An Instrument for the Measurement of the Spectrum of Thermoelastic Noise in Sapphire

by

Christopher Hollitt

Thesis submitted for the degree of

Doctor of Philosophy

in

The University of Adelaide

School of Chemistry and Physics

October, 2007

This work contains no material which has been accepted for the award of any other degree or diploma in any University or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being available for loan and photocopying.

SIGNED: .....

DATE: 8/11/2007

Supervisor: Peter Veitch.

Acknowledgments

Well, here it is; a monument to far too many days spent in a room with no windows. Though this document is long, it would be far longer if it were to include the many helpful hints and suggestions given to me by so many wonderful people over the years. For all these things I will be eternally grateful, but sadly this meagre page will have to suffice in the way of acknowledgment.

As is customary, I begin by thanking Peter Veitch, my supervisor, and Jesper Munch, head of our little band. I find that words utterly fail me in attempting to convey my regard for them. I will never be able to repay their efforts and will have to rely on the universe to do it for me. I would also like to thank Bob Vincent, Murray Hamilton and particularly Don McCoy, for showing me what a real physicist looks like and for proving that the species is not yet quite extinct.

My sincere thanks go to David McClelland and David Blair for their hospitality during my valuable visits to ANU and UWA respectively. My thanks go to John Winterflood, Daniel Shaddock, Bram Slagmolen, Karl Baigent and particularly to Malcolm Gray for their invaluable advice during those trips and subsequently.

Returning to Adelaide, I cannot sing the praises of the super technicians loudly enough. To Blair Middlemiss, Trevor Waterhouse, Graham Eames and Bob Nation, each of whom taught me an astonishing amount, my most sincere thanks and admiration. To Bob Hurn I offer my most humble gratitude, for his wisdom, for allowing my invasion and most of all, for his friendship. My thanks also to Carmel Palumbo for her astonishing ability to decipher the incomprehensible.

To David Ottaway and Damien Mudge, who both taught me so much about the realities of life in the lab, my thanks. To my fellow travellers, Aidan Brooks, Matthew Heintze, Alex Hemming, Tim Hill, David Hosken and Shu-yen Lee, thankyou all; in addition to all of your help, you made each day far more pleasant than they would otherwise have been. It has been an honour and a pleasure.

My most profound thanks also go to my wonderful family for their support over the years and for rarely telling me how crazy I was. Finally, and most importantly, my thanks to Melanie Johnston-Hollitt, for showing me the world and for making it worth living in.

Alle baten helpen, zei de muis, en er piste en zee.

-Old Dutch proverb

Abstract

Length variations within the test masses of interferometric gravitational wave detectors limit their sensitivity over an important part of their operating frequency range. Even when perfectly decoupled from their environment, a test mass will exhibit fluctuations in its geometry due to a number of intrinsic noise sources. The best known of these noise sources is thermal noise, which had long been expected to limit detector sensitivity in the frequency range between about 10 Hz and 100 Hz. More recently it has become apparent that there is a menagerie of additional sources of test mass length noise. Arguably the most important of these additional noise sources is thermoelastic noise, which is caused by a non-zero coefficient of expansion coupling localised thermodynamic variations in test mass temperature into geometry changes. This noise source has particular importance because it is expected to be the dominant noise source in sapphire, a material being considered for use in gravitational wave detectors of the future.

To date, understanding of the family of length noise sources has been based on extensive theoretical work, with only limited experimental verification of the resultant models. If the models are to be employed with confidence to design future detector test masses, then more comprehensive verification is required. Such verification would ideally arise from the direct measurement to the spectra of the different noise sources for a wide range of test mass configurations.

This thesis describes the design and construction of an instrument intended to measure the length noise present in optical materials. The detection of thermoelastic noise in sapphire provides the initial motivation for the instrument, though the intent is to build a system capable of robustly characterising a variety of noise sources in multiple materials. Theoretical modelling is used to design a system that is expected to be dominated by thermoelastic noise.

The system measures length noise by monitoring the resonance frequency of a monolithic cavity with a laser, a novel technique that renders the system insensitive to environmental noise and to radiation pressure noise when compared to alternate techniques. A cylindrical optical substrate is prepared by polishing and coating mirrors onto its end faces and it is then suspended from a vibration isolation tower

using a niobium flexure. A stabilised non-planar ring oscillator (NPRO) Nd:YAG laser is then locked to the cavity and used to extract cavity length changes. Although an intrinsically low noise laser is used, additional noise reduction is required to meet the system noise specification. The laser is thus intensity stabilised by monitoring its output and frequency stabilised by locking the laser to a long reference cavity, yielding a system that will be capable of directly measuring the spectrum of sapphire thermoelastic noise for frequencies less than 400 Hz.

List of Symbols

Several symbols will be used repeatedly throughout this thesis to represent specific quantities or parameters. The following is a list of these symbols and short descriptions for the readers convenience. This list is not exhaustive but every effort has been made to maintain conformity in the symbols used here. Wherever possible standard symbols and notation have been used. Unless explicitly noted in the text, all quantities have been expressed using SI units.

α_l	Coefficient of linear thermal expansion
κ	Thermal conductivity
λ	Wavelength
ϕ	Loss angle
ρ	Density
σ	Poisson's ratio
σ_T	Variance in temperature
ω	Angular frequency
ω_c	Critical (or corner) angular frequency

c	Speed of light
C_{FB}	Finite Beam size constant for radiation pressure effects
C_{FTM}	Geometry constant for thermoelastic noise in a finite test mass
C_v	Specific heat with consant volume
D	Thermal Diffusivity
E_o	Young's modulus
f	Frequency
J	Low frequency thermolastic noise shape function
f_c	Critical (or corner) frequency
h	Plank's constant
$\hbar$	Plank's constant / 2π
I	Thermal noise beam geometry constant = 1.873 22 for a gaussian beam
k_B	Boltzmann's constant
l	Physical length
L	Optical path length
n	Refractive index
n_0	Refractive index of air at one atmosphere
N_0	Number of air molecules per unit volume at one atmosphere
P	Power
R	Test mass radius
r_0	Laser spot size
t_i	Interaction time
T	Temperature
$\bar{v}_o$	Mean air molecule velocity at one atmosphere

Contents

Chapter 1

Introduction

The development of narrow linewidth lasers has revolutionised the field of precision metrology, allowing previously unattainable levels of sensitivity to be reached. Such is the success of laser metrology that it is now being employed in an attempt to make the most sensitive measurement ever: the detection of gravitational radiation using interferometry. This measurement requires unprecedented performance, as the required strain sensitivity is equivalent to sensing the width of a human hair in the distance from Earth to the nearest star.

Initially, the performance of measurement systems was limited by technical noise or by environmental noise coupled into the systems. Such is the sensitivity of more recent designs that vibrations intrinsic to the materials used to build the devices have now become significant. The breathtaking requirements of advanced gravitational wave interferometry have uncovered a menagerie of unexpected vibration sources that are present in all optical media.

