 Introduction & Motivation

Any interferometric topology formed from more than two partially transmissive and reflective optics can be considered a *coupled cavity* system, where the behaviour of one part of the system can influence the state of another. Gravitational wave interferometers are typically comprised of numerous coupled cavities (depending on the specific detector topology); where the distinct cavities are coupled through the laser light, the nature of the resonant condition in each cavity, and (potentially) the opto-mechanical effect of the light's radiation pressure on the nominally free mirror masses.

An optical cavity (or combinations thereof) in a detector often represents an important degree of freedom that gives – or combines to form – the signal output of the detector. As any effective detector must be held at its operating point, all of the

degrees of freedom within the coupled optical system must be controlled.

In *Section 2.3.1* the concept of using a single set of phase modulation sidebands to control a single optical cavity was introduced, and subsequently implemented in the experiments detailed in *Chapter 3* and *Chapter 4*.

Current, and planned second generation gravitational wave interferometers, use techniques such as *multiple* sets of phase modulation sidebands to effectively control all relevant length degrees of freedom [70][1]. Methods such as *gain hierarchy* are employed, where inter-cavity coupling is suppressed by the relative gains of the control servos. However, such systems can require extensive effort to be expended in the control system design, in order to de-couple the various length sensing signals.

Additionally, as the circulating light power increases in future detectors ($\sim 1\,\mathrm{MW}$ in Advanced LIGO), the need to de-couple the control signals of different parts of the interferometer becomes more critical. An increase in circulating power brings with it an increase in the radiation pressure exerted on the mirrors – pushing them away from their resonant rest positions – resulting in a rigid *optical spring* effect that can dominate the mirror motion and complicate any control system. These optical spring effects have already been seen in large-scale interferometers [71].

The goal of any de-coupling scheme is to diagonalise the relevant *control matrix*. This matrix represents the degree to which the sensing schemes for the separate degrees of freedom are sensitive to each other. A representation of a control matrix for two generic coupled degrees of freedom (Φ_A and Φ_B) and their respective error signals (A and B) is shown in *Table 5.1*.

Ideally, the (normalised) diagonal components will equal unity, whereas the off-diagonal components will be zero – signifying no sensitivity to the other degrees of freedom, and an orthoganalised control scheme. In the case of a gravitational wave

[1] Angular degrees of freedom (the mutual alignment between the mirror faces) are typically controlled with differential wavefront sensing techniques.

	Φ_A	Φ_B
A	$\frac{dA}{d\Phi_A}$	$\frac{dA}{d\Phi_B}$
B	$\frac{dB}{d\Phi_A}$	$\frac{dB}{d\Phi_B}$

Table 5.1: A generic control matrix for two variables Φ_A and Φ_B and their respective error signals – A and B.

interferometer, the components Φ_A and Φ_B would typically represent an optical phase or physical cavity length, with A and B then being the demodulated error signals used to hold these quantities close to their respective operating points.

To test control systems and de-coupling schemes, it is desirable to set up a simplified system as a proof of concept. The simplest relevant coupled optical system consists of three freely suspended mirrors, forming a single interferometer arm coupled to a recycling cavity – see *Figure 5.1* – hereafter termed a *three-mirror coupled cavity*.

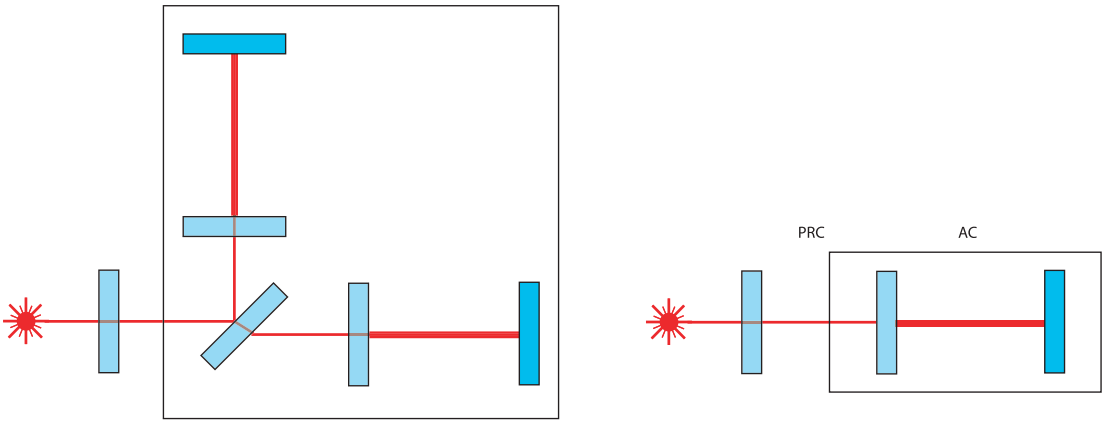

Figure 5.1: A full-scale gravitational wave interferometer consists of a sequence of coupled optical cavities. To simplify the study of appropriate control schemes, a simplified three-mirror coupled cavity system was used; effectively representing a single interferometer arm (AC) with a power recycling cavity (PRC).

5.2 Experimental Layout

Utilising the JIF interferometer prototype in its power recycling configuration (with a reflective optic in the *PRM* position, in place of the blank optic used for the experiments in *Chapter 4*), a three-mirror coupled cavity system was formed comprising the power recycling mirror (*PRM*), inboard test-mass (*ITM*) and end test-mass (*ETM*)[2]: recall *Figure 3.3* and *Table 3.1* for an overview of the system's layout and optical properties, and see *Figure 5.2* for a schematic of the experimental setup.

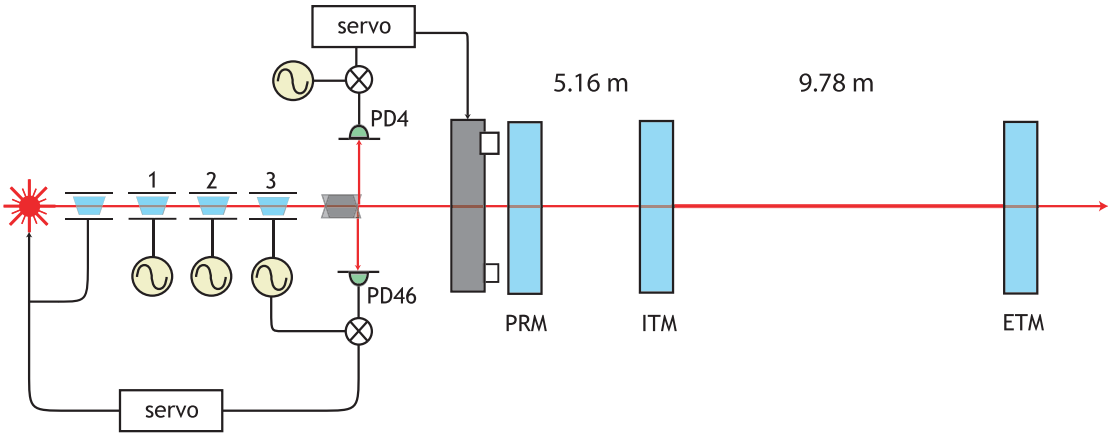

Figure 5.2: The three-mirror coupled cavity experiment setup. A short (5.16 m) power recycling cavity and a long 9.78 m arm cavity were formed by three suspended mirrors. One electro-optic modulator (EOM) forms part of the laser frequency stabilisation scheme, while three more provide modulation sidebands for the length sensing scheme.

Three electro-optic modulators generated modulation sidebands for the control of the system: EOM_1 (with reference to *Figure 5.2*) generated modulation components at 14.525 MHz, EOM_2 at 10 MHz and EOM_3 at 46.526 MHz respectively. The control scheme and choice of frequencies is detailed further in *Section 5.3*.

The return light from the three-mirror coupled cavity was separated by a Faraday isolator in the beam path, and directed towards two resonant InGaAs photodiodes, similar to the one detailed in *Appendix B.3*. In conjunction with TFM-3H mixers, local oscillator signals and appropriate low-pass filtering, these photodiodes provided the demodulated error signals for both constituent cavities. The 46.526 MHz modu-

[2]Additionally, two steering mirrors were present within the *PRC* to direct the light around the folded interferometer arm.

lation frequency was provided by an Agilent 8648A signal generator, while the other frequencies were generated by Agilent 33120A models.

The *AC* was controlled through feedback to the laser frequency over three paths, similarly to the method described in *Chapter 3* and *Chapter 4*. Here however, the low frequency feedback was governed by the temperature control of the laser crystal, controlled through a low frequency integrator circuit, acting below ~ 0.6 Hz. Additionally, some features in the electronic control servo for the laser (see *Appendix B.6*) were switched to allow for the higher effective finesse (and therefore lower cavity bandwidth) of the coupled system – see *Figure 5.3* and *Figure 5.4* for the optical transfer functions of the standalone arm cavity and the three-mirror coupled system respectively. The error signal for the control of this cavity was derived using the standard *PDH* method of imposing modulation sidebands (at 46.526 MHz) and using a demodulated signal as the error-point, as detailed in *Chapter 4*.

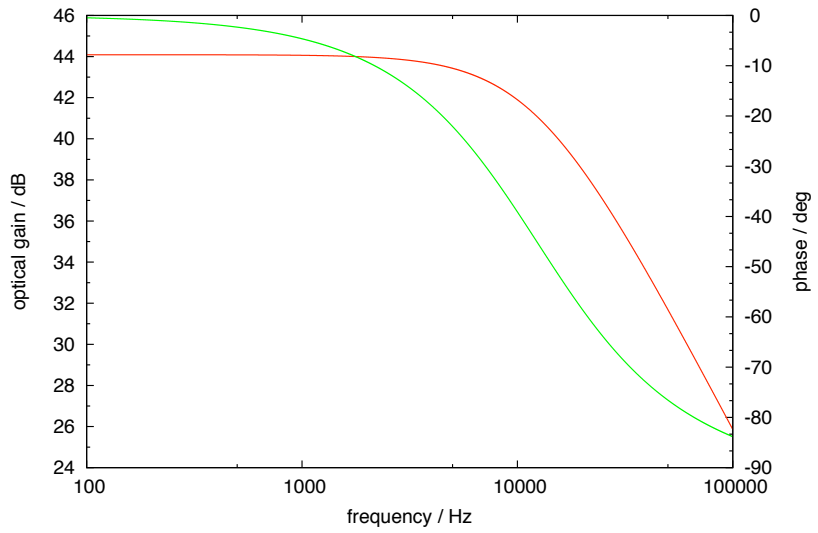

Figure 5.3: Optical transfer function for the arm cavity: optical gain is shown in red, and phase in green.

The power recycling cavity was controlled by feeding back to the position of the *PRM*, through a coil/magnet actuator system designed to operate up to a unity gain frequency of ~ 500 Hz. The error signal for this feedback path was derived from the demodulation of the beat between the 10 MHz sidebands and the 14.525 MHz sidebands,

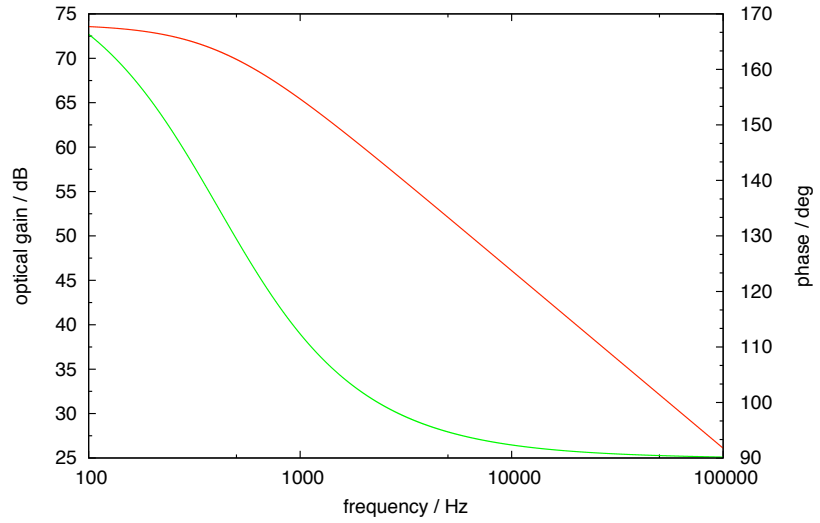

Figure 5.4: Optical transfer function for the three-mirror coupled cavity: optical gain is shown in red, and phase in green.

with the local oscillator signal for this process provided by a separate signal generator (phase-locked to the 10 MHz oscillator) operating at the beat frequency – 4.525 MHz. The control scheme and choice of frequencies is covered in more detail in *Section 5.3*.

5.3 Control Scheme

5.3.1 Overview

As gravitational wave interferometers become more complex, it becomes necessary to increase the complexity of their control systems. As single cavities can be controlled by *PDH* sensing techniques, it is natural to extend these methods when dealing with multiple cavities; employing perhaps separate modulation frequencies for each cavity to be controlled.

For a single cavity, the error signal is constructed from the interference between the carrier light – which exhibits a small relative phase shift on returning from a cavity that has undergone a disturbance – and a set of static (with respect to the carrier) phase modulation sidebands: see *Section 2.3.1*. A suitable set of signals for controlling a coupled cavity system can be constructed from a superposition of any set of such

"static" and "sensitive" components. Coupling effects arise when two or more separate length sensing schemes share common sensing components.

To *de-couple* the sensing and control scheme for the three-mirror coupled cavity, the following system was employed:

An amplitude modulated sideband field (at 14.525 MHz) was chosen to be resonant in the power recycling cavity alone, thus removing any common dependence on the carrier light phase. The *PRC* length control signal was then derived from the interference of this field with a set of 10 MHz phase modulated sidebands, that were resonant in neither the arm or power recycling cavities.

Using a phase modulated sideband field at the same frequency generated an unsuitable control signal, offset from zero when far from a resonance, and without the desired bi-polarity around the operating point – see *Figure 5.5*.

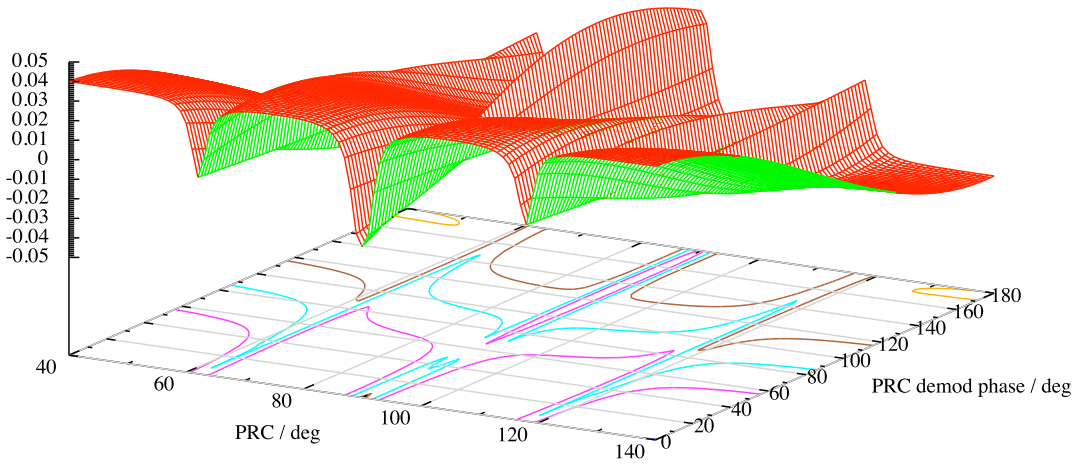

Figure 5.5: Modelled *PRC* error signal using phase modulated sidebands at 14.525 MHz. A minimum, rather than a zero-crossing, appears at the operating point.

The 14.525 MHz amplitude modulated component could be adjusted to be either a pair of amplitude modulation sidebands, or a single (upper or lower) sideband where appropriate, using a novel modulation setup – detailed in *Section 5.4*.

The positions of the key modulation frequencies with respect to the cavity free spectral ranges, and each other, can be seen in *Figure 5.6*. The central carrier compo-

nent is seen to be resonant in the arm cavity, while its control sidebands lie outside the linewidths of both cavities. Simultaneously, the 14.525 MHz field is seen to be resonant in the power recycling cavity.

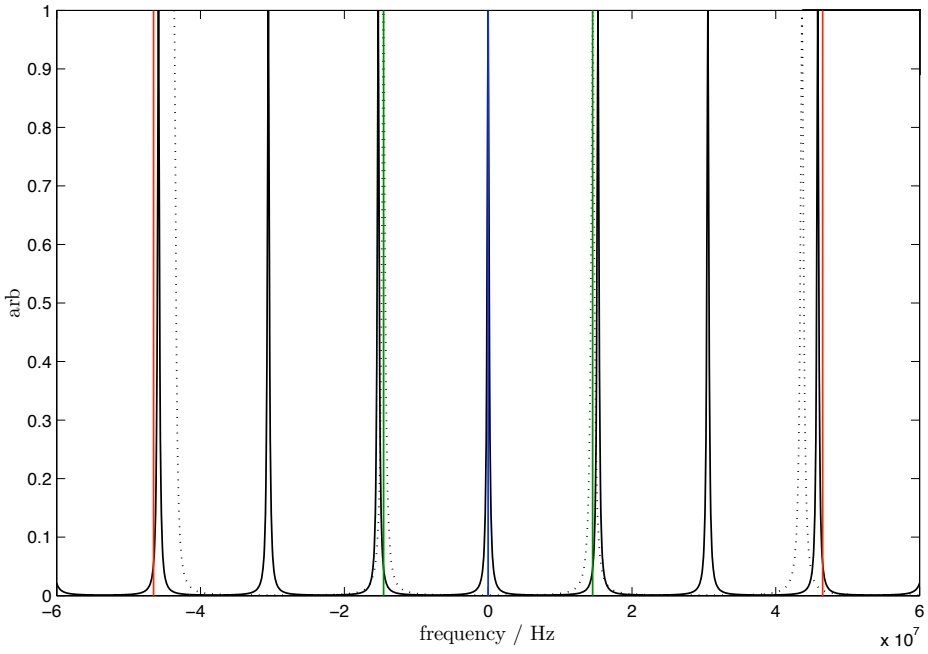

Figure 5.6: The resonant features of the three-mirror coupled cavity experiment. The free spectral ranges of the main arm cavity (solid black) and the power recylcing cavity (dashed black) are shown with the carrier position shown in blue, 14.525 MHz sidebands in green and 46.526 MHz sidebands in red. The carrier is seen to be resonant in the AC while being anti-resonant in the PRC – in which the amplitude modulated components *are* resonant.