While some researchers have found it frustrating that these effects have surfaced to hamper their efforts, for others their emergence opens a new opportunity to develop fundamental materials science. Our nascent ability to measure intrinsic vibrations in different media will lead to an increased understanding of the internal mechanics of the materials and will allow the verification of largely untested theoretical models. The dominant noise sources will be different in each material as they exhibit a variety of dependencies on material size and shape, on the power and shape of the laser beam used to interrogate them, on temperature and on frequency.

To unravel these dependencies, an extensive battery of tests will be needed to explore a large parameter space for each material of interest. These tests will enable the confirmation or refinement of current theoretical models. The emergence of a picture of these noise sources promises to provide a fascinating window into the surprisingly exotic low-level vibrations in optical media. Only then will metrologists be able to confidently optimise future instrument designs and continue to improve instrumental sensitivity into the future.

This thesis describes the development of an instrument capable of directly measuring vibration spectra in optical materials. In the first instance, the instrument is intended to measure a noise source of immediate interest to the gravitational wave detection community. More generally though, it is intended that the system be versatile enough to undertake the more basic characterisation of different materials under a wide range of conditions to help develop a thorough understanding of intrinsic displacement noise in optical media.

1.1 Interferometric Gravitational Wave Detectors

The detection of gravitational waves remains one of the great outstanding endeavors of modern physics. These elusive ripples in space-time, first predicted by Einstein over half a century ago, are so weak that no verified direct detection has yet been made. The sensitivity required in gravitational wave detection is extreme, yet there is an expectation that with modern laser metrology techniques successful detection may be possible. Several teams around the world are now working to build giant laser interferometers capable of measuring the tiny strains required for unambiguous detection of gravitational waves.

Successful detection of gravitational waves will not only provide an important verification of general relativity, but will allow the development of a new branch of observational astronomy. A variety of possible gravitational wave signatures for different astrophysical objects have been predicted through numerical simulation [1,2]. The shapes of the expected signals are broadly determined by the class of object, but show features that depend on the details of the astrophysical interaction. Detection

of a gravitational wave signal will thus allow astrophysicists to constrain properties of the objects under study. This work is expected to help unlock some of the mysteries of a variety of astrophysical phenomena, such as black hole dynamics, neutron star collisions and supernovae. Beyond these predictions, it is the unexpected that is most exciting. As has historically been the case, the introduction of a new type of telescope can be expected to reveal a whole range of new objects in the universe.

Maximising the sensitivity of a gravitational wave detector is crucial to their success as astrophysical observatories. Greater sensitivity lends the instruments the ability to detect weaker signals, allowing detection of signals originating farther from the Earth. The volume of space accessible to a detector, and hence its event rate, increases as the cube of its effective range. Consequently, even minor increases of sensitivity lead to dramatic increases in expected detection frequency.

While a full description of an interferometric gravitational wave detector is well beyond the scope of this thesis, a general overview will enable us to appreciate the mechanisms that limit a detector's performance. A simplified overview of an interferometric gravitational wave detector is shown in figure 1.1. The detector consists of a Michelson interferometer that is augmented with optical cavities in each arm. The detector continuously compares the lengths of the two arm cavities. The mirrors in the arms are thus known as *test masses*, as they sample variation in the space-time metric.

Gravitational waves are quadrupole in nature, so they cause simultaneous contraction and expansion of space in perpendicular directions. There are two possible polarizations of gravitational radiation, denoted by the symbols + and x, which have directions of fluctuation oriented at 45° [3]. Any gravitational radiation signal can be decomposed into these two polarizations.

Extant interferometric detectors are sensitive to just one of these polarizations, so the degree to which an incident wave will couple with a detector will depend on its orientation. When radiation of the required polarization reaches a detector it causes expansion in one interferometer arm and contraction in the other. This leads to a difference in phase shift for the light travelling in each arm and thus disturbs the normal destructive interference at the interferometer output. The presence of

light at the interferometer output port thus indicates the passage of a gravitational wave.

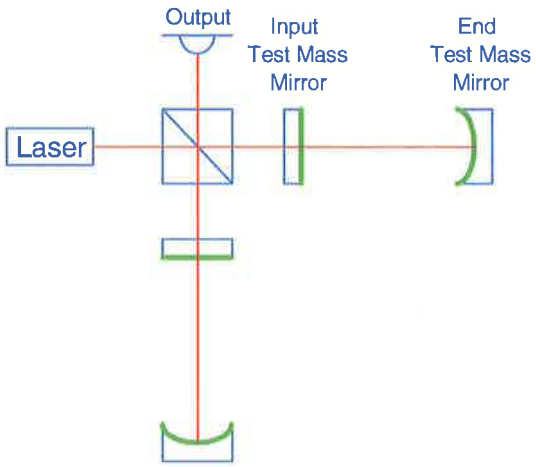

Figure 1.1: Schematic of a simple gravitational wave detector. Green lines indicate reflective surfaces.

Gravitational radiation is expected to display a range of temporal behaviours, depending on the source of the radiation [1,2]. This temporal variation of the gravitational radiation leads to time dependant reversals of the expansion and contraction, so the two arms will contract and expand in time with the incoming radiation. This results in a time dependant signal at the interferometer output. Because of limitations due to environmental and technical noise, ground based interferometers are most sensitive to gravitational radiation with frequencies between about 10 Hz and 1 kHz.

There are several classes of noise sources that can mask a gravitational wave signal by producing a signal at the interferometer output:

- Environmental vibrations that are coupled into the interferometer,
- Optical or technical noise from the laser and the various electronic systems that run the detector,
- Intrinsic vibrations produced within the test masses of the interferometer.

Figure 1.2 taken from [4] shows an estimate of the sensitivity of the Advanced LIGO detector, with several contributions to detector noise shown. Noise sources

form each of the above categories are included; seismic noise is an external noise sources coupled into the instrument, optical noise arises from the optical measurement system and substrate and coating thermal noise in the test mass are intrinsic to the test masses themselves. The figure highlights the role that test mass thermal noise plays in determining detector sensitivity in the region between 20 Hz and 400 Hz.

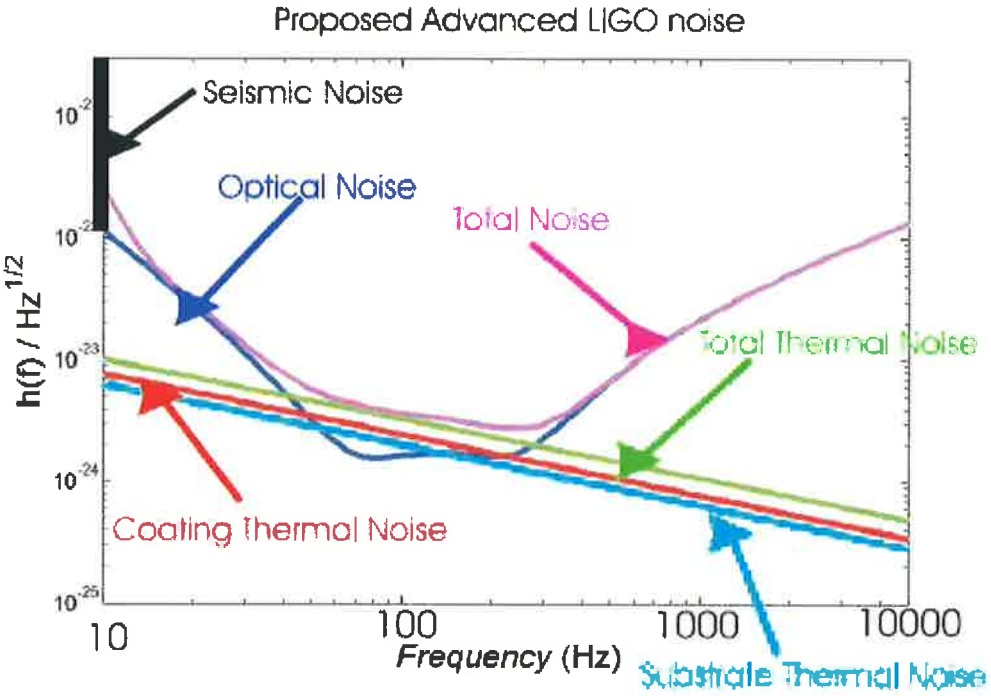

Figure 1.2: An estimate of the sensitivity of Advanced LIGO.

Noise that is intrinsic to the test masses themselves is expected to dominate interferometer sensitivity in the middle of their detection band; the region where they are most sensitive. Any improvement in detector sensitivity in this frequency band will result in an increase in the detectors effective range and can thus can be expected to produce significant increases in the instruments science output. The need to increase out understanding of test mass intrinsic noise provides the motivation for this thesis, which seeks to measure the spectrum of test mass length noise directly.

Consider a single arm cavity of an interferometer, as shown in figure 1.3. Ideally

the arm cavity would monitor the distance between two test points in space, such as is defined by the separation of the mirror's centres of mass, ΔL_{cm}. However, the test masses have finite size, so in practice we are limited to measuring the distance between their reflective surfaces rather than between their centres of mass. Any noise source which alters the thickness of the test masses will thus lead to variation of the cavity length and hence an undesired signal at the interferometer output.

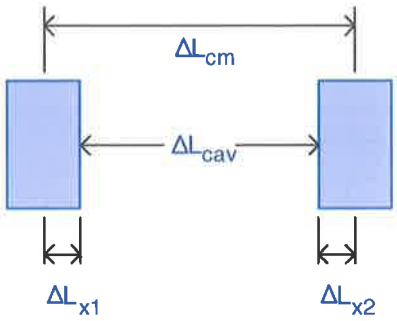

Figure 1.3: Diagram illustrating the relationship between the different sources of physical length noise and the cavity length variation.

Reduction in the magnitude of the test mass length noise sources ΔL_{x1} and ΔL_{x2} is thus an important challenge if the sensitivity of interferometric gravitational wave detectors is to be improved. Such sources of noise cannot be reduced by any form of isolation, but must instead must be managed by an informed choice of test mass material, geometry and coatings.

1.2 Sources of Test Mass Length Noise

Unfortunately, there are a number of different effects that can cause length noise in an interferometer test mass. One particular noise source, thermal noise, has long been recognised as having importance in determining the sensitivity of gravitational wave detectors, but more recently it has been recognised that there are several more exotic noise sources that are also important. One of these, thermoelastic noise, is sufficiently large that it can dominate over thermal noise for some materials.

1.2.1 Thermal Noise

The study of thermal noise began as a result of the discovery of Brownian motion and was initially studied in electrical systems [5,6]. Because of this early history, thermal noise is sometimes referred to as *Nyquist* or *Johnson* noise. Thermal noise occurs because the molecules making up a test mass at finite temperature are always subject to random thermal motion. The macroscopic consequences of the random molecular motion in a body will be influenced by the geometry and properties of the body. Molecular action that is well suited to the body will be sustained and reinforced, while other motion will not very effective at producing macroscopic motion. Thus we can regard thermal noise as a stochastic noise source that drives the mechanical admittance of the material body in which it is present. Following [7], we can write an expression for the mechanical length noise present in a body as the product of its mechanical admittance and its thermal excitation.

$$S_{x,Th}(f) = \frac{k_B T}{\pi^2 f^2} \Re e \{Y(f)\} \tag{1.1}$$

where
- k_B is Boltzmann's constant
- T is the temperature of the material
- $Y(f)$ is the spectrum of the objects mechanical admittance

Thermal noise in the context of resonant bar gravitational wave detectors has been well studied [8] (and references therein). This work showed that the shape of the resultant displacement noise was found to be dependant on the resonance structure of the body in question and its damping mechanism [9]. Thermal noise in mirror substrates has been also been studied extensively for application in gravitational wave research - see for example [10] and [7]. In the later reference the spectrum of thermal length noise below the first longitudinal test mass resonance as measured by a Gaussian probe beam is given as

$$S_{x,Th}(f) = 4k_B T \frac{1-\sigma^2}{\pi^3 E_o r_o} I\phi \times \frac{1}{f} \tag{1.2}$$

where
- σ is Poisson's ratio for the material
- E_0 is Young's modulus of the material
- I is equal to 1.87322
- ϕ is the loss angle of the material
- r_0 is the spot size of the probe laser beam

The constant I arises from the Gaussian beam used to sample the surface position of the test mass. The shape of the spectrum below test mass resonance is determined by the behaviour of the loss angle, ϕ [9]. If the material exhibits viscous damping then ϕ is proportional to frequency and consequently the spectrum of thermal noise is independent of frequency. It the dominant material damping is structural damping then ϕ is constant and the thermal noise spectrum is proportional to $1/f$ in this region.

The analysis above describes the thermal noise produced below mechanical resonance of the body, a situation which is appropriate for discussion of test mass thermal noise as the resonances in question are above the detection band for gravitational waves. The equation shows that the use of intrinsically low loss materials is preferable, as are fabrication and mounting techniques that do not increase losses.

1.2.2 Thermoelastic Noise

While thermal noise was always seen as a problem for gravitational wave interferometers, the same cannot be said of thermoelastic noise. Thermoelastic noise has long been known to the gravitational wave community, where it was important in the context of bar detectors [9, 11]. However, its importance to interferometers was not appreciated until a landmark paper by Braginsky [12]. At that time, sapphire was being considered as a replacement material in the test masses of future detectors because of its extremely low loss, which results in sapphire having a lower level of thermal noise than fused silica. A comparison made between the two materials in 1999 is shown in figure 1.2, which is taken from [13]. However, [12] revealed that sapphire has a higher level of thermoelastic noise than silica.