5.3.2 Modulation Frequencies

In addition to the restrictions imposed by the positions of the respective cavity resonances, the modulation frequencies were required to fulfill certain criteria. The demodulated error signals were required to be zero at the desired condition (when the carrier was resonant in the arm cavity), and the frequencies were required to be kept relatively low (≤ 100 MHz) to ease the design and construction of the associated electronics. They were also required to not be direct multiples of each other, to minimise any adverse coupling between signal harmonics.

A model of the system was written using the frequency domain interferometer sim-

ulation software FINESSE [53], which is capable of simulating error signals, transfer functions and light powers for arbitrary combinations of typical optical and electrical interferometer components. Using the FINESSE model of the system, it was clear that any modulation frequency suitably removed from the carrier resonance – or a cavity free spectral range – would provide a suitable demodulated control signal for the arm cavity.

Initially the modelling was performed using macroscopic cavity lengths that later proved incorrect, namely 5.35 m and 9.67 m for the *PRC* and *AC* respectively. Using these lengths the modulation frequency 46.526 MHz was chosen for the arm cavity control, as it fulfills the criteria outlined above, and it was initially thought that this frequency would provide an optimal control signal gradient. The modelled demodulated error signal around this frequency region is shown in *Figure 5.7*.[3] [4]

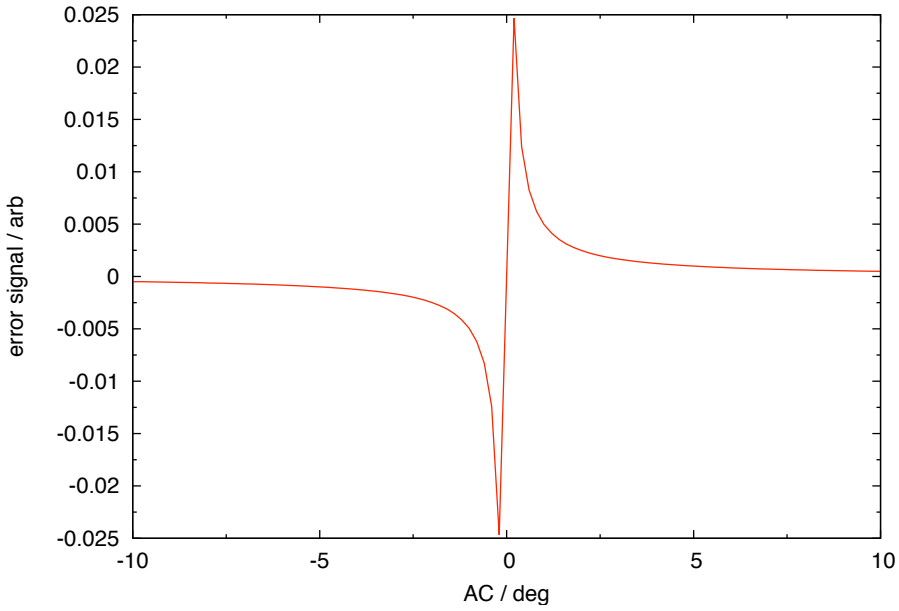

Figure 5.7: Modelled error signal for the arm cavity.

Similarly, the modulation frequency 13.937 MHz was initially chosen to provide the control signal for the power recycling cavity, in conjunction with the 10 MHz modulation field, whose frequency was chosen effectively arbitrarily on the condition that it

[3] The error signal is shown against the cavity tuning in degrees, which is equivalent to $\delta l \left(\frac{360}{\lambda}\right)$, where δl is a change in the cavity length and λ is the laser wavelength.

[4] Modulation indices of $m = 0.3$ (for all sidebands) are used for these simulations.

was non-resonant in either cavity.

The modulation frequency for the resonant field in the power recycling cavity is equivalent to half the cavity's free spectral range, as a consequence of the need to have the carrier simultaneously anti-resonant. A simulated error signal for the power recycling cavity can be seen in *Figure 5.8*.

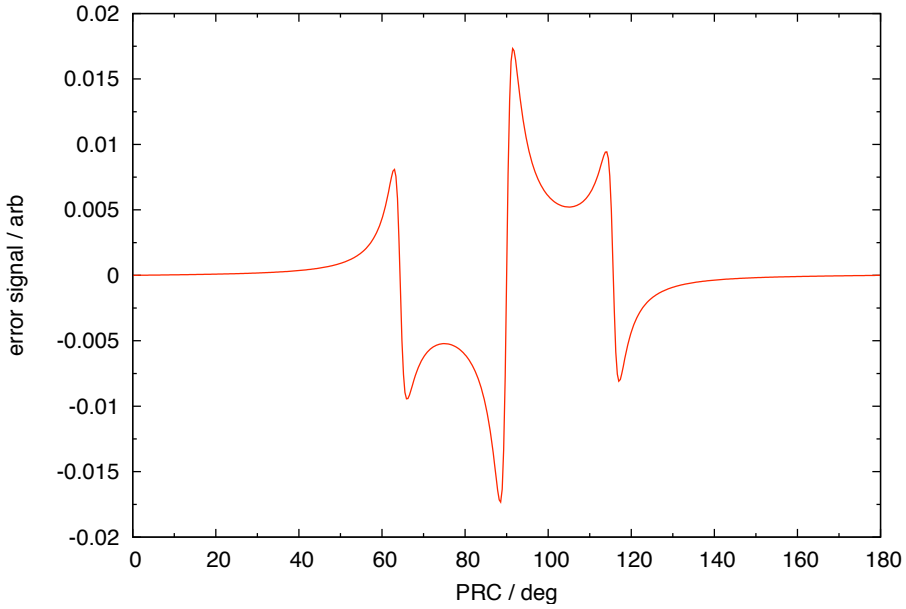

Figure 5.8: Modelled error signal for the power recycling cavity.

On correcting the erroneous lengths, subsequent remodelling indicated that the arm cavity control signal was relatively insensitive to changes in the arm length – see *Figure 5.9*. As such, the modulation frequency was not changed.

The gradient of the control signal for the power recycling cavity however was seen to vary significantly with the length of the cavity – see *Figure 5.10*. To compensate for this, the modulation frequency was changed to 14.525 MHz, and the resonant circuit that coupled electrical power to the electro-optic modulator (see *Appendix B.7*) was retuned accordingly.

The different length dependencies of the two signals can be understood simply by considering that the arm cavity control sidebands need only lie *outside* the arm cavity linewidth, whereas the sideband field that provides the control signal for the power

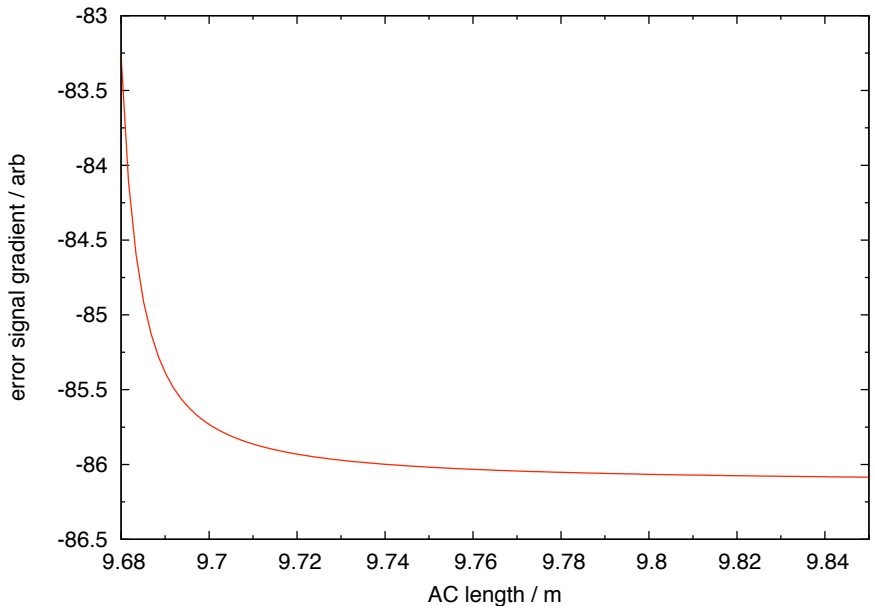

Figure 5.9: The gradient of the *AC* error signal as a function of the cavity length, keeping the modulation frequency constant at 46.526 MHz. Large changes in the cavity length have relatively little effect on the gradient of the error signal.

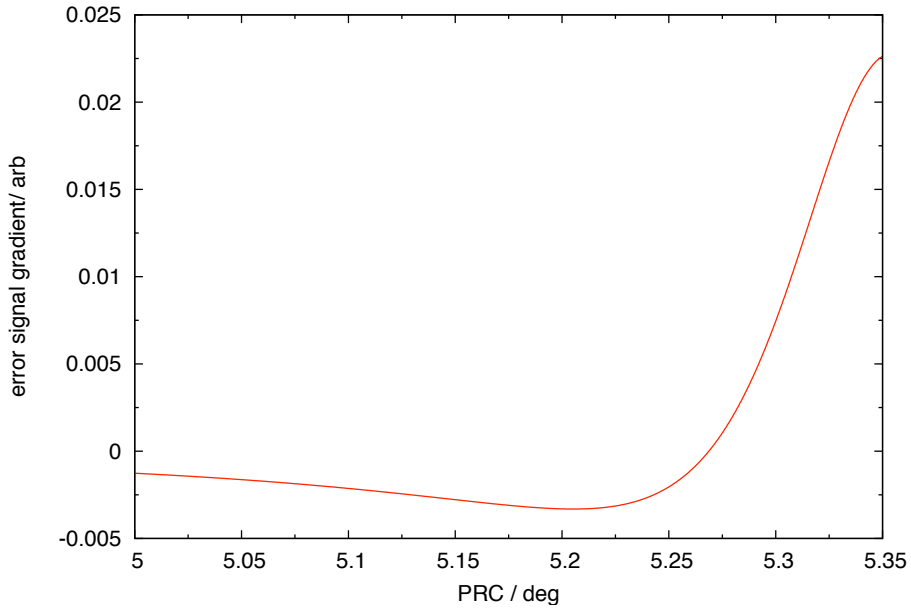

Figure 5.10: The gradient of the *PRC* error signal as a function of the cavity length, keeping the modulation frequency constant at 13.937 MHz.

recycling cavity must lie *within* the *PRC* linewidth. Thus any change in macroscopic cavity length (and hence free spectral range) will have a greater impact on the power recycling cavity signal.

Using the final, corrected, lengths and modulation frequencies for the system, it was possible to construct a graphical representation of the de-coupled control matrix, by plotting the respective error signals against simulated motion in either part of the coupled cavity system. This can be seen in *Figure 5.11* to *Figure 5.14*; the control signals exhibit bi-polar responses to their respective degrees of freedom, and minimal responses to each other. The model control matrix is then given by the (normalised) gradients – taken at the zero crossing – of these four signals.

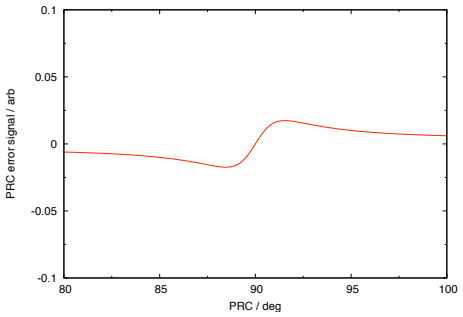

Figure 5.11: *PRC* error signal response to motion of the *PRM*.

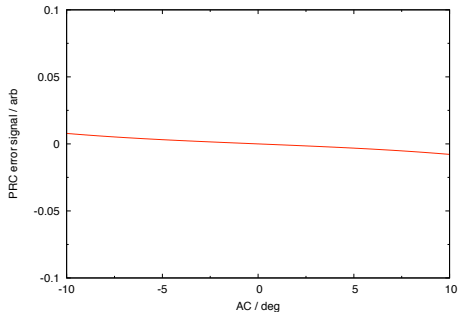

Figure 5.12: *PRC* error signal response to motion of the *ITM*.

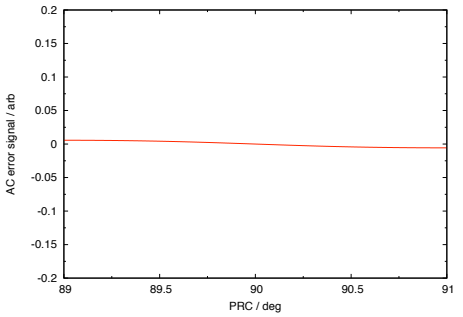

Figure 5.13: *AC* error signal response to motion of the *PRM*.

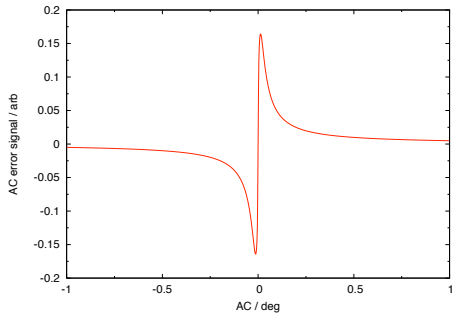

Figure 5.14: *AC* error signal response to motion of the *ETM*.

5.4 Sideband Generation

In *Section 5.3* it was shown that an amplitude modulated component would provide a suitable error signal for a de-coupled three mirror system. Amplitude modulation sidebands are typically generated by combining two phase modulated beams, requiring the use of two separate electro-optic modulators. Here, a novel, single modulator system was used, which was also capable of generating *single sideband* modulation, where either the upper or lower modulation sideband alone was present.

5.4.1 Background

The polarisation state of laser light can be passively manipulated through the use of optical *wave-plates*. These are typically bi-refringent optical elements, that exhibit two possible propagation axes with different refractive indices – termed the *ordinary* and *extraordinary* axes[5]. For a linearly polarised input, rotating the wave-plate around the optical axis alters the proportion of the incident light that is polarised along each orthogonal crystal axis.

As each component polarised along the two axes will experience a different phase retardation due to the difference in refractive indices, cutting the wave-plate to a specific thickness ensures that the light exiting the wave-plate will be composed from two light fields with a fixed relative phase lag. The thickness of the waveplate determines the phase difference accrued, and the rotation angle of the wave-plate determines the relative proportions of these two fields.

Two common forms of wave-plate are the *quarter wave-plate* and the *half wave-plate*. The former retards one propagating component (with respect to the other) by a phase equivalent to a quarter of a wavelength. At the specific orientation angle where the amount of light polarised along the ordinary and extraordinary axes are equal, the output light will be circularly polarised. If the wave-plate is aligned such that all the

[5]These are also sometimes referred to as *slow* and *fast* axes, where the lower refractive index indicates the fast axis.

light is polarised along either the ordinary or extra-ordinary axis, the output will be linearly polarised. Between these two conditions the output light will exhibit elliptical polarisation. The half wave-plate functions similarly, retarding one component by a phase equivalent to half the wavelength. On recombining at the output face of the optic, the orthogonal components combine to form a light field with a polarisation rotated with respect to the input field – equal to twice the angle of the rotation of the wave-plate.

An electro-optic phase modulator is an active electrical component (often a lithium niobate – $LiNbO_3$ – crystal), which operates by retarding the phase of the light passing through it; the retardation being proportional to an applied (often high frequency) voltage. This effect acts to retard the phase of the light polarised only in the actuation plane of the crystal, and the input light is set to be polarised along this plane before entering the modulator, thus creating phase modulation sidebands. By rotating the modulator with respect to the optical axis, the applied phase modulation sidebands are only applied to the carrier component that is polarised along the actuation axis. In general the modulator now acts as a voltage controllable wave-plate with a fixed rotation angle, and can in principle be made to mimic either a quarter wave-plate or half wave-plate by applying the correct voltage at DC. See *Figure 5.15* for a schematic of the modulator layout.

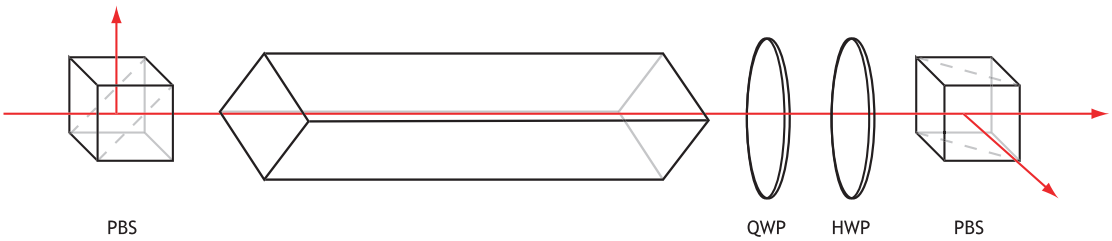

Figure 5.15: The flexible optical modulation setup. Linearly polarised light is phase modulated along one axis as it passes through a rotated phase modulator. Two wave-plates and an output polariser then allow selected components of the modulated light to pass through to the system.

5.4.2 Modulation Process

Light entering the system is first linearly polarised (if it is not already so) by a polarising beam-splitter cube, before passing through the rotated modulator. As the modulator is rotated at 45° with respect to the polarisation of the light, half of the light will travel along each axis, with one component experiencing the phase modulation imposed by an electrical input signal at 14.525 MHz.

It can be instructive to view the phase modulation process graphically in terms of phasors – see *Figure 5.16*.

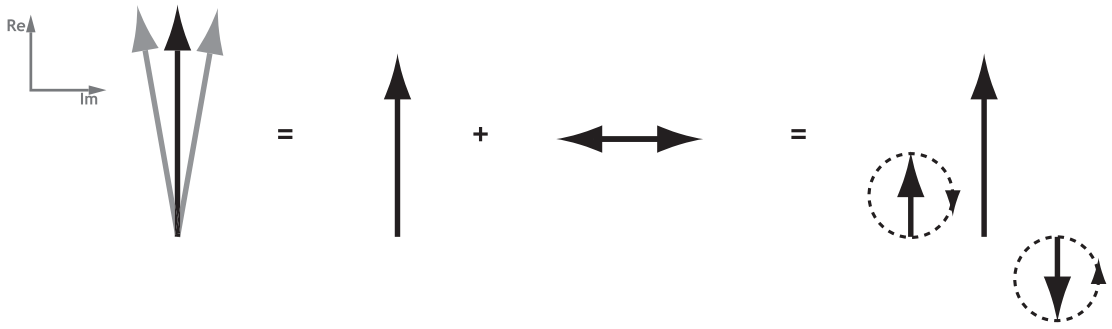

Figure 5.16: A phasor representation of phase modulation. The modulation can be decomposed into a static carrier (i.e. the axes of the Argand plane rotate at the angular frequency of the carrier light) and oscillating phasor in quadrature. This oscillating component can be further decomposed into two phasors with opposing phases, rotating in opposite directions.