Thermoelastic noise is the fluctuation associated with thermoelastic damping in a material. Regions inside every material undergo statistical variations in tempera-

ture. In materials having non-zero thermal expansion coefficients, these fluctuations produce localised volume changes. As this volume change is established, heat flows from the warmer regions to the cooler to balance the thermal gradient. This heat is then lost, in the sense that it does not return to its original location when the material resumes its previous volume. Thermoelastic noise will be discussed in more detail in section 2.2.1, where a mathematical description of thermoelastic noise is presented.

1.3 Methods for Measuring Intrinsic Length Noise

1.3.1 Measurement using an Interferometer

Since long-baseline interferometers are limited by test mass intrinsic length noise over part of their frequency band, it seems that they are a logical measurement system for such noise sources. However, long baseline interferometers do not make good candidates for an extensive noise study because they lack the flexibility to tailor parameters of the optical system such as laser spot size. They are also ill-suited to frequent changes of the test masses to test different materials and cavity configurations.

One obvious alternative is to measure length noise with a dedicated, laboratory scale interferometer. Such an interferometer could be optimised for the measurement task by judicious choice of spot sizes and cavity lengths.

Using an interferometer has the great advantage that many noise sources that could limit the sensitivity of the experiment will be common mode to the two arms. As an example, laser frequency noise will be common mode to both interferometer arms and therefore the system should be first order insensitive to such variations. Ideally, laser intensity noise and low frequency seismic noise are also common mode to the two arm cavities and therefore their effects should be relatively small in an interferometer. However, the benefit of this common mode rejection is limited by the achievable matching of the two arms. For example, if the seismic isolation systems of the two arms cavities were matched to 1 part in 100, then common mode rejection would only reduce the seismic noise level by a factor of 100. Difficulties

in matching the seismic isolation systems and the cavity optics requires that efforts still be made to significantly reduce these noise sources. While the use of an interferometer certainly helps by relaxing the requirements for the noise mitigation systems, it does so at the cost of increased optical and electronic complexity.

Nevertheless, several experiments are underway to measure length noise using laboratory scale interferometers. An interferometer based experiment at the University of Tokyo has recently produced first results [14]. This project has measured thermal noise in BK7 glass, thermoelastic noise in CaF_2 and coating thermal noise on fused silica substrates. The sensitivity of the experiment is $1{\times}10^{-15}\ m/\sqrt{Hz}$ at 100 Hz and improves at higher frequencies. Sensitivity at frequencies lower than 100 Hz is limited by seismic noise and by technical noise in the system electronics. An earlier experiment at the University of Tokyo measured thermal noise in a cantilever in the frequency range between 50 Hz and 500 Hz [15] using interferometric techniques.

The Thermal Noise Interferometer (TNI) at the California Institute of Technology was specifically designed to measure length noise in fused silica and sapphire substrates. The TNI is also being used to measure length noise in test mass coatings. Length noise spectra in fused silica and sapphire have been reported by the TNI team in [16] and [17] for fused silica and sapphire respectively. The TNI has also been used to measure test mass coating noise [18, 19]. The current noise floor of the TNI is approximately $1{\times}10^{-15}\ m/\sqrt{Hz}$ at 100 Hz, dropping to the order of $1{\times}10^{-18}\ m/\sqrt{Hz}$ beyond around 500 Hz. Figure 1.4 shows the noise floor of the TNI and is taken from [20]. As can be seen the sensitivity rapidly degrades below 100 Hz, presumably because of seismic noise contamination.

A length noise experiment using the $40m$ prototype interferometer at Caltech has also been proposed [21]. This experiment would have the advantage of more closely mimicking the conditions in a real interferometer. Most notably, the longer arm cavities allow the use of larger spot sizes than benchtop scale interferometers and could therefore access a part of parameter space that is arguably more relevant to full scale interferometers.

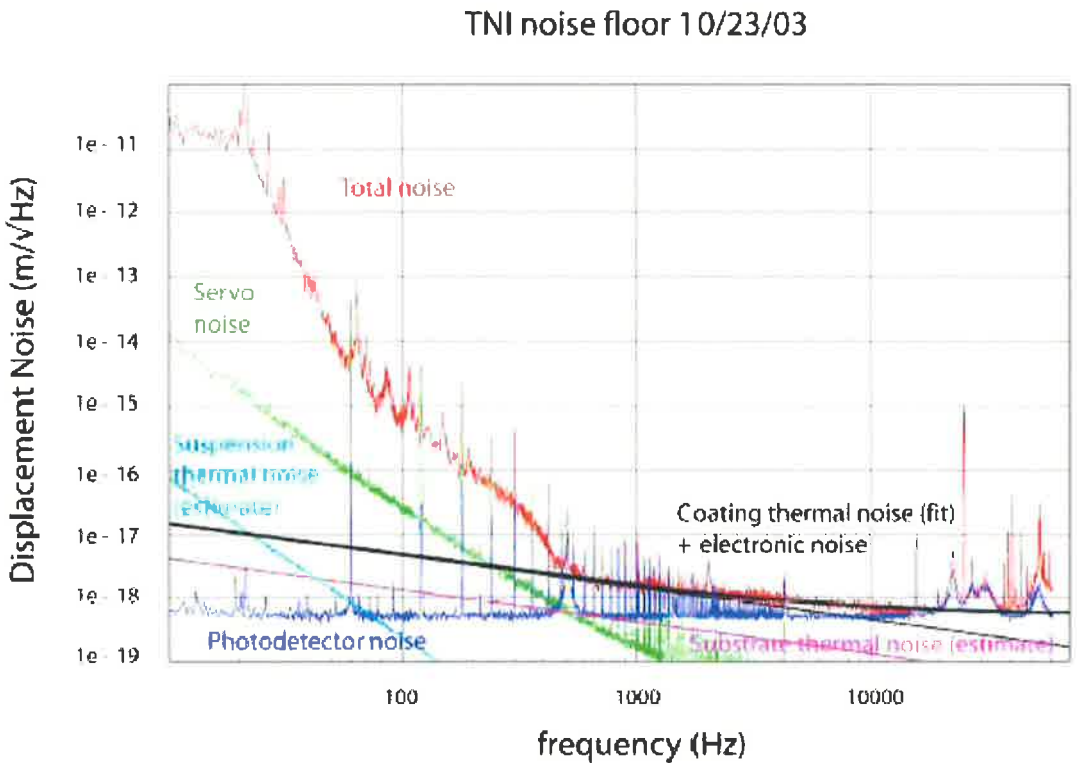

Figure 1.4: Noise floor of the TNI with contributions shown.

1.3.2 Measurement using a Fabry-Perot

A Fabry-Perot cavity has a resonant frequency that depends on the length of the cavity. Thus, monitoring of the cavity resonant frequency allows information about cavity length noise to be inferred. A Fabry-Perot system is notionally simpler than a Michelson interferometer, but it lacks the interferometer's common mode rejection of laser frequency noise.

In a Fabry Perot cavity, two (or more) mirrors are arranged so that light is reflected back upon itself in a closed path. When the phase change due to the light's round trip path length in the cavity is a multiple of 2π radians, the light arrives back at its starting point in phase. This leads to perfect constructive interference of the light; the light under these conditions is said to be resonant in the cavity. The

resonant frequencies are those that satisfy the expression

$$f_r = \frac{mc}{2L} \tag{1.3}$$

where
- f_r is the frequency of the resonant mode
- m is a (large) integer
- c is the speed of light
- L is the optical length of the cavity

If the laser frequency is close, but not equal to, the ideal value needed to acquire a 2π phase shift during a round trip then it will still be supported by the cavity, but the constructive interference will not be as strong. A complete analysis [22] yields cavity transmission and reflection equations of

$$\begin{aligned} T_{cav} &= \frac{I_t}{I_{inc}} = \frac{(1-R)^2}{(1-R)^2 + 4R\sin^2(\delta\phi/2)}, \\ R_{cav} &= \frac{I_r}{I_{inc}} = \frac{4R\sin^2(\delta\phi/2)}{(1-R)^2 + 4R\sin^2(\delta\phi/2)} \end{aligned} \tag{1.4}$$

where
- T_{cav} is the cavity (power) transmission coefficient
- R_{cav} is the cavity (power) reflection coefficient
- I_{inc} is the power incident on the cavity
- I_t is the power transmitted through the cavity
- I_r is the power reflected from the cavity
- R is the (power) reflectivity of the cavity mirrors
- $\delta\phi$ is the round trip phase shift

It can be seen from these expressions that when $\delta\phi$ is equal to $m2\pi$, the transmission is maximised, so we obtain $T = 1$ and $R = 0$. Thus, when incoming light is at a resonant frequency of the cavity, the transmission of the cavity is a maximum. As the frequency of incoming light moves away from a resonance condition the transmission through the cavity changes. This change in reflection and transmission properties of the cavity allows the relative frequency difference between the incoming light and cavity resonance to be detected. Detailed schemes for perform-

ing this measurement will be presented in chapter 4; for now we will assume the existence of such methods for measuring the resonant frequency of the cavity.

The change in cavity resonant frequency that arises due to a change in length can be determined by differentiating equation 1.3 with respect to the path length, L, giving

$$\begin{aligned} \frac{\partial f_r}{\partial L} &= -\frac{mc}{2n}\frac{1}{L^2} \\ &= -f_r\frac{1}{L} \\ ie. \frac{\partial f_r}{f_r} &= -\frac{\partial L}{L} \end{aligned} \tag{1.5}$$

Thus, we can say that a source of length noise having spectrum $S_x^{1/2}$ will cause resonant frequency noise with spectrum

$$\begin{aligned} S_f^{1/2} &= \frac{f_r}{L}S_x^{1/2} \\ S_f^{1/2} &= \frac{c}{\lambda L}S_x^{1/2} \end{aligned} \tag{1.6}$$

While equations 1.5 and 1.6 explain how cavity length noise leads to variations in resonant frequency, they also reveal the drawback to the approach, namely that variations in laser frequency are indistinguishable from changes in cavity length. Essentially, using a Fabry-Perot to measure length noise is rather like making a measurement of the cavity length with a ruler marked in units of wavelengths of light. If the ruler stretches or shrinks then the result of the measurement changes also. Thus, if the scheme is to be effective then the probe laser must have good frequency stability. This situation can be contrasted with the case of an interferometer, where variations in laser frequency are common mode to both arms and hence cancel to some degree in the interferometer output.

One standard method for reducing the frequency noise in a laser is to phase-lock it to a stable reference cavity. If we assume that the laser can be locked to a reference cavity perfectly, then the frequency stability imparted to the laser is given

by equation 1.5, where S_x in this case is the length noise of the reference cavity. Key to the operation of this topology is that the reference cavity be much longer than the test cavity so that any reference cavity length noise induces a laser frequency shift that is smaller than that caused by length noise in the test cavity [A. Giazotto, personal communication] [23].

Consider a length change δL_{ref} in the reference cavity. According to equation 1.5 this will produce a change in probe laser frequency given by

$$\delta f_{probe} = f\frac{\delta L_{ref}}{L_{ref}}.$$

Similarly a change in test cavity length will lead to a change in its resonant frequency given by

$$\delta f_{test} = f\frac{\delta L_{test}}{L_{test}}.$$

The smallest length noise that could be unambiguously detected in the test cavity will be that which causes a frequency noise fluctuation equal to the reference cavity variation:

$$\begin{aligned} \delta f_{test} &= \delta f_{probe} \\ f\frac{\delta L_{test}}{L_{test}} &= f\frac{\delta L_{ref}}{L_{ref}} \qquad (1.7) \\ \delta L_{test} &= \frac{L_{test}}{L_{ref}}\delta L_{ref} \qquad (1.8) \end{aligned}$$

Thus we see that the significance of reference cavity length noise is reduced by a factor of the ratio of the two cavity lengths and thus the reference cavity need not be constructed from a material having losses as low as the test cavity. The reference cavity length noise, δL_{ref}, will be largest in regions near mechanical resonances of the reference cavity and in these regions it is possible for it to dominate the length noise of the test cavity. Consequently the acoustic Q of the reference cavity must be high to ensure that these resonances are narrow and thus contaminate a relatively small part of the measurement system's frequency range.

This technique is being employed by several groups for the measurement of length

noise. A project at the Universität Hannover is using this topology to measure thermal noise in a test cavity consisting of two suspended mirrors and to study radiation pressure induced length noise [24, 25]. A group from the Australian National University has also employed this topology to measure thermal noise in a flexure on which one of the cavity mirrors is mounted [26, 27]. Very recently a group at the VIRGO observatory has used this technique to measure the displacement noise of a 10 mm cavity suspended from the VIRGO seismic filter and shown the noise to be thermal in origin above 3 Hz [28].

1.4 Monolithic Test Cavity Measurement System

In this thesis we describe a novel extension of the Fabry-Perot approach to measuring length noise. We aim to measure length noise in a single test cavity and use a long reference cavity to stabilise the laser frequency. However, in our system the test cavity is formed from a single piece of material which has reflective coatings on *both* faces. With the new approach the optical length of a single piece of material is probed rather than the distance between two separate test masses. At frequencies below the first mechanical resonance of the test cavity the two cavity faces move together when the mass is excited seismically or by radiation pressure. This common mode behaviour ensures that seismic noise cannot produce length noise in the test cavity. It should be noted that the common-mode behaviour is inherent in the design and takes no "effort" to achieve. This can be contrasted with the case of an interferometer where the two arms must be carefully manufactured to be matched if optimal common mode rejection is to be achieved. Thus a monolithic test cavity retains the optical simplicity of a Fabry-Perot cavity, while having excellent common-mode rejection sought by a Michelson interferometer.