Here the input light is depicted as a vector in the Argand plane with a length proportional to the amplitude of the light, where the phase is then given by the angle of the vector with respect to the real axis. The phase modulated light can be viewed as comprising a (reduced amplitude) carrier component, and an imaginary oscillatory part, which can be further decomposed into two counter-rotating (at the rate of the modulation frequency) phasors.

At the output of the modulator, the structure of light can be depicted as in *Figure 5.17*, with the phase modulation sidebands now polarised at 45° with respect to the carrier light.

The light then passes through a quarter wave-plate. This optic can be rotated such

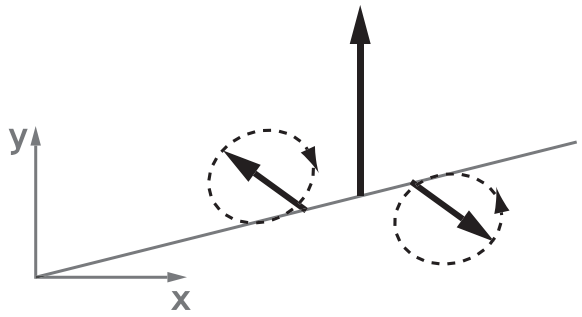

Figure 5.17: A phasor depiction of the output light from the modulator. The sideband field is polarised at 45° with respect to the carrier field.

that either the carrier becomes circularly polarised – in which case the sidebands will become linearly polarised, and orthogonal with respect to each other – or the carrier remains linearly polarised, in which case the sidebands become circularly polarised with opposing chirality. Rotation angles in-between these two states will lead to all components being elliptically polarised. In fact, as some power is inevitably lost in one of the carrier components – required to generate the sidebands – the reconstructed carrier field at the output face of the modulator will not in general be polarised exactly at 45° with respect to the sideband field, and a degree of ellipticity will typically be present in either or both fields after the quarter wave-plate.

On passing through the half wave-plate, the components can be rotated in polarisation so that they either pass through the polarising beam-splitter, or are rejected.

In summary: the modulator imposes phase modulation sidebands that are polarised at 45° to the carrier light; the quarter wave-plate induces a degree of circular polarisation to either – or both – the sidebands and carrier; the half wave-plate and output polarising beam-splitter then allow these components to be selectively passed through to the cavity system, or rejected. An unavoidable consequence of this scheme is that half of the input light is typically rejected.

5.4.3 Analysis

This process can be analysed more thoroughly using the *Jones calculus* notation for the polarised light and optics involved. In this notation, a monochromatic beam of light is represented by,

$$E_{\text{arb}} = \begin{bmatrix} E_x \\ E_y \end{bmatrix} \tag{5.1}$$

where E_{arb} is an arbitrarily polarised light field, and E_x and E_y represent the electric field component vectors along those respective axes. If the light is linearly polarised along one axis (here the x axis), it can then be represented by,

$$E_{\text{lin}x} = E_0 e^{i\omega_0 t} \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{5.2}$$

Polarising optical components can also be concisely represented in this notation. A quarter wave-plate (where the extraordinary axis is aligned along the x axis) is represented by,

$$QWP = e^{i\frac{\pi}{4}} \begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix}, \tag{5.3}$$

and a half wave-plate by

$$HWP = \begin{bmatrix} i & 0 \\ 0 & -i \end{bmatrix}. \tag{5.4}$$

In general, a voltage controllable wave-plate, can be represented by,

$$VWP = \begin{bmatrix} e^{i\alpha} & 0 \\ 0 & e^{-i\alpha} \end{bmatrix}, \tag{5.5}$$

where α is the angle of phase retardation induced by the electro-optic process.

Finally, for an optic X rotated by an angle θ around the optical axis, its resultant

Jones matrix $X(\theta)$ will be,

$$X(\theta) = R(-\theta)\, X\, R(\theta), \tag{5.6}$$

where the rotation matrix $R(\theta)$ is,

$$R(\theta) = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}; \tag{5.7}$$

and the output state of a beam (E_{out}) encountering first an optic X and then Y is given by the multiplication of the Jones matrix elements as follows:

$$E_{\text{out}} = Y\, X\, E_{\text{in}}\ . \tag{5.8}$$

In general, an oscillatory optical modulation can be expressed as,

$$E_{\text{mod}} = E_0 e^{i\omega t}\left[M_{\text{c}} + iM_{\text{u}}e^{i\omega t} + iM_{\text{l}}e^{-i\omega t}\right], \tag{5.9}$$

where M_c represents the amplitude of the carrier post-modulation, and M_{u} and M_{l} represent the amplitude of the upper and lower sidebands respectively.

The electric field after the rotated modulator – E_{mod} – is given by,

$$E_{\text{mod}} = E_{\text{in}}\, R\left(-\frac{\pi}{4}\right)\, VWP(\alpha)\, R\left(\frac{\pi}{4}\right). \tag{5.10}$$

This then leads to the following three terms for the amplitudes of the three resultant components of the light; the carrier, the upper sideband and the lower sideband,

$$M_{\text{c}} = J_0(m), \tag{5.11}$$

$$M_{\mathrm{l}} = J_1(m) R\left(\frac{-\pi}{4}\right) \begin{bmatrix} e^{-i\alpha} & 0 \\ 0 & e^{i\alpha} \end{bmatrix} R\left(\frac{\pi}{4}\right), \tag{5.12}$$

$$M_{\mathrm{u}} = J_1(m) R\left(\frac{-\pi}{4}\right) \begin{bmatrix} e^{i\alpha} & 0 \\ 0 & e^{-i\alpha} \end{bmatrix} R\left(\frac{\pi}{4}\right). \tag{5.13}$$

General expressions for the complex final fields before the output polariser can be given in terms of α, θ_0 the rotation angle of the modulator, and θ_{H} and θ_{Q} (the rotation angles of the half and quarter wave-plates respectively) [72]. By setting conditions for the desired outputs (e.g. that both the real and imaginary parts of M_{u} be zero, for lower single sideband modulation), expressions for the required angles of the optics can be found, using a symbolic mathematical application such as MAPLE™ to simplify the multiplication of the Jones matrices.

The most relevant expressions are as follows:

For upper or lower single sideband modulation,

$$\theta_{\mathrm{Q}} = \frac{1}{2} \arctan\left[\frac{\sin(4\theta_0)(\cos(2\alpha) - 1)}{\cos(2\alpha)cos(4\alpha) - 1 - cos(4\theta_0) - cos(2\alpha)}\right], \tag{5.14}$$

$$\theta_{\mathrm{H}} = -\frac{1}{2} \arctan\left[\frac{\tan(\alpha)\sin(2\theta_0 - 2\theta_{\mathrm{Q}})}{\tan(\alpha)\cos(2\theta_0 - 2\theta_{\mathrm{Q}}) \pm 1}\right]. \tag{5.15}$$

And, for the amplitude modulation case,

$$\theta_{\mathrm{H}} = \arctan\left[\frac{\sin(\theta_{\mathrm{Q}})}{\cos(\theta_{\mathrm{Q}}) + 1}\right]. \tag{5.16}$$

In this case the rotation angle of the quarter wave-plate was essentially arbitrary, and only the relative orientation between the quarter and half wave-plates was relevant.

5.4.4 Results

The power recycling cavity error signals for these situations were then modelled. The error signal for the amplitude modulation case can be seen in *Figure 5.18*, and the corresponding error signal for the (upper) single sideband modulation case can be seen in *Figure 5.19*.

These modelled signals were then verified in the experimental setup. By misaligning the *ETM*, and allowing the *PRM* to swing freely through resonances, the demodulated error signals were plotted: see *Figure 5.20* and *Figure 5.21* for the measured error signals using amplitude modulation and single sideband modulation respectively.

For the amplitude modulation case, the locking signal shows a clear qualitative agreement with the simulated case. In addition to possessing the "correct" features that correspond to resonances of the carrier light and the modulated components, several additional resonances are seen. These are thought to be due to higher order spatial modes partially resonating in the system. These could never be fully suppressed, despite taking care to align the optics.

Similarly, for the single sideband modulation case, the measured error signal strongly resembles the model. Spurious features due to higher order modes are present, in addition to extra noise thought to have arisen due to the incomplete cancellation of the lower sideband, and the inherently smaller signal size associated with the single sideband scheme.

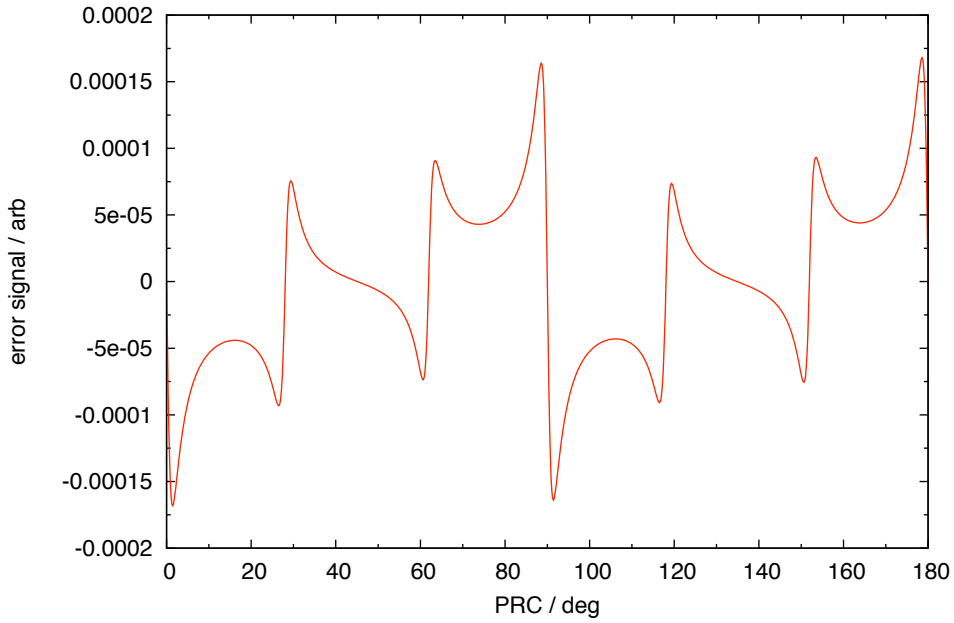

Figure 5.18: Modelled error signal for the power recycling cavity, using amplitude modulation sidebands at 14.525 MHz. In addition to the central bi-polar slope due to the sidebands, resonances from the carrier and upper and lower beat frequencies ±4.525 MHz are present.

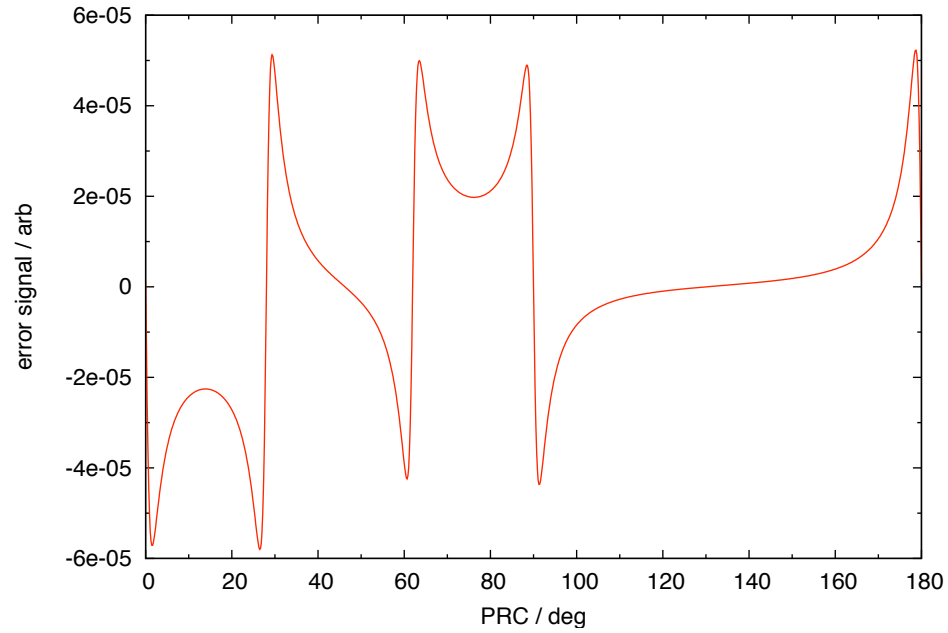

Figure 5.19: Modelled error signal for the power recycling cavity, using a single sideband modulation at 14.525 MHz. In addition to the central bi-polar slope due to the sidebands, resonances from the carrier and upper and lower beat frequencies ±4.525 MHz are present. The signal is asymmetric due to the lack of a lower sideband resonance.

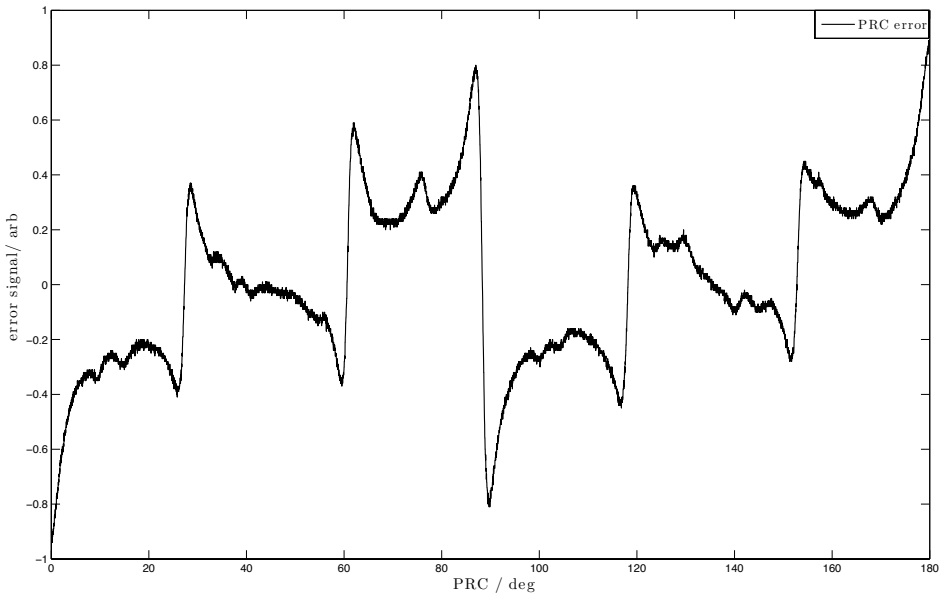

Figure 5.20: Measured demodulated error signal for the power recycling cavity using amplitude modulation sidebands. Additional features are due to higher order spatial modes partially resonating in the system.

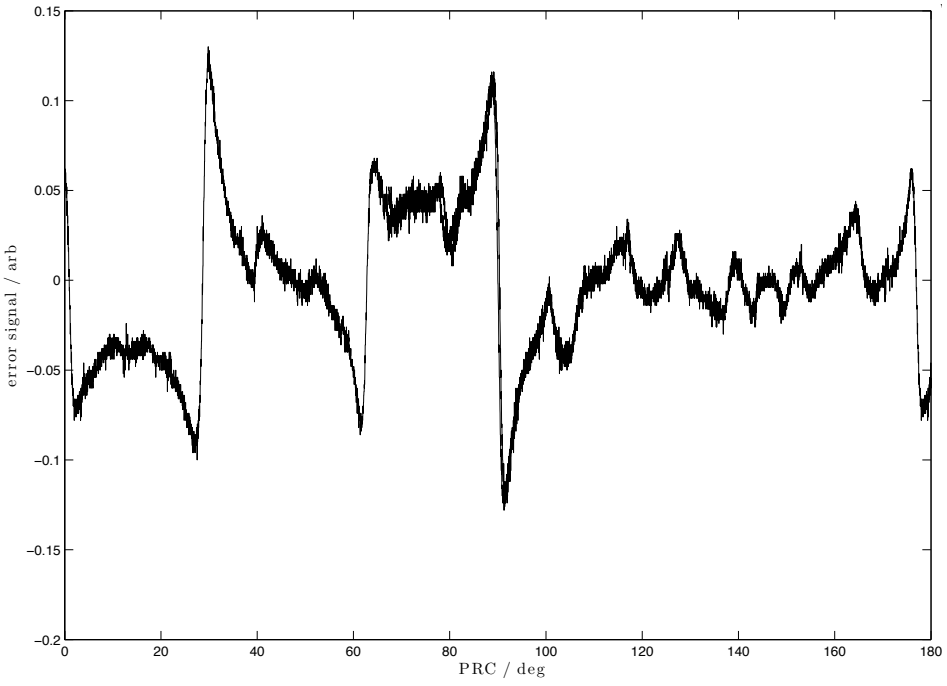

Figure 5.21: Measured demodulated error signal for the power recycling cavity using single sideband modulation. Additional features are due to higher order spatial modes and incomplete cancellation of the lower sideband field.

5.5 Error Signal Optimisation

In *Section 5.3.2* it was noted that the power recycling cavity length sensing scheme was sensitive to changes in the macroscopic length of the *PRC*. Equivalently, this can be expressed as sensitivity to offsets in modulation frequency; where the nominal frequency used was 14.525 MHz. In order to obtain an optimised error signal for the power recycling cavity, it was necessary to match the modulation frequency to the length of the cavity, which could alter with temperature fluctuations, or due to local suspension shifts caused by cycling the vacuum system.

To match the modulation frequency to the true cavity length, a set of error signal sweeps (where the system moves freely through a single resonance) were performed at different demodulation phases of the 4.525 MHz local oscillator. For these measurements, the *ETM* was again misaligned and the *AC* control sidebands turned off, so only the power recycling cavity resonances appeared. The modulation scheme used was the amplitude modulation variant, with dual sidebands.

The "true" modulation frequency was found by altering the modulation frequency and demodulation phase variables, and comparing the resultant demodulated error signals to those modelled in FINESSE.

The first step in this process was to record error signal sweeps for different values of 4.525 MHz demodulation phase. A set of example sweeps (both modelled and measured) is shown in *Figure 5.22*.

The shapes of the resonance features that appeared in the error signal (specifically those due to the carrier and the amplitude modulation sidebands) over different demodulation phases were compared to the model and used to characterise the modulation frequency setup.