The advantages and consequences of using a monolithic test cavity are discussed further below.

1.4.1 Reduction of Seismic Noise Sensitivity

The influence of seismic noise is common mode to the two faces of a monolithic test cavity. The efficiency of coupling of seismic noise into a monolithic cavity is thus much smaller than it would be if the two cavity mirrors were suspended separately. This consideration was very important in the design of this experiment due to the limited space available for the construction of a seismic isolation system. This experiment was designed to fit in a pre-existing vacuum chamber (see appendix C), which could not house a seismic isolation system of sufficient size to make any other topology feasible. The choice of a monolithic cavity greatly reduced the complexity of the required isolation system, as will be discussed in chapter 3.

1.4.2 Reduction of Radiation Pressure Noise Sensitivity

Variations in the intensity of the probe laser lead to variations in the photon pressure experienced by the mirrors of an optical cavity. In a system where the mirrors are suspended on high admittance suspension systems, this radiation pressure noise can move the mirrors relatively easily, resulting in a source of length noise. When a monolithic test cavity is used, the radiation pressure noise must strain the crystal if it is to produce length noise. In general, this is much harder to do than it is to move a suspended optic, so the length changes will be much smaller in a monolithic cavity. As a consequence we would expect that the requirements on laser intensity noise will be significantly smaller because of the use of a monolithic cavity.

1.4.3 Common-mode rejection of Suspension Noise

One potential problem in the measurement of intrinsic length noise is the possibility of mechanical noise from elsewhere in the system coupling into the motion of the cavity mirrors. Consider, for example, thermal noise in the suspension system that holds independent test masses in a cavity. Thermal noise in the two suspensions would be uncorrelated and therefore result in a change in cavity length.

The use of a monolithic cavity ensures that only length noise produced *inside* the test mass will be measured. All other noise sources will be common mode to the

two cavity surfaces and hence not able to change the cavity length. Thus we can be confident that the results of this experiment will reflect length noise intrinsic to the test cavity material.

1.4.4 Convenience in Extensive Trials

If intrinsic length noise sources are to be studied in detail it would be advantageous to measure the noise spectra in a large number of different test cavities. One could imagine conducting a series of test with differing materials, coating and laser spot sizes for example. The use of a monolithic test cavity means that to run a test on a new material only one optic would need to be remounted and aligned. This would lead to considerable efficiencies if a large number of samples were to be tested.

1.4.5 Inability to Tune Cavity Resonance

One potential disadvantage of using a monolithic cavity is that it is very hard to change its length. This means that it is difficult to tune the resonance frequency of the cavity during operation of the experiment, which removes one degree of freedom from the design process. We will see in chapters 4 and 5 that this restriction causes some practical complications in the operation of the completed experiment.

1.4.6 Presence of Refractive Index Noise

Light inside a monolithic cavity is continually passing through optical material. In general, we expect this material to show refractive index noise, which will lead to noise in the optical path length of the cavity. This is of course, indistinguishable from genuine changes in physical cavity length. Thus the use of a monolithic cavity allows an additional noise source that would not be present were a conventional cavity (or an interferometer) used. However, this noise source is of concern only if it is sufficiently large to mask the desired noise signal. Thermorefractive length noise will thus be treated as a noise source and will be discussed in section 2.3.2.2. It should be noted however, that refractive noise could be made observable with the instrument described in this thesis by judicious choice of operating conditions. This

possibility allows for the future measurement of this noise source.

1.5 Thesis Outline

The project consists of four main phases, each of which was to be completed before the next phase commenced.

1. Initial system design,
2. Bench-top testing,
3. Migration to vacuum and
4. Data collection.

This thesis primarily describes the first two phases of the project, although some preparations for the third phase are also described.

In this introductory chapter we have provided an overview of the problem at hand, namely that intrinsic length noise in optical media can limit the sensitivity of ultra-high precision metrology. We have motivated the study through a description of noise sources of particular importance to interferometric gravitational wave detectors, but the problem is of wider importance. This chapter has also included the description of the ways in which we might set out to measure length noise in optical media and discussed the advantages and disadvantages of each. The basic premise of this project described herein was to use a monolithic test cavity. The advantages and consequences of this topology were briefly discussed. Most notably, the use of a monolithic cavity vastly reduces the effects of several otherwise troublesome noise sources. A second major advantage of the monolithic cavity is that it allows us to confidently identify the measured noise as being within the cavity itself; we do not have to worry about other forms of noise being coupled in through suspension systems.

In chapter 2 we turn to more detailed analysis of the noise levels that we might expect to see in a laboratory scale experiment. The theoretical prediction for the spectrum of thermoelastic noise is discussed in some detail and as a result the

required sensitivity of the experiment is determined. The remainder of the chapter then deals with an analysis of the magnitude of various noise sources that might mask the sought-after thermoelastic signal . We show in this chapter that a sufficiently well stabilised laser should be able to detect thermoelastic noise in sapphire over the frequency band between 10 Hz and 1 kHz, the range where this noise source is most important in full scale gravitational wave detectors.

Chapter 3 is a discussion of the seismic isolation system developed for the experiment. The design and construction of the various components of the system are described, including the seismic preisolator, the suspension system that holds the test cavity and its damping system. Measurements to show the effectiveness of the setup are then presented.

The optical and electronic components of the measurement system are described in chapter 4. After an overview of the overall experimental topology, the chapter opens with a description of the subsystem used to lock the laser to the reference cavity. The design and construction of the reference cavity are discussed and it is then characterised. The Pound-Drever-Hall sensor used to measure frequency differences between the laser and the cavity resonance is then presented. The two available frequency actuators on the laser are then characterised before finally a compensator is designed to close a feedback loop around the laser stabilisation subsystem. We examine the performance of this locking system and show that it will be of sufficient quality to meet the required sensitivity.

We then turn our attention to the system used to read out the length noise in the test cavity. This necessarily begins with a description of the optical design of the test cavity. A dummy cavity was used so that experiments could commence before the sapphire test cavities had been manufactured. The design and performance of this dummy cavity are therefore also outlined. The feedback system used to lock the frequency of the reference and thus that of the probe laser to the test cavity at low frequencies is then discussed, including presentation of the fringe side locking sensor, the available actuators and the required compensator. From this discussion it emerges that intensity stabilisation of the laser will be required, a subject that is discussed more fully later in the chapter. Two practical difficulties arose during

the testing of this subsystem of the experiment, one with mode-matching the test cavity, the other with locking the stabilised laser to the test cavity. The reasons for these problems are discussed and their ramifications for the project are discussed in detail.