The traces seen in *Figure 5.22* were performed using a modulation frequency of 14.735 MHz, and it can be seen that the gradient of the resonance feature due to the carrier (denoted by **c**) changes sign between traces **C** and **D**, while the gradient of the central resonance due to the modulation sidebands (denoted by **14**±) changes sign at a *different* demodulation phase – between traces **F** and **H**.

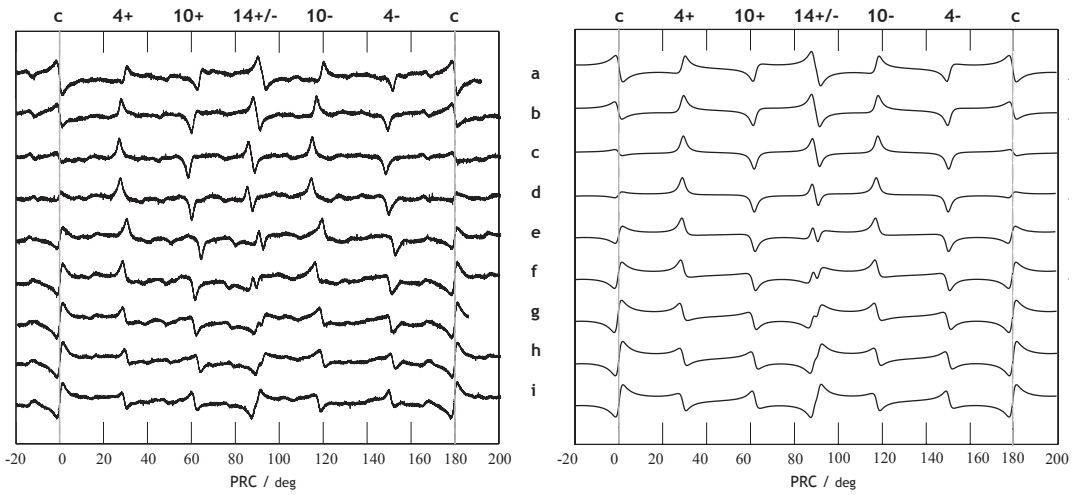

Figure 5.22: Measured (left) and modelled (right) error signal sweeps for the power recycling cavity. Here, the sideband modulation frequency was offset to 14.735 MHz and the sweeps were performed in 20° increments. The resonances are denoted as follows: **c** is the carrier resonance, **10**± is the resonance due to the 10 MHz phase modulation field, **4**± is the resonance due to the beats between the 10 MHz field and the $\sim$ 14 MHz field. The central feature (**14**±) represents the locking signal, due to the $\sim$ 14 MHz sidebands. The modulation indices were $m = 0.6$ for the 10 MHz sidebands and $m = 0.03$ for the amplitude modulated sidebands.

The sign change of the central feature was marked with a "ripple" feature (most visible in traces **F** to **H**), indicating that the upper and lower amplitude modulation sidebands were not simultaneously resonant in the power recycling cavity – implying that the modulation frequency was too high or too low for the physical cavity length. A detail of this region for various demodulation phases can be seen in *Figure 5.23*.

By adjusting the modulation frequency so that the carrier and central $\sim$ 14 MHz ripple feature changed sign at the same demodulation phase, the modulation frequency could be matched to the cavity length with a resolution of $\sim$ 200 kHz. For too high a frequency the $\sim$ 14 MHz feature sign change occurred at a lower demodulation phase than the carrier sign change, and vice-versa. By setting the demodulation phase of the sidebands to be that of the carrier sign change point, and then adjusting the modulation frequency of the sidebands and local oscillator, the carrier and central $\sim$ 14 MHz features could be made to change sign at the same demodulation phase. This served as a useful coarse adjustment step – judging the state of the ripple feature was somewhat subjective, thus limiting the useful accuracy.

A finer optimisation could then be achieved by adjusting the demodulation phase to

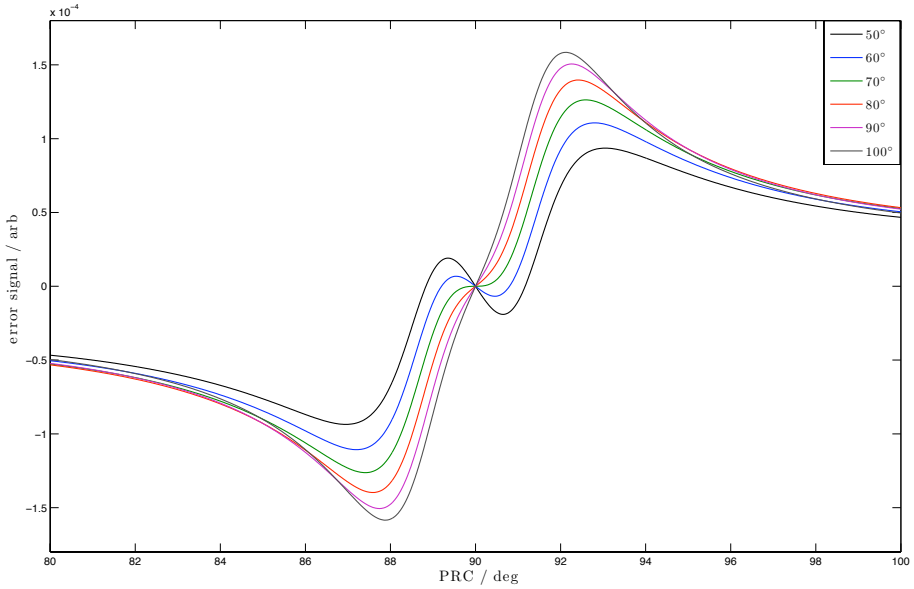

Figure 5.23: Detail of the *PRC* error signal for various demodulation phases, where the modulation frequency is offset by $\sim 200\,\text{kHz}$.

the point where the gradient of the carrier resonance changed sign, and then adjusting the modulation frequency until the ripple features overlapped and were minimised – see *Figure 5.24* and *Figure 5.25*. This procedure allowed frequency optimisation to $\sim 1\,\text{kHz}$, and was performed with the error signal 90° out of phase, as this allowed the ripple features to be minimised more easily.

Ultimately this process was limited by the presence of higher order spatial modes – present due to small misalignments between the three mirrors – whose features on the error signal could not easily be distinguished from those of the carrier or modulation sideband resonances.

5.6 Coupling Measurements

Using the methods for sideband generation described in *Section 5.4*, the methods for error signal optimisation described in *Section 5.5*, and the FINESSE model developed for the system, the appropriate control matrices for both the amplitude and single sideband modulation cases could be modelled, and evaluated experimentally.

After manually aligning all three mirrors, the three-mirror coupled cavity system

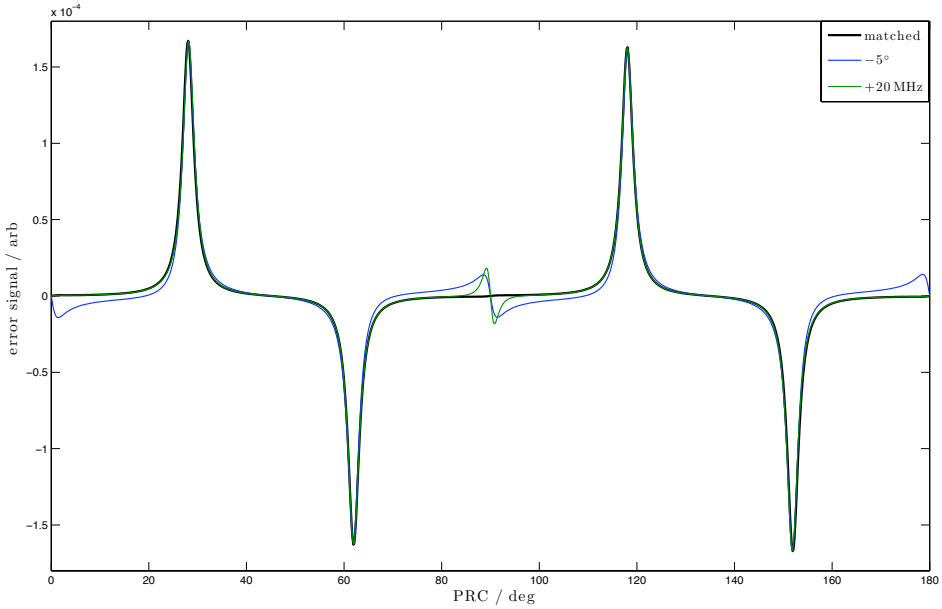

Figure 5.24: Modelled *PRC* error signals, shown out of phase, where the **10**± and **4**± features are prominent. A change in demodulation phase causes features from the carrier and central signal to appear, whereas a change in modulation frequency only causes a central ripple.

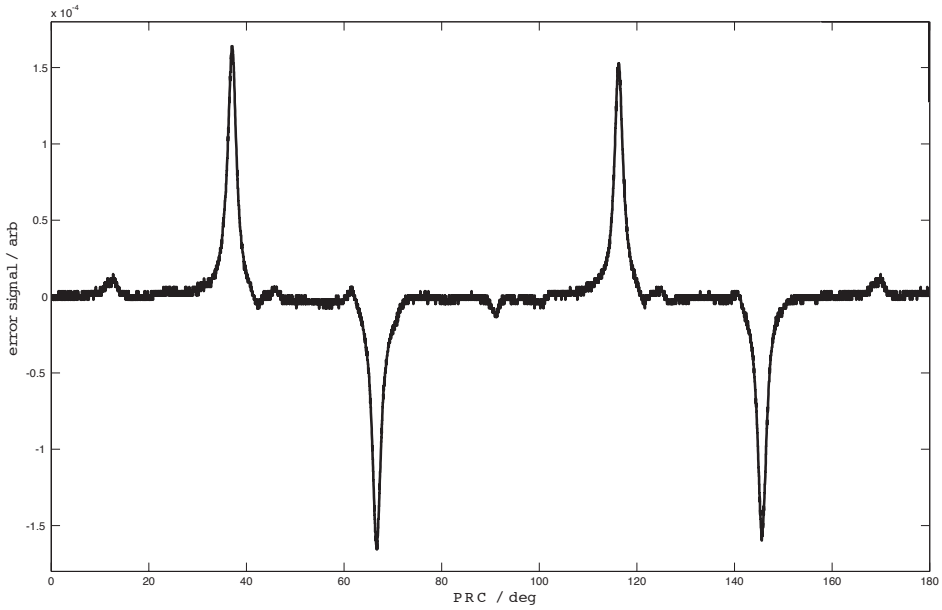

Figure 5.25: A measured, frequency optimised, error signal for the *PRC*, shown out of phase where the resonance features of the carrier and the amplitude modulation sidebands have been minimsed by optimising the modulation frequency. Modelled plots are shown in *Figure 5.24*.

could be locked (simply a matter of enabling the two locking servos and waiting several seconds until the mirrors were all simultaneously close to their locking points) and control signals injected to deduce the required control matrices.

Using a Stanford SR785 spectrum analyser, a signal source was injected into either the arm cavity, or power recycling cavity servo at a single frequency (6.016 kHz [6]), immediately after the appropriate error point. With the system locked, the level of the injected signal peak was then measured at the respective error-points. *Figure 5.26* and *Figure 5.27* show example spectra measured at the *AC* and *PRC* error-points respectively, for a 30 mVPk injection into the *AC* feedback path.

In calculating the control matrices, this measured signal level was then adjusted to account for the optical transfer functions of the cavities, and the open loop gain response of the arm cavity servo – the signal injection lay outwith the control bandwidth of the power recycling servo, and so did not require this correction.

To calibrate the relative responses of the two actuators involved (that is, the PZT and the coil/magnet actuator driving the motion of the *PRM*) the power recycling cavity was locked to the main carrier light, and a signal injected into the *PRC* servo. The transfer function gain at this frequency was measured, before a similar signal was then injected into the PZT path, and again measured at the *PRC* error point, thus providing a measure of the relative effects of the two actuation paths.

[6]This specific frequency was chosen to coincide with the location of a spectrum analyser measurement bin.

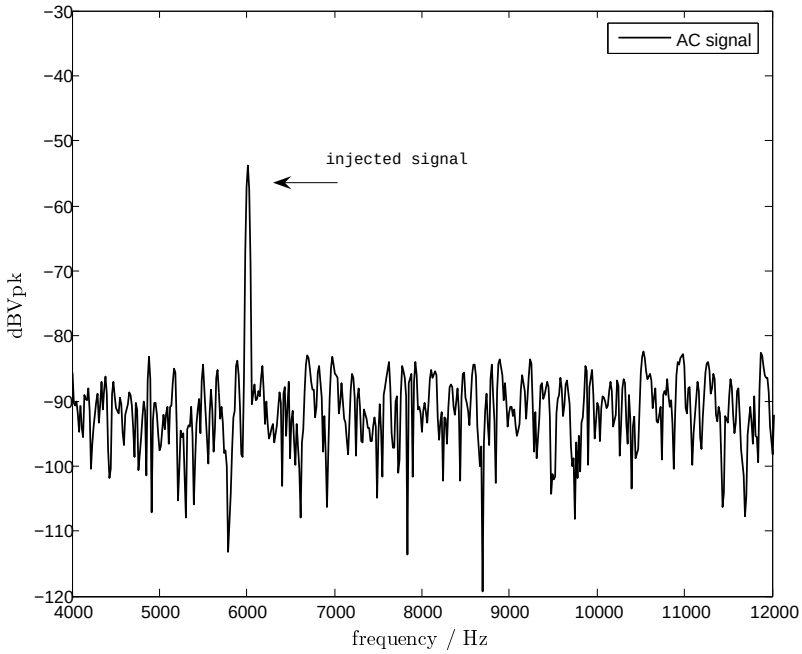

Figure 5.26: The *AC* error-point spectrum, with a 30 mVPk signal at 6.016 kHz injected after the *AC* error-point.

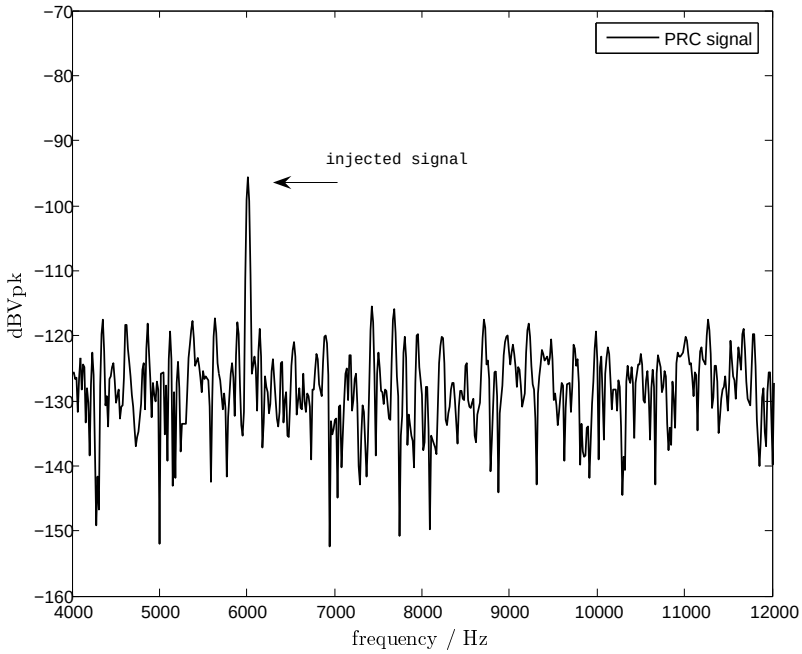

Figure 5.27: The *PRC* error-point spectrum, with a 30 mVPk signal at 6.016 kHz injected after the *AC* error-point.

5.6.1 Amplitude Modulation Scheme

In general it was possible to adjust the power recycling cavity error signal's demodulation phase to either maximise the error signal, or alternatively to minimise the error signal's coupling to the other degree of freedom. The zero crossing of the power recycling cavity error signal always appears at the carrier anti-resonance point, and so the carrier should be fully resonant in the arm cavity, regardless of the chosen demodulation phase – see *Figure 5.28*.

The optimal demodulation phase (for the PRC) was chosen to minimise the coupling to the arm cavity signal[7]. This was achieved through injecting a test signal, and adjusting the demodulation phase until this injected signal was minimised in the arm cavity feedback spectrum.

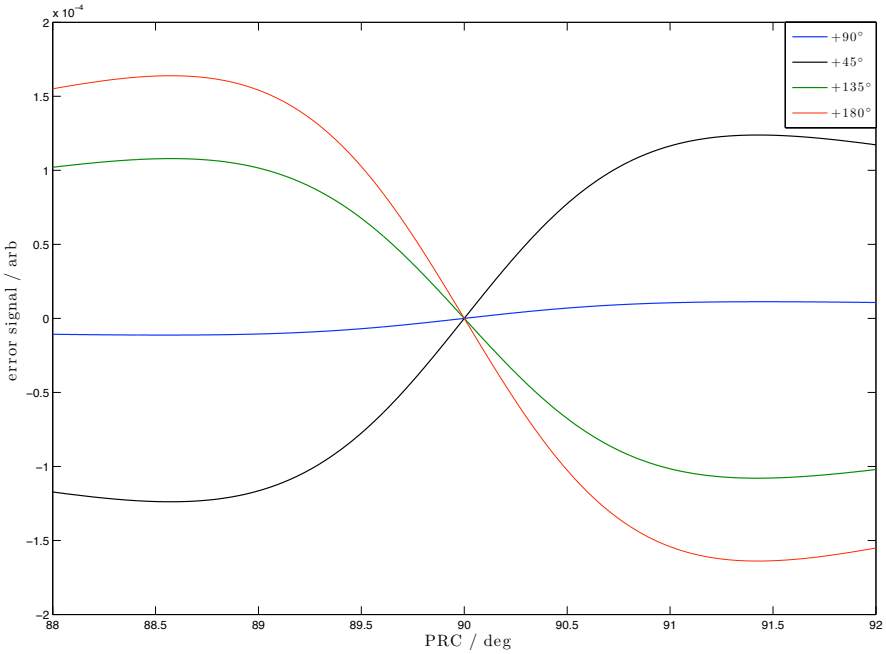

Figure 5.28: The *PRC* error signal for various demodulation phases, for the amplitude modulation scheme.

Using the chosen demodulation phase, FINESSE was used to calculate the appropriate control matrix, and the methods outlined above used to measure the control matrix of the real system – see *Table 5.2*. Note that the control matrices presented

[7]Note the demodulation phase for the arm cavity was trivially set to maximise the *AC* error signal.

here are normalised as a whole – and not by each row individually – to allow more direct comparison of the two modulation methods.

	Φ_{PRC}	Φ_{AC}
PD4	6.4×10^{-6}	1.4×10^{-6}
PD46	9.3×10^{-2}	1

	Φ_{PRC}	Φ_{AC}
PD4	1.3×10^{-5}	1.9×10^{-11}
PD46	2.5×10^{-3}	1

Table 5.2: Measured (left) and modelled (right) normalised control matrices for the amplitude modulation control scheme. All entries are normalised to the arm cavity length sensing signal – the lower right entry.

As required, the non-diagonal terms are smaller than the diagonal ones. The power recycling cavity error signal's response to motion of the *PRC* is seen to be less than modelled, and – more importantly – the coupling between this signal and motion from the *AC* is seen to be a factor of 10^5 larger than modelled. This can be explained due to the fact that this signal was optimised by adjusting the demodulation phase. In practice, this was achievable with a resolution of $\sim 0.1°$. A finer phase adjustment would feasibly have led to better de-coupling between the degrees of freedom.

The discrepancies in the other matrix elements are thought to originate from residual misalignments between the optics, and the presence of higher order spatial modes. Even though considerable time was spent manually aligning the optics in an attempt to minimise this effect, more accurate and stable alignment could certainly improve the results.

5.6.2 Single Sideband Scheme

For the single sideband modulation scheme, the zero crossing of the power recycling cavity error signal is *not* always at the carrier anti-resonant point for arbitrary demodulation frequencies – see *Figure 5.29*. As such, the demodulation phase that coincided with the carrier's anti-resonance was chosen, fulfilling the condition of having the maximum amount of light resonant in the arm cavity, while potentially degrading the level of de-coupling.

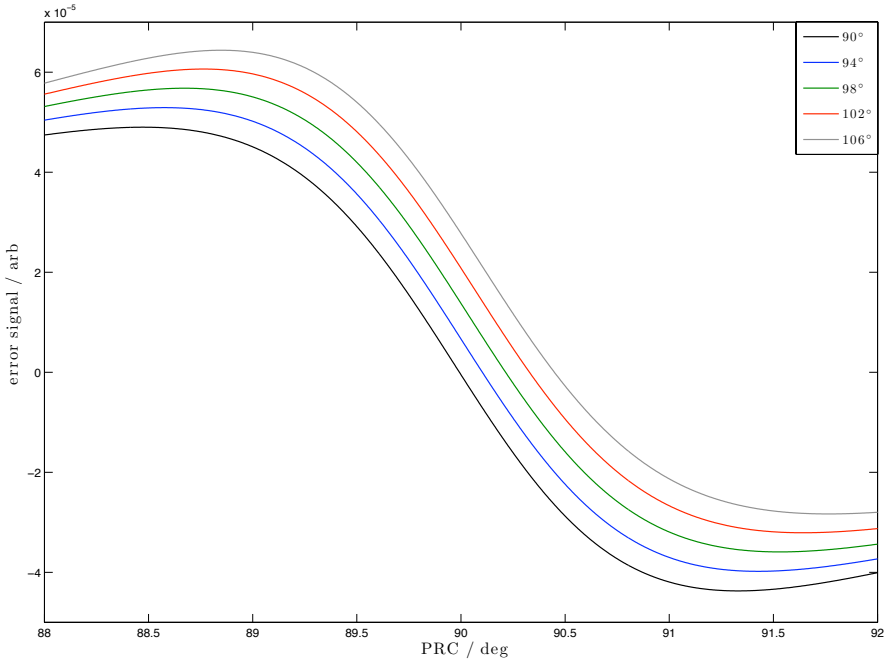

Figure 5.29: The *PRC* error signal for various demodulation phases, for the single sideband modulation scheme.

Again, the FINESSE model was used to calculate a theoretical control matrix, while the experimental matrix was measured identically to the amplitude modulation scheme case. The resultant modelled and measured control matrices are presented in *Table 5.3*.

	Φ_{PRC}	Φ_{AC}
PD4	2.1×10^{-6}	9.0×10^{-5}
PD46	3.2×10^{-2}	1

	Φ_{PRC}	Φ_{AC}
PD4	1.9×10^{-6}	2.8×10^{-5}
PD46	2.5×10^{-3}	1

Table 5.3: Measured (left) and modelled (right) normalised control matrices for the single sideband control scheme. All entries are normalised to the arm cavity length sensing signal – the lower right entry.

Here, the measured matrix elements more closely resemble those from the model. However, it can be seen that in both the model and the measured cases, the *AC* contribution to the *PRC* error signal (the top right element in the matrix) is larger than the contribution from the power recycling cavity itself. As such, although this scheme more closely matches the prediction, it provides a lesser degree of de-coupling than

the amplitude modulation scheme. The discrepancies in the matrix elements are again thought to be due to problems of mis-alignment, and the setting of the demodulation phase. One advantage of the single sideband modulation scheme is that the demodulation phase setting (chosen to maximise the power in the AC) was less sensitive and easier to achieve than that in the amplitude modulation scheme – set by adjusting the demodulation scheme to minimise the size of an injected signal.

5.7 Conclusions

The central goal of these experiments was to effectively de-couple the control signals associated with the two key degrees of freedom of a three-mirror coupled cavity system.

While this was broadly achieved, the degree of de-coupling achieved was ultimately limited by the presence of alignment errors in the system [73] and the finite available adjustability of the demodulation phases.

The amplitude modulation scheme exhibited the best degree of de-coupling, whereas the single sideband modulation scheme required less fine adjustment of the demodulation phase, and agreed with the simulated model more closely.

Despite the observed discrepancies, these results provide a useful proof of concept for similar techniques to be used in future detectors, or other complex coupled optical systems. In particular, it is assumed that the observed alignment issues would either be eradicated or greatly reduced by an auto-alignment system, present on all full-scale detectors.

In addition, a novel method for the flexible generation of different modulation components was developed, modelled and successfully implemented; and the utility and accuracy of the simulation software with respect to a real, complex, optical system was verified.

Chapter 6

Conclusions

Interferometric gravitational wave detectors are currently beginning to reach their design sensitivities and return data that is of scientific and astrophysical importance. While this represents a significant achievement, it will be necessary to improve detector sensitivities even further in order to perform astronomy in any meaningful sense.

Current detectors are limited by several fundamental noise sources. Arguably the most important noise source arises from thermal effects. In order to reduce – or at least quantify – this noise contribution, it will be necessary to measure it directly. While measurements of mechanical resonances can provide useful insights into the phenomenon, only a direct thermal noise measurement system can unambiguously quantify this noise source in the relevant measurement band.

The first part of this work involved the design, modelling, construction and initial com-

missioning of just such an interferometric measurement. All the necessary subsystems required to isolate and control the measurement optics were successfully installed. A flame fibre pulling machine was built, allowing the successful construction of two quasi-monolithic fused silica suspensions.

Initial commissioning of the measurement system was performed, and calibrated sensitivity spectra for the thermal noise measurement cavity recorded, utilising a flexible digital control system.

Although the measured displacement spectrum was not limited by the thermal noise of the desired test optic, there is no evidence to suggest that further, more rigorous, commissioning work and "noise hunting" will not locate and eradicate the adverse noise sources, resulting in a measurement system operating at, or near, its design goals. When fully functional, the system should provide an easily adaptable facility for the direct measurement of the thermal noise associated with test-masses, in the gravitational wave detection band. One future extension of the system already planned, is to evaluate the thermal noise levels of a silicate bonded composite mass – proposed for use in future detectors.

As interferometric gravitational wave detectors develop into second and third generation devices, the input laser power can be expected to increase, and the optical topologies can be expected to grow more complex. These developments will lead to the optical cavities in future devices becoming strongly inter-coupled, through optical, and opto-mechanical means.

The second major part of this work concerned the design and implementation of a novel extension to the standard *Pound-Drever-Hall* optical cavity control scheme, in an effort to de-couple the control of a three-mirror coupled cavity system. The system was successfully designed, and modelled; and appropriate electronics for the con-

trol scheme installed. Measurements of the relevant control matrices for two control schemes (one involving amplitude modulation sidebands, the other a single sideband modulation scheme) were performed and compared to those modelled in software.

The results for the amplitude modulation scheme exhibited the best de-coupling between the relevant degrees of freedom, while the single sideband control scheme results more closely resembled the modelled values. Although both control schemes were ultimately limited by practical issues (such as mirror alignment), the three-mirror coupled cavity system was successfully and reliably locked using both methods, and the experiment provided a proof of concept for the potential use of similar extensions to the *Pound-Drever-Hall* technique in future detectors, or other complex optical systems.

Finally, these two relatively dissimilar experiments were successfully performed in the 10 m JIF prototype, effectively aiding in the commissioning of the facility and demonstrating its utility as an interferometric test-bed for future gravitational wave detectors.

Appendix A

Material Properties

Symbol	Value/Units	Description
α_{SiO_2}	5.5×10^{-7} K^{-1}	Thermal expansion coefficient for SiO_2 [74]
$\alpha_{Ta_2O_5}$	3.6×10^{-6} K^{-1}	Thermal expansion coefficient for Ta_2O_5 [74]
Y_{SiO_2}	7.2×10^{10} Pa	Young's modulus for SiO_2 [49]
$Y_{Ta_2O_5}$	1.4×10^{11} Pa	Young's modulus for Ta_2O_5 [74]
σ_{SiO_2}	0.17	Poisson ratio for SiO_2 [48]
$\sigma_{Ta_2O_5}$	0.23	Poisson ratio for Ta_2O_5 [74]
n_{SiO_2}	1.45	Refractive index for SiO_2 [74]
$n_{Ta_2O_5}$	2.03	Refractive index for Ta_2O_5 [74]
$\phi^{bulk}_{SiO_2}$	5×10^{-8} rad	Loss angle for bulk SiO_2 [49]
$\phi^{layer}_{SiO_2}$	1×10^{-4} rad	Loss angle for thin layer SiO_2 [74]
$\phi^{layer}_{Ta_2O_5}$	3.8×10^{-4} rad	Loss angle for thin layer Ta_2O_5 [74]
C_{SiO_2}	1.64×10^{6} JK^{-1}m^{-3}	Specific heat capacity for SiO_2 [74]
$C_{Ta_2O_5}$	2.1×10^{6} JK^{-1}m^{-3}	Specific heat capacity for Ta_2O_5 [74]
κ_{SiO_2}	1.38 Wm^{-1}K^{-1}	Thermal conductivity for SiO_2 [74]
$\kappa_{Ta_2O_5}$	33 Wm^{-1}K^{-1}	Thermal conductivity for Ta_2O_5 [74]
ρ_{SiO_2}	2200 kgm^{-3}	Density of SiO_2 [49]
β_{SiO_2}	1.5×10^{-5} K^{-1}	Temperature dependence of n_{SiO_2}[1] [47][75][76]
$\beta_{Ta_2O_5}$	1.21×10^{-4} K^{-1}	Temperature dependence of $n_{Ta_2O_5}$ [47]
γ_{SiO_2}	1.52×10^{-4} K^{-1}	Temperature dependence of Y_{SiO_2} [77]

[1] the value given in [47] is -1.5×10^{-5} K^{-1}. The negative sign is considered erroneous as [75] uses the positive value cited above for Advanced LIGO noise estimations, and [76] gives roughly the same value of $\sim +1.25\times10^{-5}$ K^{-1}.

Appendix B

Electronics

Alignment

To allow remote control of the alignment of the mirrors inside the vacuum enclosure, all the suspensions involved in the thermal noise experiment had four 10 mm × 5 mm magnets (made from NeFeB) bonded to the rear sides of their upper stages. A DC force was then applied to these magnets by passing current through coils placed behind them. The circuit used to supply these signals is seen below in *Figure B.1*.

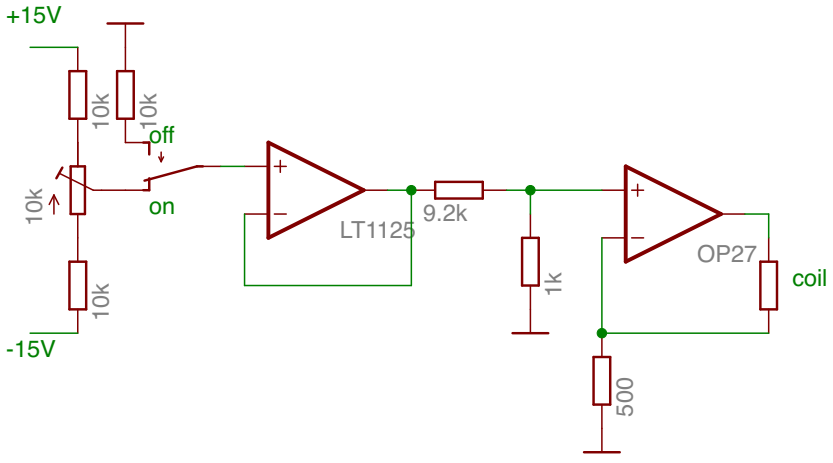

Figure B.1: Alignment control circuit: the suspensions were aligned by applying force to the upper stages by driving a current through a coil of resistance ∼ 10 Ω. Variations of this circuit provided alignment control over all seven suspensions.

A fine (ten turn) potentiometer provided an input level of between ±5 V, followed by a buffer amplifier, a voltage divided and an output stage. The BS suspension used the circuit as pictured above. The other metal suspensions required increased range, and so had the 1 kΩ resistor before the output stage removed. The silica suspensions required yet more range, so (additionally) the resistor to ground in the output stage was reduced from 500 Ω to 242 Ω.

The BS required less adjustable range, as it had the longest optical path to the output, while the silica suspensions required extra range as it was not desirable to manually shift the suspensions – readily achievable with the all-metal ones – so any initial suspension offsets had to be compensated electronically.

Coil Driver

The feedback signal to lock the cavity was provided in a similar way to the alignment control: via an electromagnetic force applied to magnets bonded to the ETM_{TNE} mass. The optic had four 5 mm × 3 mm magnets bonded to its rear side, while the RM suspension lower stage held four coils.

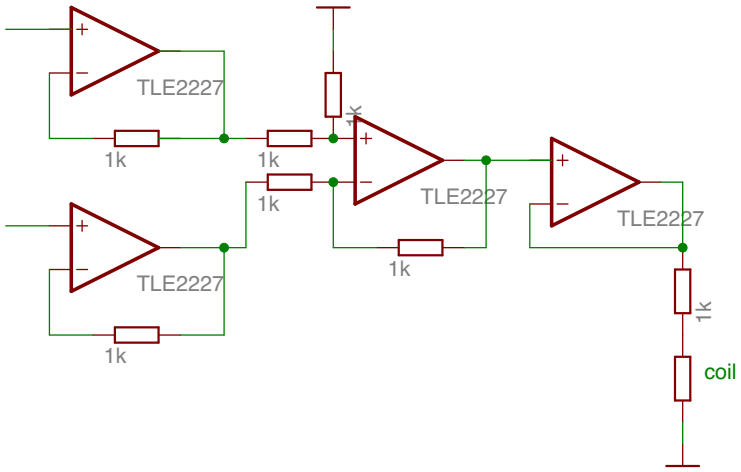

Figure B.2: TNE coil driver circuit: four of these circuits provided the feedback signals to the four coils attached to the lower RM mass, enabling the test cavity to be locked. Note that the op-amps used were changed to OPA227s, due to stability problems with the TLE2227 models.

The coil driver circuit used was adapted from the circuit used to provide the coil driver signals for the larger masses used in the main arm cavity. This contained a differential receiving stage, one input of which was simply grounded for use in the *TNE* setup. The coils were connected in series with a 1 kΩ resistor to reduce the current through the feedback coils to more manageable levels.

Additionally, a four-way splitter box was in place prior to the coil-driver circuit. This contained four potentiometers, used to optionally adjust the voltage applied to each coil separately. Note that these were all set to unity gain for the measurements performed.

Photodiode

The signal used to lock the thermal noise test cavity (and therefore measure the displacement noise) comes from the interference of the main carrier light with the imposed phase modulation sidebands, which must be in the radio-frequency range to lie sufficiently outside any cavity linewidths. Typically, small modulation indices are used, so the detection photodiode must be optimised to detect small signals at high frequencies. In addition, it is also useful to be sensitive to changes in the DC light power, as this allows monitoring of the overall light level, and can facilitate triggering methods[1]. A passive resonant *LC* circuit can fulfill these requirements – see *Figure B.3* below.

The resonant circuit is formed between a variable inductor (here marked as 3 μH) and the parasitic capacitance of the photodiode ($\sim$ 85 pF) along with any stray capacitance present. A high-pass *RC* filter at the entrance to the high frequency path and a low-pass *LR* filter in the DC path effectively separate the small high frequency signal component from the overall light level. The photodiode used was an EG&G InGaAs model C30642GH with a 2 mm^2 active area and photocurrent generation efficiency of $\sim$ 0.78 A/W, and was reverse biased at 7 V. The 1 n capacitor provided a ground to the RF path of the circuit, while presenting an effectively infinite impedance at DC.

[1]The time taken to successfully lock a cavity can be reduced by activating the feedback when the cavity is transiently resonating. This can be achieved by monitoring the DC return light and triggering the feedback control to coincide with a preset threshold.

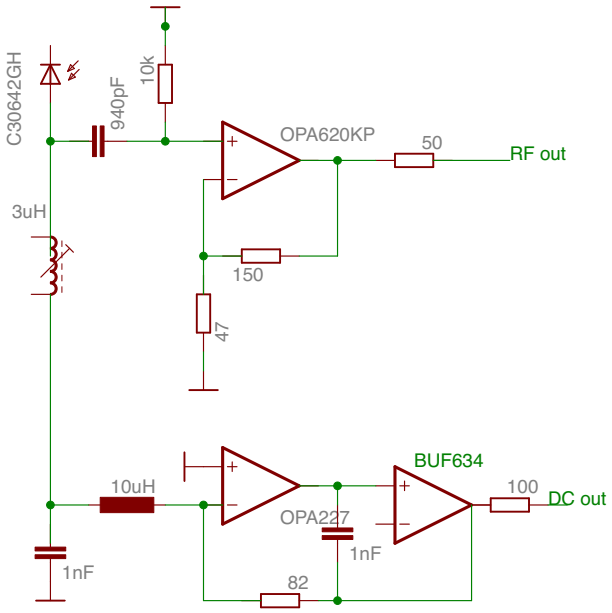

Figure B.3: *TNE* photodiode circuit. The resonant high frequency path was tunable around the 10 MHz target by approximately ±1 MHz. A low noise, high frequency OPA620 op-amp was used in the RF path, while an OPA227 and BUF634 combination comprised the DC path.