To finish chapter 4, the intensity stabilisation scheme used in the experiment is described. This begins with a brief discussion of the sensor and then the available actuators are characterised. The chosen actuator is found to be non-linear, a result that makes the design of the compensator slightly more complicated than it would otherwise be. The resultant compensator is then discussed. Finally, the performance of the intensity locking loop is presented and it is found to meet the sensitivity requirements.

In chapter 5, initial results taken with the sapphire test cavity are presented. While this data is very preliminary, it does show that the experiment is already within an order of magnitude of its required sensitivity.

Chapter 6 returns to the problems uncovered in setting up the sapphire test cavity readout scheme and discusses ways in which they might be circumvented. A supplementary technique to temporarily allow locking of the system during mode matching is first described. This scheme has poor noise performance, so is used only during mode matching. It is, however, expected to allow modematching of the test cavity to be performed quickly and efficiently.

Attention then turns to the problem of effectively acquiring lock between the stabilised light and the test cavity. A number of alternative strategies are discussed and it emerges that thermal tuning of the reference cavity is the most feasible method. Consequently the thermal properties of the reference cavity are modelled using a finite element package and the results are presented. The chapter then concludes by presenting the design of an auto-locking system to automatically move the reference cavity and laser into resonance with the test cavity resonance and then lock the two together.

Chapter 7 concludes the body of the thesis by summarising the work completed so far and outlining the work still to be done.

Chapter 2

Physical Noise Sources

2.1 Introduction

Having established our overall experimental topology in chapter 1 we are now able to develop the experimental design more fully.

In this chapter we will first estimate the expected level of thermoelastic noise in the sapphire test cavity (2.2.1). This will enable us to decide on a test cavity geometry and to arrive at the design sensitivity in section 2.2.2.

In section 2.3, we examine the range of noise sources that might compromise our ability to measure the thermoelastic noise. In section 2.3.2, we first examine noise sources that are present in the test cavity. This analysis will show that these noise sources are smaller than thermoelastic noise between 1 Hz and 1 kHz and will thus not prove troublesome.

In section 2.3.3, we will show that the free-running laser frequency noise is too high to allow the experiment to reach the design sensitivity. The laser will therefore need to be stabilised using a reference cavity. In this chapter we will assume that the stabilisation scheme does not produce technical noise but we will analyse noise sources in the reference cavity. Technical noise will be analysed in chapter 4. Section 2.3.4 will show that the stabilised laser light will be sufficiently stable to allow the design sensitivity to be reached in the final phase of this experiment, although the sensitivity of the preliminary experiment described in this thesis is expected to be limited by acoustic excitation and gas damping of the reference cavity and by gas

pressure fluctuations inside the reference cavity.

2.2 Sensitivity Specification

2.2.1 Expected Thermoelastic Noise Spectrum

Thermoelastic damping is a well known loss mechanism in mechanical systems [9, 11]. Thermoelastic damping occurs because mechanical strain and temperature are coupled through a non-zero thermal expansion coefficient. Thermal gradients are set up whenever a material is deformed. In the presence of thermal gradients heat flows from the warmer regions to the cooler. This process is irreversible at low frequencies, because once the energy has relocated it is unavailable for the restoration of the material to its original form. This loss of energy leads to damping of any mechanical strain in the material.

The fluctuation-dissipation theorem states that for every source of dissipation there will be a corresponding fluctuation [29]. Thus, there is a noise source associated with thermoelastic damping and it is known as thermoelastic noise. This source of noise and its role in gravitational wave detection was first discussed by Braginsky et al. in [12] and it was rapidly realised that thermoelastic noise could prove an important barrier to achieving ultimate sensitivity in gravitational wave detectors.

The temperature in a volume of material exhibits random statistical fluctuations having variance [30]

$$\sigma_T^2 = \frac{k_B T}{\rho C_v V},$$

where,
- σ_T^2 is the variance in temperature,
- k_B is Boltzmann's constant,
- T is temperature,
- ρ is density,
- C_v is specific heat at constant volume,
- V is volume.

These temperature fluctuations couple into strain in the material through the

thermal expansion coefficient, producing volume fluctuations in the material.

It is convenient to regard a macroscopic body as being made up of an ensemble of isolated thermal domains, each undergoing random volume fluctuations. These independent domains have a size that is determined by the thermal properties of the material in question and the observing frequency; if there is insufficient time for heat to diffuse from one cell to another then they are thermally distinct. Averaging over a large number of cells should lead to a cancellation of the noise source, so we should expect that the noise would diminish if we looked at a larger volume or a larger area of the body's surface. This intuition is borne out by analysis; [12] provides a thermoelastic displacement noise spectrum, $S^{*}_{TE}(\omega)$, at the surface of the material:

$$S^{*}_{TE}(\omega) = \frac{8(1+\sigma)^2\alpha_l^2\kappa k_B T^2}{\sqrt{2\pi}C_v^2\rho^2 r_o^3} \times \frac{1}{\omega^2} \tag{2.1}$$

where
- σ is Poisson's ratio for the material,
- κ is the thermal conductivity of the material,
- α_l is the linear thermal expansivity of the material,
- r_o is the spot size of the probe laser beam.

The analysis of [12] assumed infinite test masses. In [31] the treatment was extended to account for finite sized test masses. This was done by introducing the scaling factor, C_{FTM}, which is of order unity:

$$S^{**}_{TE}(\omega) = C^2_{FTM} S^{*}_{TE}(\omega) \tag{2.2}$$

Because of the complexity of the analytic expression for C^2_{FTM}, it is presented in a graphical form in [31].

In the analysis of Cerdonio et al. [32] the analysis was extended further to include the non-adiabatic case:

$$S^{***}_{TE}(\omega) = \frac{8(1+\sigma)^2\alpha_l^2\kappa k_B T^2}{\sqrt{2\pi}C_v^2\rho^2 r_0^3} \times \frac{J(\omega)}{\omega_c^2} \tag{2.3}$$

$$\text{where } \omega_c = \frac{1}{2\pi}\frac{\kappa}{\rho C r_o^2}$$

$$\text{and, } J(\omega) = \left(\sqrt{\frac{2}{\pi}}\right)\int_0^\infty du \int_{-\infty}^\infty dv \left\{ \frac{u^3 e^{\frac{-u^2}{2}}}{(u^2+v^2)[(u^2+v^2)^2 + \left(\frac{\omega}{\omega_c}\right)^2]} \right\}$$

This spectrum is identical with the Braginsky expression for frequencies above ω_c. At low frequencies however, the spectrum has a slope of approximately $1/\sqrt{f}$ rather than $1/f^2$. This approximation to the low frequency slope will be used throughout the remainder of this thesis.