The impedance (Z_{RF}) of the high frequency path – used to convert the signal current from the photodiode into a voltage – is given by [78],

$$Z_{\mathrm{RF}} \simeq Q\sqrt{\frac{L}{C}}, \tag{B.1}$$

where Q is a measure of the strength of the resonance, and L and C are the inductance of the variable inductor, and C the capacitance of the photodiode.

The Q was measured by placing a wound wire coil (connected to a signal generator) above the circuit, and monitoring the output pin of the photodiode with an oscilloscope. The circuit could first be tuned to the correct resonant frequency by setting the output of the signal generator to 10 MHz and adjusting the tunable inductor until a maximum output was achieved. The Q could then be found by varying the signal generator output frequency about the resonance, and noting the photodiode output – see *Figure B.4*. The Q factor was then the resonant frequency of the peak, divided by the width at $\frac{1}{\sqrt{2}}$ of its maximum.

The signal voltage is proportional to both the incident photocurrent and the impedance Z_{RF}. The impedance will also have a thermal Johnson noise associated with it, proportional to $\sqrt{Z_{\mathrm{RF}}}$. It is desirable then to maximise this conversion impedance (and therefore the Q) in order to increase the signal-to-noise ratio of the detection process.

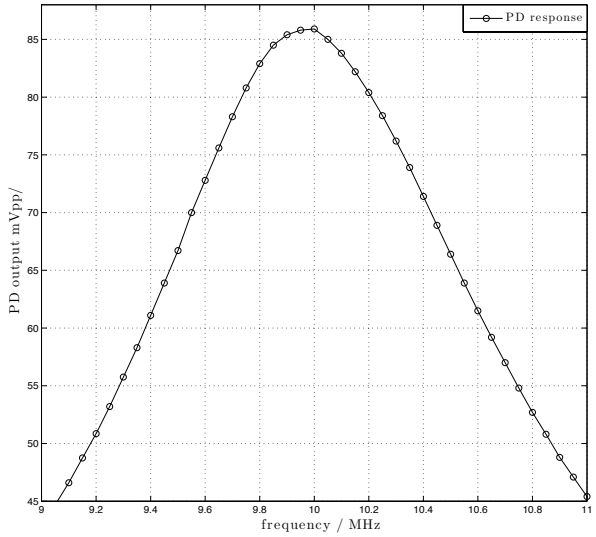

Figure B.4: Response of the TNE resonant photodiode circuit around 10 MHz.

The Q measured for the thermal noise experiment photodiode was $Q \sim 8$, giving a value of $Z_{\mathrm{RF}} = 1.5$ kΩ

Note that the resistors in the gain stage of the RF path are neglected in this analysis, but as their resistances are considerably lower than that of the conversion impedance, their relative Johnson noise will be correspondingly less significant.

It is now possible to estimate an equivalent photocurrent – I_{NE} – due to the noise in this resonant conversion process [78] by evaluating the following expression:

$$I_{\mathrm{NE}} = \frac{1}{2e}\left(\frac{4k_{\mathrm{b}}T}{Z_{\mathrm{RF}}} + I_{\mathrm{op}}^2 + \frac{V_{\mathrm{op}}^2}{Z_{\mathrm{RF}}^2}\right). \tag{B.2}$$

Here, e is the charge on an electron, k_{b} is Boltzmann's constant, T is the temperature, I_{op} is the input current noise of the op-amp used, and V_{op} is the input voltage noise of the same op-amp. Using the appropriate values for the OPA620KP op-amp

used in the high frequency path[2] this was calculated to be $I_{\mathrm{NE}} = 60\ \mu\mathrm{A}$.

The voltage noise of the output can then be expressed as [79][80],

$$\xi_V^2 = 2eZ_{\mathrm{RF}}^2 \left(I_{\mathrm{DC}} + I_{\mathrm{NE}}\right), \tag{B.3}$$

where I_{DC} is the DC photocurrent. For the noise associated with the detection system to be lower than that of the photocurrent shot noise, I_{NE} must be smaller than I_{DC}. As $I_{\mathrm{DC}} = 0.78 \times 70{\times}10^{-3} = 5.46{\times}10^{-2}$ A, when a nominal 70 mW is incident on the photodiode, this condition is satisfied.

Returning to the DC path, the relatively large and constant value of I_{DC} necessitates the use of the BUF634 in conjunction with the OPA227, to act as a high-current buffer amplifier.

Sample traces from the photodiode outputs (after the RF output has been demodulated and low-pass filtered) can be seen in *Figure B.5*.

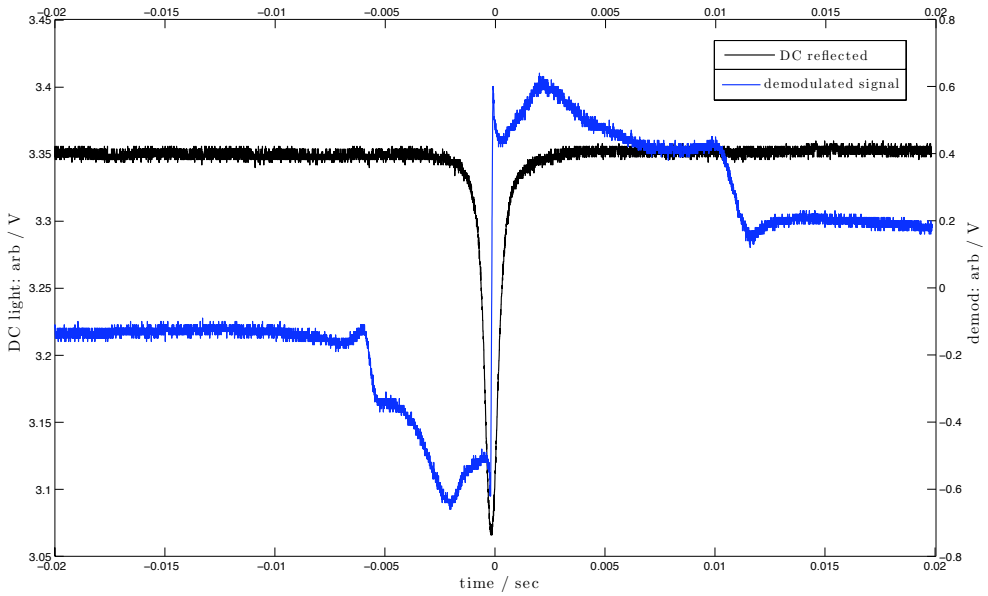

Figure B.5: The reflected DC light from the cavity, and associated demodulated error signal – as measured with the thermal noise resonant photodiode – is shown here as the cavity passes through a single resonance.

[2] $I_{\mathrm{op}} = 2.3{\times}10^{-12}\ \mathrm{A}/\sqrt{\mathrm{Hz}}$ and $V_{\mathrm{op}} = 2.3{\times}10^{-9}\ \mathrm{V}/\sqrt{\mathrm{Hz}}$.

Differential Signals: Sending & Receiving

To reduce electrical pick-up noise accrued when traversing the lab, differential "send/ receive" boxes were constructed. In *send* mode, these would take an input signal and output it through two channels, one of which was inverted. In receive mode, the circuit would subtract these two channels, thus subtracting any common noise, and outputting both a positive and negative version of the signal, each with $\times 2$ gain. See *Figure B.6.*

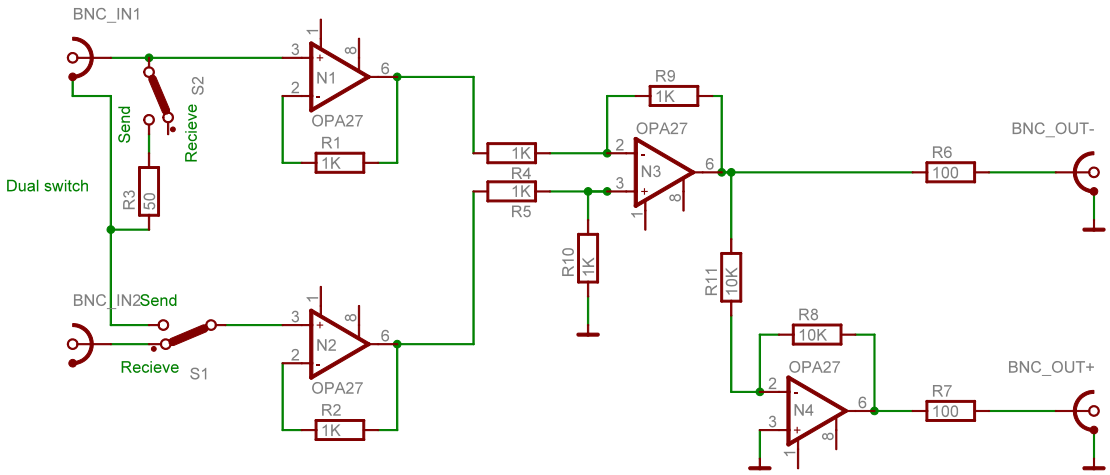

Figure B.6: Circuit diagram for the differential send / receive boxes.

Whitening & De-whitening

Converting a signal from the analogue to digital domain (e.g. on entering the dSPACE system) inevitably results in some loss of information, termed *quantisation noise*. The amplitude spectral density noise associated with converting an analogue voltage into its digital representation is [81]

$$\xi_{\mathrm{ADC}} = \sqrt{\frac{\Delta^2}{12 f_{\mathrm{N}}}} \ \mathrm{V}/\sqrt{\mathrm{Hz}}, \tag{B.4}$$

where Δ is the minimum representable voltage (i.e. the voltage range of the input divided by the resolution of the conversion process) and f_{N} is the Nyquist frequency, equivalent to half the sampling frequency.

This expression can be rewritten in terms of the input voltage range V_{range}, the time

between discrete samples t_s and the number of bits associated with the analogue to digital conversion – b – as follows:

$$\xi_{\mathrm{ADC}} = \frac{V_{\mathrm{range}}}{2^b}\frac{t_s}{6} \quad \mathrm{V}/\sqrt{\mathrm{Hz}}. \tag{B.5}$$

This represents not a noise *limit*, but an effect that is actively imposed on the signal. As it is a white noise spectrum that is applied to all measurement frequencies, it will act to limit the dynamic range of the feedback at higher frequencies, where the feedback signal is typically smaller. To combat this, the signal is *whitened* before entering the dSPACE system, and then *de-whitened* on exiting. The whitening filter – see *Figure B.7* – acts to (approximately) reshape the error signal towards a flat spectrum, while the de-whiteneing filter – see *Figure B.8* – performs the inverse operation.

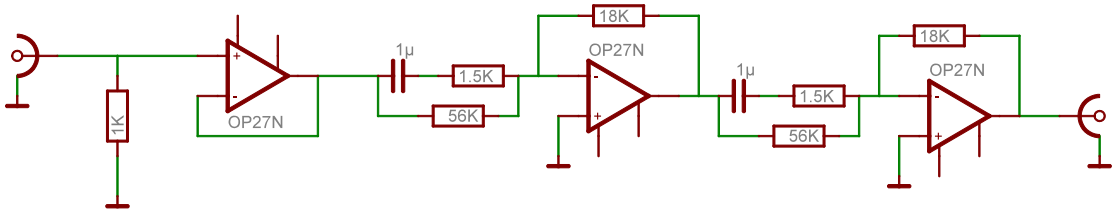

Figure B.7: Circuit diagram for the dSPACE whitening filter.

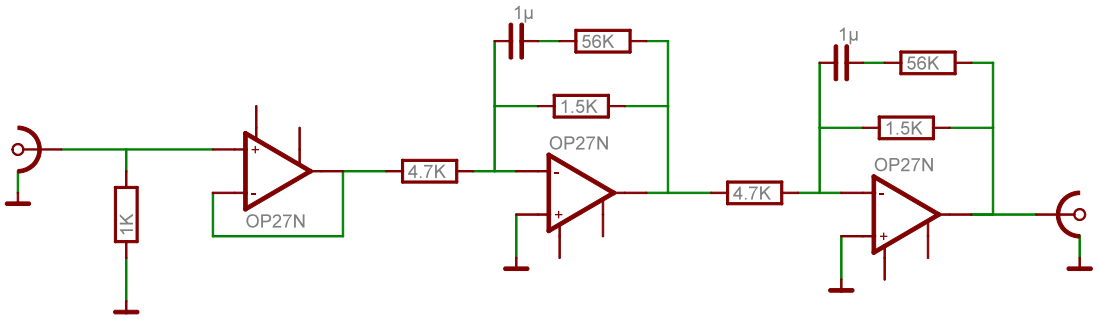

Figure B.8: Circuit diagram for the dSPACE de-whitening filter.

Frequency Stabilisation Servo

The laser frequency stabilisation servo was designed to feed back to the laser frequency through two primary paths: the PZT bonded to the laser, and an EOM phase modulator in the path of the laser beam. The PZT path was the "slow" path, governing frequencies up to ~ 10 kHz, with the EOM element controlling feedback at frequencies above this, up to the unity gain frequency of the system. At very low frequencies (≤ 1 Hz), either the ETM position feedback was used (via a coil-driver and reaction mass set-up), or the temperature control feature of the laser was implemented.

The two paths consist of numerous stages, designed to impose frequency filters such that – in conjunction with the optical cavity and feedback element responses – the system is controlled stably up to ~ 100 kHz. Other features notable in the diagram – *Figure B.9* – are potentiometers to control the relative gains of the two paths, trimming potentiometers for the removal of unwanted DC offsets in the system and an additional buffer amplifier at the input to allow unobtrusive monitoring of the input error point. Additionally a Stanford pre-amp was placed after the PZT output to allow extra adjustability, and a high voltage amplifier was used to send the feedback signal to the EOM.

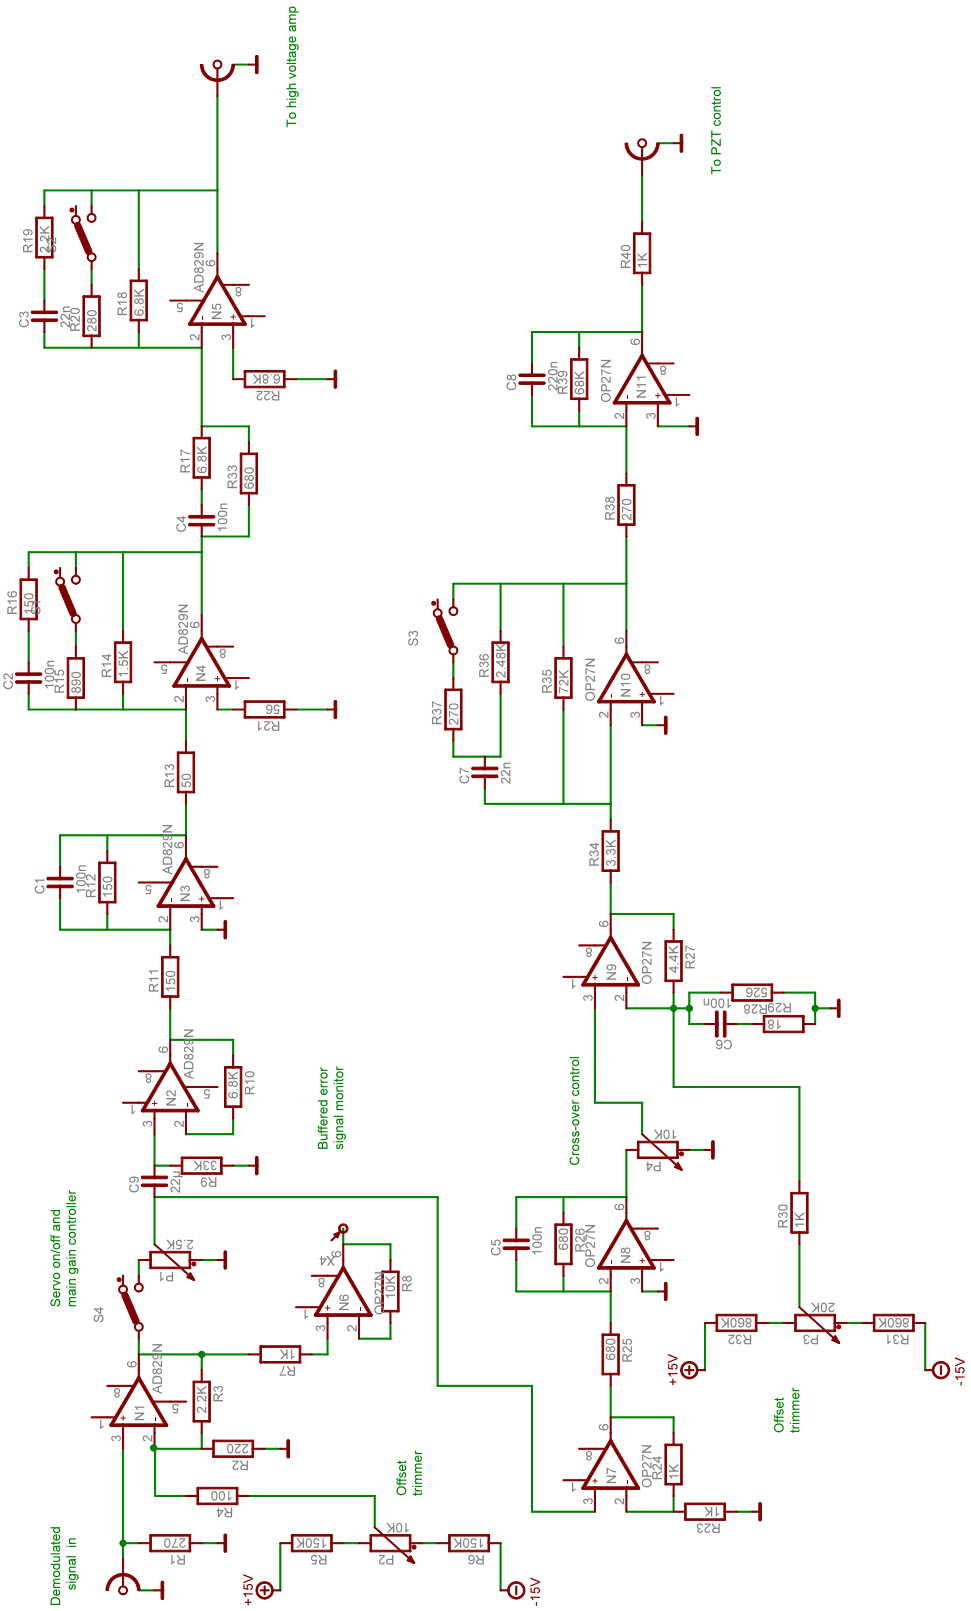

Figure B.9: Circuit diagram for the laser feedback servo, with outputs for the PZT on the laser crystal and the feedback EOM.

Miscellany

Modulators

To achieve modulation sidebands of any reasonable amplitude, large voltages must be applied across the electro-optic modulators. A typical half-wave voltage (that is, the voltage required to induce a phase shift of π radians) is of the order of several hundred volts. The large voltages and high frequencies involved require the power transfer circuit to be suitably impedance matched, to both reduce unwanted electrical pick-up in other electronics, and to maximise the efficiency of the sideband generation.

An auto-transformer circuit – see *Figure B.10* – was used, with a resonant LC circuit formed by the capacitance of the EOM itself (and an optional trimming capacitor for adjustment) and an air-wound coil inductor. The input impedance was matched to the 50 Ω source by adjusting the input tapping point on the coil.

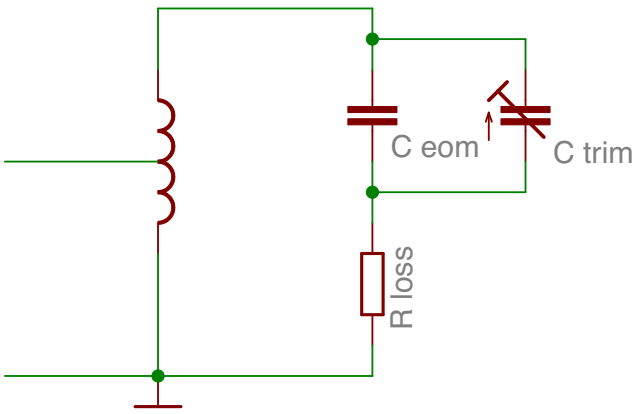

Figure B.10: Auto-transformer circuit for matching power into electro-optic modulators.

Low-Pass Filter

When dealing with demodulated signals, it is frequently desirable to low-pass filter them to effectively remove residual high frequency noise from pre-demodulation, the LO input to the mixer, and stray high frequency pick-up. The filter in *Figure B.11* is an effective *T-filter* design that provides > 90 dB of filtering at 10 MHz when matched

into a 50 Ω load, as is shown in *Figure B.12*.

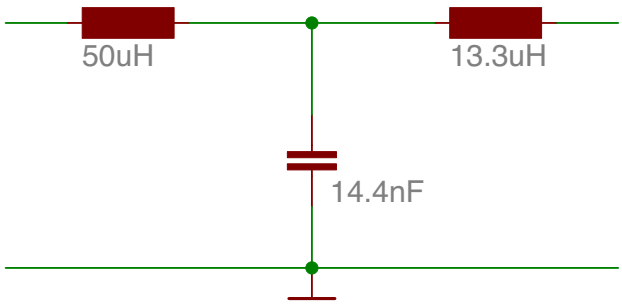

Figure B.11: Circuit design for the 300 kHz low-pass T-filter.

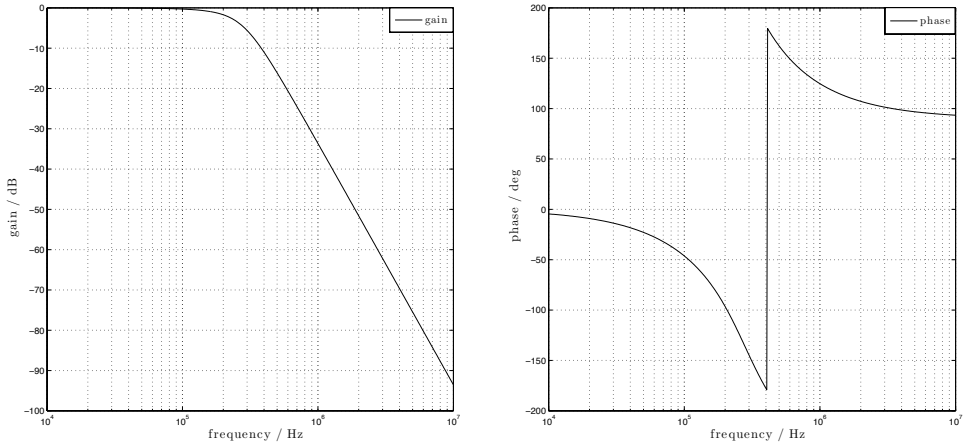

Figure B.12: Modelled gain and phase response of the low-pass T-filter.

Zener Clamp

When attempting to first feed back to a system, where the appropriate gain levels are not know, it can be useful to restrict the feedback signal to known safe levels. However, reducing the gain can alter the characteristics of the feedback loop. Using two Zener diodes, back-to-back in a Zener "clamp" setup, as pictured in *Figure B.13* restricts the maximum voltage to $\pm V_{zener}$ where V_{zener} is the Zener voltage of the diode, and the

current is then given by $I_{\max} = \frac{V_s - V_{\text{zener}}}{R}$; where V_s is the input voltage and R is the resistor pictured. This circuit acts to restrict the maximum output of a system without adversely affecting the gain.

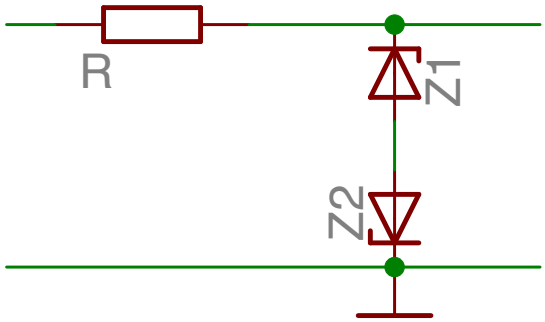

Figure B.13: Zener "clamp" circuit, to restrict maximum signal without adversely affecting gain.

Appendix C

Measuring Loop Transfer Functions

When constructing a control system to stabilise a generic plant, it is desirable to know both the open and closed-loop transfer functions of the system. However, frequently – and specifically in the case of an interferometer – it is not always easy to measure the transfer functions of each stage in the loop and simply find their product. An optical cavity feedback system, such as the *Pound-Drever-Hall* method described in *Section 2.3.1*, is only linear close to the operating point of the cavity: i.e. the system must be locked (the loop must be closed) for any meaningful characterisation to take place.

A method for obtaining both the open and closed-loop transfer functions for such a system, using a standard spectrum analyser and software, is presented below.

Open-Loop Transfer Function

In order to measure an open-loop transfer function of a servo system while the system is locked, a suitable injection and measurement stage must be found. Preferably, this should be the inverting input of an inverting amplifier stage, where the amplifier stage presents a flat frequency response over the measurement range.

A swept-sine measurement from the input of the stage to its output (taking care not to cause any saturation) should then be performed. This measurement gives,

$$T_1 = \frac{G_x}{1 - GH}, \tag{D.1}$$

where G_x is the transfer function of the measurement stage, and GH is the open-

loop transfer function of the complete system.

The same measurement should then be performed with the loop unlocked, and the inputs to the electronics grounded so as to avoid injecting large transient signals due to cavity resonances. This measurement gives,

$$T_2 = G_x, \tag{D.2}$$

and the open-loop transfer function GH is then given by,

$$GH = 1 - \frac{T_2}{T_1}. \tag{D.3}$$

Closed-Loop Transfer Function

To obtain a closed-loop transfer function the measurements required are exactly the same as above, but with differing analysis. The closed-loop transfer function $\frac{GH}{1-GH}$ is given by,

$$\frac{GH}{1-GH} = \frac{T_1}{T_2} - 1, \tag{D.4}$$

where T_1 is the swept-sine response of an injection stage in the electronics while the loop is locked, and T_2 is the same measurement with the loop open.

Practical Method

The analysis above assumes that the quantities T_1 and T_2 are vectors of complex numbers, representing the gain and phase of the measurement. Depending on the specific analyser and methods used, this may or may not be the case. In the case where a *Stanford SRS785* analyser (or similar) is used, the following method will yield the appropriate information.

The measurement traces should be saved in the analyser's internal *".78D"* format, via the *Display to Disk* option. The *srt785.exe* command line application can then be used to convert the internal format into a three column ASCII delimited file consisting

of [*frequency, real, imaginary*] data. An example command being,

```
srt785 /Oasc /D /Cx,r,i SRS001.78D T1.asc
```

On importing the data into MATLAB, the *real* and *imaginary* parts can be combined into a complex number, and the appropriate transfer functions constructed as follows:

```
load T1.ASC -ASCII
load T2.ASC -ASCII

freq = T1(:,1);

T1z = T1(:,2) + i*T1(:,3);
T2z = T2(:,2) + i*T2(:,3);

OLTF =  1 - (T2z./T1z);

OLTF_gain = abs(OLTF);
OLTF_gain_dB = 20*log10(OLTF_gain);
OLTF_phase = angle(OLTF);

CLTF = (T1z./T2z) - 1;

CLTF_gain = abs(CLTF);
CLTF_gain_dB = 20*log10(CLTF_gain);
CLTF_phase = angle(CLTF);
```

Plotting the transfer function vectors against the frequency vector defined in the first line then yields the required transfer function plots.

Bibliography

[key] ADS: *Astrophysics Data System* reference DOI: *Digital Object Identifier* ISBN: *International Standard Book Number* URL: *Internet address*

[1] D. N. Spergel, R. Bean, O. Doré, M. R. Nolta, C. L. Bennett, J. Dunkley, G. Hinshaw, N. Jarosik, E. Komatsu, L. Page, H. V. Peiris, L. Verde, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. *Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications for Cosmology.* The Astrophysical Journal Supplement Series, **170**:377-408, June 2007.

ADS: 2007ApJS..170..377S
DOI: 10.1086/513700

[2] F. K. Baganoff, Y. Maeda, M. Morris, M. W. Bautz, W. N. Brandt, W. Cui, J. P. Doty, E. D. Feigelson, G. P. Garmire, S. H. Pravdo, G. R. Ricker, and L. K. Townsley. *Chandra X-Ray Spectroscopic Imaging of Sagittarius A* and the Central Parsec of the Galaxy.* The Astrophysical Journal, **591**:891-915, July 2003.

ADS: 2003ApJ...591..891B
DOI: 10.1086/375145

[3] A. Einstein. *Die Grundlage der allgemeinen Relativitätstheorie.* Annalen der Physik, **354**:769-822, 1916.

ADS: 1916AnP...354..769E
DOI: 10.1002/andp.19163540702

[4] A. Einstein. *Zur Elektrodynamik bewegter Körper.* Annalen der Physik, **322**:891-921, 1905.

ADS: 1905AnP...322..891E
DOI: 10.1002/andp.19053221004

[5] P. R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors.* World Scientific Publishing, 1994.

ISBN: 978-981-02-1820-1

[6] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation.* W.H. Freeman and Company, 2000.

ISBN: 978-0-7167-0344-0

[7] B. F. Schutz. *A first course in general relativity*. Cambridge University Press, 1985.

ISBN: 978-0-521-27703-7

[8] G. Woan. *The Cambridge Handbook of Physics Formulas*. Cambridge University Press, 2000.

ISBN: 978-0-521-57507-2

[9] S. M. Carroll. *The Cosmological Constant*. Living Reviews in Relativity, **4**(1), 2001.

URL: http://www.livingreviews.org/lrr-2001-1

[10] B. F. Schutz. *The detection of gravitational waves*. In Proceedings of the 1995 Houches School on Astrophysical Sources of Gravitational Radiation. Springer Berlin, 1995.

[11] B. F. Schutz. *Determining the Hubble constant from gravitational wave observations*. Nature, **323**:310, September 1986.

ADS: 1986Natur.323..310S
DOI: 10.1038/323310a0

[12] *Gravitational Waves. Proceedings of the 1994 Snowmass Summer Study on Particle and Nuclear Astrophysics and Cosmology*. World Scientific Publishing, 1995.

[13] D. G. Blair, I. S. Heng, E. N. Ivanov, and M. E. Tobar. *Present Status of the Resonant Mass Gravitational Wave Antenna, NIOBE*. In Second Edoardo Amaldi Conference on Gravitational Wave Experiments, 1998.

[14] M. Visco. *The EXPLORER and NAUTILUS Gravitational Wave Detectors and Beyond*. In General Relativity and Gravitational Physics, volume 751 of *American Institute of Physics Conference Series*, 2005.

[15] L. Conti. *The AURIGA second scientific run and the dual detector of gravitational waves*. Nuclear Instruments and Methods in Physics Research A, **518**:236-239, February 2004.

ADS: 2004NIMPA.518..236C
DOI: 10.1016/j.nima.2003.10.070

[16] M. P. McHugh, Z. Allen, W. O. Hamilton, W. W. Johnson, and G. Santostasi. *The ALLEGRO gravitational wave detector.*. International Journal of Modern Physics D, **9**:229-232, June 2000.

ADS: 2000IJMPD...9..229M
DOI: 10.1142/S0218271800000190

[17] P. Astone, D. Babusci, L. Baggio, M. Bassan, M. Bignotto, M. Bonaldi, M. Camarda, P. Carelli, G. Cavallari, M. Cerdonio, A. Chincarini, E. Coccia, L. Conti, S. D'Antonio, M. de Rosa, M. di Paolo Emilio, M. Drago, F. Dubath, V. Fafone, P. Falferi, S. Foffa, P. Fortini, S. Frasca, G. Gemme, G. Giordano, G. Giusfredi, W. O. Hamilton, J. Hanson, M. Inguscio, W. W. Johnson, N. Liguori, S. Longo, M. Maggiore, F. Marin, A. Marini, M. P. McHugh, R. Mezzena, P. Miller, Y. Minenkov, A. Mion, G. Modestino, A. Moleti, D. Nettles, A. Ortolan, G. V. Pallottino, R. Parodi, G. Piano Mortari, S. Poggi, G. A. Prodi, L. Quintieri, V. Re, A. Rocchi, F. Ronga, F. Salemi, G. Soranzo, R. Sturani, L. Taffarello, R. Terenzi, G. Torrioli, R. Vaccarone, G. Vandoni, G. Vedovato, A. Vinante, M. Visco, S. Vitale, J. Weaver, J. P. Zendri, and P. Zhang. *Results of the IGEC-2 search for gravitational wave bursts during 2005*. Physical Review D, **76**(10):102001, November 2007.

ADS: 2007PhRvD..76j2001A
DOI: 10.1103/PhysRevD.76.102001

[18] S. E. Whitcomb. *Ground-based gravitational-wave detection: now and future*. Classical and Quantum Gravity, **25**(11):114013, June 2008.

ADS: 2008CQGra..25k4013W
DOI: 10.1088/0264-9381/25/11/114013

[19] D. Sigg and the LIGO Science Collaboration. *Status of the LIGO detectors*. Classical and Quantum Gravity, **23**:51, April 2006.

ADS: 2006CQGra..23S..51S
DOI: 10.1088/0264-9381/23/8/S07

[20] F. Acernese, P. Amico, M. Alshourbagy, F. Antonucci, S. Aoudia, P. Astone, S. Avino, D. Babusci, G. Ballardin, F. Barone, L. Barsotti, M. Barsuglia, F. Beauville, S. Bigotta, S. Birindelli, M. A. Bizouard, C. Boccara, F. Bondu, L. Bosi, C. Bradaschia, S. Braccini, A. Brillet, V. Brisson, D. Buskulic, E. Calloni, E. Campagna, F. Carbognani, F. Cavalier, R. Cavalieri, G. Cella, E. Cesarini, E. Chassande-Mottin, N. Christensen, C. Corda, A. Corsi, F. Cottone, A.-C. Clapson, F. Cleva, J.-P. Coulon, E. Cuoco, A. Dari, V. Dattilo, M. Davier, M. del Prete, R. DeRosa, L. Di Fiore, A. Di Virgilio, B. Dujardin, A. Eleuteri, I. Ferrante, F. Fidecaro, I. Fiori, R. Flaminio, J.-D. Fournier, S. Frasca, F. Frasconi, L. Gammaitoni, F. Garufi, E. Genin, A. Gennai, A. Giazotto, G. Giordano, L. Giordano, R. Gouaty, D. Grosjean, G. Guidi, S. Hebri, H. Heitmann, P. Hello, S. Karkar, S. Kreckelbergh, P. La Penna, M. Laval, N. Leroy, N. Letendre, B. Lopez, M. Lorenzini, V. Loriette, G. Losurdo, J.-M. Mackowski, E. Majorana, C. N. Man, M. Mantovani, F. Marchesoni, F. Marion, J. Marque, F. Martelli, A. Masserot, M. Mazzoni, L. Milano, F. Menzinger, C. Moins, J. Moreau, N. Morgado, B. Mours, F. Nocera, C. Palomba, F. Paoletti, S. Pardi, A. Pasqualetti, R. Passaquieti, D. Passuello, F. Piergiovanni, L. Pinard, R. Poggiani, M. Punturo, P. Puppo, K. Qipiani, P. Rapagnani, V. Reita, A. Remillieux, F. Ricci, I. Ricciardi, P. Ruggi, G. Russo, S. Solimeno, A. Spallicci, M. Tarallo, M. Tonelli, A. Toncelli, E. Tournefier, F. Travasso, C. Tremola, G. Vajente, D. Verkindt,

F. Vetrano, A. Viceré, J.-Y. Vinet, H. Vocca, and M. Yvert. *Status of Virgo detector*. Classical and Quantum Gravity, **24**:381, October 2007.

ADS: 2007CQGra..24..381A
DOI: 10.1088/0264-9381/24/19/S01

[21] D. Tatsumi, R. Takahashi, K. Arai, N. Nakagawa, K. Agatsuma, T. Yamazaki, M. Fukushima, M.-K. Fujimoto, A. Takamori, A. Bertolini, V. Sannibale, R. DeSalvo, S. Márka, M. Ando, K. Tsubono, T. Akutsu, K. Yamamoto, H. Ishitsuka, T. Uchiyama, S. Miyoki, M. Ohashi, K. Kuroda, N. Awaya, N. Kanda, A. Araya, S. Telada, T. Tomaru, T. Haruyama, A. Yamamoto, N. Sato, T. Suzuki, and T. Shintomi. *Current status of Japanese detectors*. Classical and Quantum Gravity, **24**:399, October 2007.

ADS: 2007CQGra..24..399T
DOI: 10.1088/0264-9381/24/19/S03

[22] K. Kuroda and the LCGT Collaboration. *The status of LCGT*. Classical and Quantum Gravity, **23**:215, April 2006.

ADS: 2006CQGra..23S.215K
DOI: 10.1088/0264-9381/23/8/S27

[23] Advanced LIGO team. *Advanced LIGO Reference Design*. Technical report. Ligo, May 2007.

URL: http://www.ligo.caltech.edu/docs/M/M060056-08/

[24] The Virgo Collaboration. *Advanced Virgo Conceptual Design: VIR-042A-07*. Technical report. Virgo, October 2007.

URL: http://wwwcascina.virgo.infn.it/advirgo/docs/AdV_Design.pdf

[25] B. Willke, P. Ajith, B. Allen, P. Aufmuth, C. Aulbert, S. Babak, R. Balasubramanian, B. W. Barr, S. Berukoff, A. Bunkowski, G. Cagnoli, C. A. Cantley, M. M. Casey, S. Chelkowski, Y. Chen, D. Churches, T. Cokelaer, C. N. Colacino, D. R. M. Crooks, C. Cutler, K. Danzmann, R. J. Dupuis, E. Elliffe, C. Fallnich, A. Franzen, A. Freise, I. Gholami, S. Goßler, A. Grant, H. Grote, S. Grunewald, J. Harms, B. Hage, G. Heinzel, I. S. Heng, A. Hepstonstall, M. Heurs, M. Hewitson, S. Hild, J. Hough, Y. Itoh, G. Jones, R. Jones, S. H. Huttner, K. Kötter, B. Krishnan, P. Kwee, H. Lück, M. Luna, B. Machenschalk, M. Malec, R. A. Mercer, T. Meier, C. Messenger, S. Mohanty, K. Mossavi, S. Mukherjee, P. Murray, G. P. Newton, M. A. Papa, M. Perreur-Lloyd, M. Pitkin, M. V. Plissi, R. Prix, V. Quetschke, V. Re, T. Regimbau, H. Rehbein, S. Reid, L. Ribichini, D. I. Robertson, N. A. Robertson, C. Robinson, J. D. Romano, S. Rowan, A. Rüdiger, B. S. Sathyaprakash, R. Schilling, R. Schnabel, B. F. Schutz, F. Seifert, A. M. Sintes, J. R. Smith, P. H. Sneddon, K. A. Strain, I. Taylor, R. Taylor, A. Thüring, C. Ungarelli, H. Vahlbruch, A. Vecchio, J. Veitch, H. Ward, U. Weiland, H. Welling, L. Wen, P. Williams, W. Winkler, G. Woan, and R. Zhu. *The GEO-HF project*. Classical and Quantum Gravity, **23**:207, April 2006.

ADS: 2006CQGra..23S.207W
DOI: 10.1088/0264-9381/23/8/S26

[26] A. Freise, S. Chelkowski, S. Hild, W. Del Pozzo, A. Perreca, and A. Vecchio. *Triple Michelson Interferometer for a Third-Generation Gravitational Wave Detector*. ArXiv e-prints, **0804**:1036, April 2008.

ADS: 2008arXiv0804.1036F
DOI: 10.1088/0264-9381/26/8/085012

[27] K. Danzmann and A. Rüdiger. *LISA technology - concept, status, prospects*. Classical and Quantum Gravity, **20**:1, May 2003.

ADS: 2003CQGra..20S...1D
DOI: 10.1088/0264-9381/20/10/301

[28] R. T. Stebbins and LISA Pathfinder Science Team. *The LISA Pathfinder Mission*. In Bulletin of the American Astronomical Society, volume 38 of *Bulletin of the American Astronomical Society*, December 2007.

ADS: 2007AAS...21114009S

[29] R. A. Hulse and J. H. Taylor. *Discovery of a pulsar in a binary system*. The Astrophysical Journal Letters, **195**:L51-L53, January 1975.

ADS: 1975ApJ...195L..51H
DOI: 10.1086/181708

[30] M. J. Valtonen, H. J. Lehto, K. Nilsson, J. Heidt, L. O. Takalo, A. Sillanpää, C. Villforth, M. Kidger, G. Poyner, T. Pursimo, S. Zola, J.-H. Wu, X. Zhou, K. Sadakane, M. Drozdz, D. Koziel, D. Marchev, W. Ogloza, C. Porowski, M. Siwak, G. Stachowski, M. Winiarski, V.-P. Hentunen, M. Nissinen, A. Liakos, and S. Dogru. *A massive binary black-hole system in OJ287 and a test of general relativity*. Nature, **452**:851-853, April 2008.

ADS: 2008Natur.452..851V
DOI: 10.1038/nature06896

[31] H. Sudou, S. Iguchi, Y. Murata, and Y. Taniguchi. *Orbital Motion in the Radio Galaxy 3C 66B: Evidence for a Supermassive Black Hole Binary*. Science, **300**:1263-1265, May 2003.

ADS: 2003Sci...300.1263S
DOI: 10.1126/science.1082817

[32] F.A. Jenet, A. Lommen, S.L. Larson, and L. Wen. *Constraining the Properties of Supermassive Black Hole Systems Using Pulsar Timing: Application to 3C 66B*. The Astrophysical Journal, **606**(2):799-803, 2004.

DOI: 10.1086/383020
URL: http://stacks.iop.org/0004-637X/606/799

[33] B. Abbott and the LIGO Science Collaboration. *Upper limits on gravitational wave emission from 78 radio pulsars*. Physical Review D, **76**(4):042001, August 2007.

ADS: 2007PhRvD..76d2001A
DOI: 10.1103/PhysRevD.76.042001

[34] B. Abbott and the LIGO Science Collaboration. *Beating the Spin-Down Limit on Gravitational Wave Emission from the Crab Pulsar*. The Astrophysical Journal Letters, **683**:L45-L49, August 2008.

ADS: 2008ApJ...683L..45A
DOI: 10.1086/591526

[35] S. Goßler. *The suspension systems of the interferometric gravitational-wave detector GEO 600*. PhD thesis. Universität Hannover, 2004.

[36] C. I. E. Torrie. *Development of Suspensions for the GEO 600 Gravitational Wave Detector*. PhD thesis. University of Glasgow, January 2001.

[37] H. B. Callen and T. A. Welton. *Irreversibility and Generalized Noise*. Physical Review, **83**:34-40, July 1951.

ADS: 1951PhRv...83...34C
DOI: 10.1103/PhysRev.83.34

[38] H. B. Callen and R. F. Greene. *On a Theorem of Irreversible Thermodynamics*. Physical Review, **86**:702-710, June 1952.

ADS: 1952PhRv...86..702C
DOI: 10.1103/PhysRev.86.702

[39] R. F. Greene and H. B. Callen. *On a Theorem of Irreversible Thermodynamics. II*. Physical Review, **88**:1387-1391, December 1952.

ADS: 1952PhRv...88.1387G
DOI: 10.1103/PhysRev.88.1387

[40] P. R. Saulson. *Thermal noise in mechanical experiments*. Physical Review D, **42**:2437-2445, October 1990.

ADS: 1990PhRvD..42.2437S
DOI: 10.1103/PhysRevD.42.2437

[41] Y. Levin. *Internal thermal noise in the LIGO test masses: A direct approach*. Phys. Rev. D, **57**(2):659–663, January 1998.

ADS: PhysRevD.57.659
DOI: 10.1103/PhysRevD.57.659

[42] D. R. M. Crooks. *Mechanical Loss and its Significance in the Test Mass Mirrors of Gravitational Wave Detectors*. PhD thesis. University of Glasgow, May 2003.

[43] G. M. Harry, A. M. Gretarsson, P. R. Saulson, S. E. Kittelberger, S. D. Penn, W. J. Startin, S. Rowan, M. M. Fejer, D. R. M. Crooks, G. Cagnoli, J. Hough, and N. Nakagawa. *Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings*. Classical and Quantum Gravity, **19**:897-917, March 2002.

ADS: 2002CQGra..19..897H
DOI: 10.1088/0264-9381/19/5/305

[44] G. M. Harry, H. Armandula, E. Black, D. R. M. Crooks, G. Cagnoli, J. Hough, P. Murray, S. Reid, S. Rowan, P. Sneddon, M. M. Fejer, R. Route, and S. D. Penn. *Thermal noise from optical coatings in gravitational wave detectors*. Applied Optics, **45**:1569-1574, March 2006.

ADS: 2006ApOpt..45.1569H
DOI: 10.1364/AO.45.001569

[45] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. *Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae*. Physics Letters A, **264**:1-10, December 1999.

ADS: 1999PhLA..264....1B
DOI: 10.1016/S0375-9601(99)00785-9

[46] S. Rowan, J. Hough, and D. R. M. Crooks. *Thermal noise and material issues for gravitational wave detectors*. Physics Letters A, **347**:25-32, November 2005.

ADS: 2005PhLA..347...25R
DOI: 10.1016/j.physleta.2005.06.055

[47] V. B. Braginsky and S. P. Vyatchanin. *Thermodynamical fluctuations in optical mirror coatings*. Physics Letters A, **312**:244-255, June 2003.

ADS: 2003PhLA..312..244B
DOI: 10.1016/S0375-9601(03)00473-0

[48] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. *Thermo-refractive noise in gravitational wave antennae*. Physics Letters A, **271**:303-307, July 2000.

ADS: 2000PhLA..271..303B
DOI: 10.1016/S0375-9601(00)00389-3

[49] A. W. Heptonstall. *Characterisation of Mechanical Loss in Fused Silica Ribbons for use in Gravitational Wave Detector Suspensions*. PhD thesis. University of Glasgow, August 2004.

[50] A. M. Gretarsson, G. M. Harry, S. D. Penn, P. R. Saulson, W. J. Startin, S. Rowan, G. Cagnoli, and J. Hough. *Pendulum mode thermal noise in advanced interferometers: a comparison of fused silica fibers and ribbons in the presence of surface loss*. Physics Letters A, **270**:108-114, May 2000.

ADS: 2000PhLA..270..108G
DOI: 10.1016/S0375-9601(00)00295-4

[51] S. A. Hughes and K. S. Thorne. *Seismic gravity-gradient noise in interferometric gravitational-wave detectors*. Physical Review D, **58**(12):122002, December 1998.

ADS: 1998PhRvD..58l2002H
DOI: 10.1103/PhysRevD.58.122002

[52] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. *Laser phase and frequency stabilization using an optical resonator*. Applied Physics B: Lasers and Optics, **31**:97-105, June 1983.

ADS: 1983ApPhB..31...97D
DOI: 10.1007/BF00702605

[53] A. Freise. FINESSE 0995*: Frequency domain INterferomEter Simulation SoftwarE*, 2007.

URL: http://www.rzg.mpg.de/~adf/

[54] E. Black. *Notes on the Pound-Drever-Hall technique*. Technical report. Ligo, April 1998.

URL: http://www.ligo.caltech.edu/docs/T/T980045-00.pdf

[55] S. Reid. *Studies of materials for future ground-based and space-based interferometric gravitational wave detectors*. PhD thesis. University of Glasgow, July 2006.

[56] J. Mizuno. *Comparison of optical configurations for laser-interferometric gravitational-wave detectors*. PhD thesis. Max-Planck-Institut Für Quantenoptik, May 2003.

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

[58] M. Perreur-Lloyd and R. Jones. SolidWorks® *designs and drawings for the JIF Thermal Noise Experiment*, March 2008.

[59] K. G. Baigent, D. A. Shaddock, M. B. Gray, and D. E. McClelland. *Laser Stabilisation for the Measurement of Thermal Noise*. General Relativity and Gravitation, **32**:399-410, March 2000.

ADS: 2000GReGr..32..399B
DOI: 10.1023/A:1001963914433

[60] *Innolight mephisto data-sheets*.

URL: http://www.innolight.de/

[61] M. V. Plissi, C. I. Torrie, M. Barton, N. A. Robertson, A. Grant, C. A. Cantley, K. A. Strain, P. A. Willems, J. H. Romie, K. D. Skeldon, M. M. Perreur-Lloyd, R. A. Jones, and J. Hough. *An investigation of eddy-current damping of multistage pendulum suspensions for use in interferometric gravitational wave detectors*. Review of Scientific Instruments, **75**:4516-4522, November 2004.

ADS: 2004RScI...75.4516P
DOI: 10.1063/1.1795192

[62] M. E. Husman, C. I. Torrie, M. V. Plissi, N. A. Robertson, K. A. Strain, and J. Hough. *Modeling of multistage pendulums: Triple pendulum suspension for GEO 600*. Review of Scientific Instruments, **71**:2546-2551, June 2000.

ADS: 2000RScI...71.2546H
DOI: 10.1063/1.1150646

[63] M. V. Plissi, K. A. Strain, C. I. Torrie, N. A. Robertson, S. Killbourn, S. Rowan, S. M. Twyford, H. Ward, K. D. Skeldon, and J. Hough. *Aspects of the suspension system for GEO 600*. Review of Scientific Instruments, **69**:3055-3061, August 1998.

ADS: 1998RScI...69.3055P
DOI: 10.1063/1.1149054

[64] S. Rowan, S. M. Twyford, J. Hough, D.-H. Gwo, and R. Route. *Mechanical losses associated with the technique of hydroxide-catalysis bonding of fused silica*. Physics Letters A, **246**:471-478, September 1998.

ADS: 1998PhLA..246..471R
DOI: 10.1016/S0375-9601(98)00533-7

[65] E. Elliffe. *Aspects of Thermal Noise Modeling in Ground-based Gravitational Wave Detectors and Developments of Hydroxide Catalysis Bonding for Space-based Gravitational Wave Detectors and other Optical Applications*. PhD thesis. University of Glasgow, June 2005.

[66] J. R. Smith, G. M. Harry, J. C. Betzwieser, A. M. Gretarsson, D. A. Guild, S. E. Kittelberger, M. J. Mortonson, S. D. Penn, and P. R. Saulson. *Mechanical loss associated with silicate bonding of fused silica*. Classical and Quantum Gravity, **20**:5039-5047, December 2003.

ADS: 2003CQGra..20.5039S
DOI: 10.1088/0264-9381/20/23/007

[67] S. Goßler, G. Cagnoli, D. R. M. Crooks, H. Lück, S. Rowan, J. R. Smith, K. A. Strain, J. Hough, and K. Danzmann. *Damping and tuning of the fibre violin modes in monolithic silica suspensions*. Classical and Quantum Gravity, **21**:923, March 2004.

ADS: 2004CQGra..21S.923G
DOI: 10.1088/0264-9381/21/5/082

[68] *Promatech Ltd.*

URL: http://www.promatech.co.uk

[69] G. Heinzel. LISO *circuit modeling program.*

URL: ftp://ftp.rzg.mpg.de/pub/grav/ghh/liso/

[70] E. Gustafson, D. Shoemaker, K. Strain, and R. Weiss. *LSC white paper on detector research and development*. Technical report. Ligo, September 1999.

URL: http://www.ligo.caltech.edu/docs/T/T990080-00.pdf

[71] O. Miyakawa, R. Ward, R. Adhikari, M. Evans, B. Abbott, R. Bork, D. Busby, J. Heefner, A. Ivanov, M. Smith, R. Taylor, S. Vass, A. Weinstein, M. Varvella, S. Kawamura, F. Kawazoe, S. Sakata, and C. Mow-Lowry. *Measurement of optical response of a detuned resonant sideband extraction gravitational wave detector*. Physical Review D, **74**(2):022001, July 2006.

ADS: 2006PhRvD..74b2001M
DOI: 10.1103/PhysRevD.74.022001

[72] B. W. Barr, S. H. Huttner, J. R. Taylor, B. Sorazu, M. V. Plissi, and K. A. Strain. *Optical modulation techniques for length sensing and control of optical cavities.* Appl. Opt., **46**(31):7739–7745, 2007.

DOI: 10.1364/AO.46.007739
URL: http://ao.osa.org/abstract.cfm?URI=ao-46-31-7739

[73] S. H. Huttner, B. W. Barr, M. V. Plissi, J. R. Taylor, B. Sorazu, and K. A. Strain. *Novel sensing and control schemes for a three-mirror coupled cavity.* Classical and Quantum Gravity, **24**(15):3825-3836, 2007.

DOI: 10.1088/0264-9381/24/15/004
URL: http://stacks.iop.org/0264-9381/24/3825

[74] D. R. M. Crooks, G. Cagnoli, M. M. Fejer, G. Harry, J. Hough, B. T. Khuri-Yakub, S. Penn, R. Route, S. Rowan, P. H. Sneddon, I. O. Wygant, and G. G. Yaralioglu. *Experimental measurements of mechanical dissipation associated with dielectric coatings formed using SiO_2, Ta_2O_5 and Al_2O_3.* Classical and Quantum Gravity, **23**:4953-4965, August 2006.

ADS: 2006CQGra..23.4953C
DOI: 10.1088/0264-9381/23/15/014

[75] *GWINC: Gravitational Wave Interferometer Noise Calculator*, 2008.

URL: http://ilog.ligo-wa.caltech.edu:7285/advligo/GWINC

[76] T. Toyoda and M. Yabe. *The temperature dependence of the refractive indices of fused silica and crystal quartz.* Journal of Physics D Applied Physics, **16**:L97-L100, May 1983.

ADS: 1983JPhD...16L..97T
DOI: 10.1088/0022-3727/16/5/002

[77] P. Willems, V. Sannibale, J. Weel, and V. Mitrofanov. *Investigations of the dynamics and mechanical dissipation of a fused silica suspension.* Physics Letters A, **297**:37-48, May 2002.

ADS: 2002PhLA..297...37W
DOI: 10.1016/S0375-9601(02)00380-8

[78] G. Heinzel. *Advanced optical techniques for laser-interferometric gravitational-wave detectors.* PhD thesis. Universität Hannover, 1999.

[79] D. Shoemaker, R. Schilling, L. Schnupp, W. Winkler, K. Maischberger, and A. Rüdiger. *Noise behavior of the Garching 30-meter prototype gravitational-wave detector.* Phys. Rev. D, **38**(2):423–432, July 1988.

ADS: PhysRevD.38.423
DOI: 10.1103/PhysRevD.38.423

[80] K. Numata. *Direct measurement of mirror thermal noise*. PhD thesis. University of Tokyo, December 2002.

[81] B. Allen and P. Brady. *Quantization Noise in Ligo Interferometers*. Technical report. Ligo, September 1997.

URL: http://www.ligo.caltech.edu/docs/T/T970128-02.pdf