Numata et al. corrected the $\left(\frac{2}{\pi}\right)$constant in the first pair of brackets above to $\left(\frac{\sqrt{2}}{\pi^{3/2}}\right)$ [14]. Although the source of the error was not identified in [14], it was due to a factor of 2π having been missed during integration in the Cerdonio analysis. Derivation of this step is performed in appendix A to confirm the Numata result.

We therefore arrive at a full expression for thermoelastic noise;

$$S_{TE}(\omega) = C_{FTM}^2 \times \frac{8(1+\sigma)^2 \alpha_t^2 \kappa k_B T^2}{\sqrt{2\pi} C_v^2 \rho^2 r_0^3} \times \frac{J(\omega)}{\omega_c^2} \tag{2.4}$$

with

$$J(\omega) = \left(\frac{\sqrt{2}}{\pi^{3/2}}\right)\int_0^\infty du \int_{-\infty}^\infty dv \left\{ \frac{u^3 e^{\frac{-u^2}{2}}}{(u^2+v^2)[(u^2+v^2)^2 + \left(\frac{\omega}{\omega_c}\right)^2]} \right\}$$

and C_{FTM}^2 as in [31].

2.2.1.1 Dependence of Thermoelastic Noise Level on Test Cavity Design

Measurement of the thermoelastic noise will be simplified if we maximise its amplitude in the test cavity. Examination of equation 2.4 shows that we can influence the level by choosing an appropriate geometry for the test cavity. Changes in the size and shape of the cavity lead to variation of the C_{FTM} parameter and also move the critical frequency where the slope of the noise spectrum changes. Changing the

curvature of the ends of the test cavity will change the level of noise by altering r_0. The choice of test cavity geometry is restricted by a number of practical constraints, however:

- The test cavity must be short relative to the reference cavity as was discussed in chapter 1;
- The test cavity must have a diameter large enough to allow it to be mounted on its suspension without introducing excessive loss. Chapter 3 discusses the test mass suspension further;
- The test cavity must include a mirror with "reasonable" curvature so that alignment of the probe laser into the cavity is not too difficult. Cavities having close to flat mirrors at both ends are more sensitive to misalignment and are therefore more difficult to deal with in practice.

The thermoelastic noise level depends on the test cavity diameter through variations in the C^2_{FTM} parameter. However, for realistic cavity designs, C_{FTM} is not very strongly influenced by the diameter of the mirror [31]. A test cavity diameter of 50 mm was therefore selected as a compromise between suspension performance and cost.

Variation of the expected thermoelastic noise in sapphire as the spot size is varied is shown in figure 2.1. The material constants required in the thermoelastic noise equation were obtained from [12] and are listed in Table 2.1. $C^2_{FTM} = 1$ was used in the graph, although the final design is likely to have a slightly different value. The low frequency approximation of $1/\sqrt{f}$ dependence described in [32] was used to produce the portion of the non-adiabatic curves below the corner frequency.

The spot size is dependant on both cavity length and cavity mirror curvature. We are constrained to choose a short test cavity so that it remains short compared to the reference cavity. However, we must make the test cavity at least 10 mm long to allow for its suspension. In addition we would like the test cavity length to be comparable to its radius so that the C^2_{FTM} factor remains close to unity. A choice of test mass length in the range 20 - 30 mm seems a reasonable compromise

Symbol	Property	Value	Units
σ	Poisson's ratio	0.29	
E_o	Young's modulus	400×10^9	Pa
κ	Thermal conductivity	40	$Wm^{-1}K^{-1}$
α_l	Thermal expansivity	5.0×10^{-6}	K^{-1}
C_v	Specific heat	790	$Jkg^{-1}K^{-1}$
ρ	Density	4×10^3	kgm^{-3}

Table 2.1: Sapphire material properties

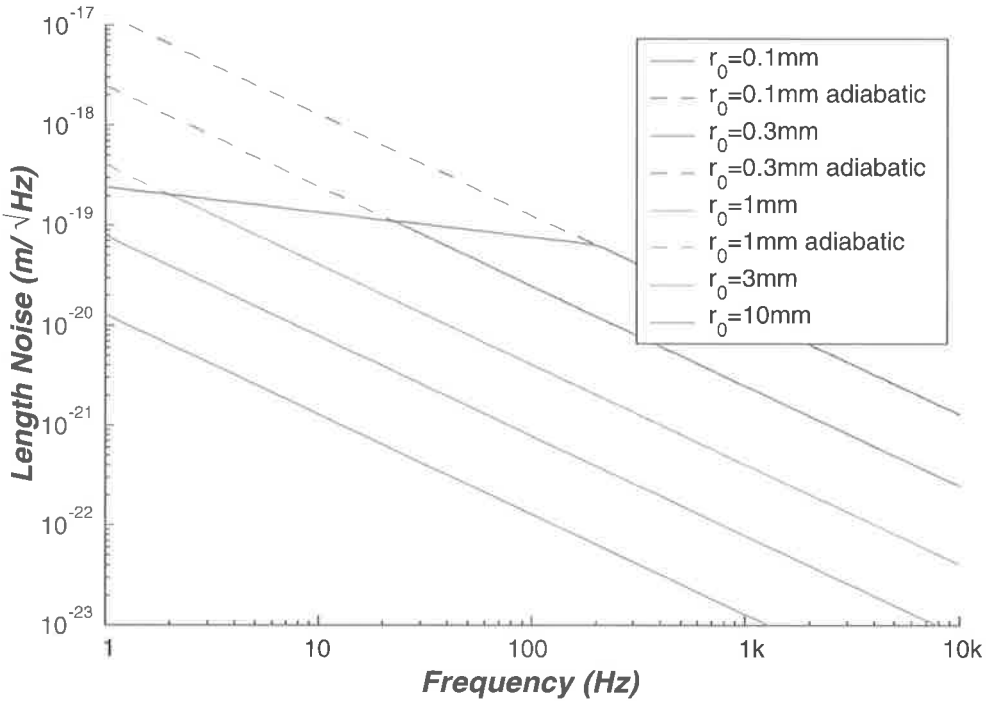

Figure 2.1: Variation in the thermoelastic noise spectrum as the spot size is varied. The adiabatic approximation to each curve is shown as a dashed line.

between these factors. As shown in figure 2.1, we should choose a small spot size to increase the thermoelastic noise level. A small spot size will also place the corner of the thermoelastic noise spectrum within the frequency region that we will study, allowing this important feature of the spectrum to be verified.

A limited range of spot sizes was feasible for our short cavity, as is shown in figure 2.2 which displays the spot size variation with mirror curvature for cavities 20 mm and 30 mm long. While the spot size is controllable through mirror curvature, we can only vary the spot size between approximately 0.15 and 0.25 mm for reasonable values of mirror curvature. We elected to use a test cavity having a length of 20 mm and a mirror curvature of 250 mm, resulting in a laser mode with radius 172

μm at the flat input coupler and 185 μm at the back mirror.

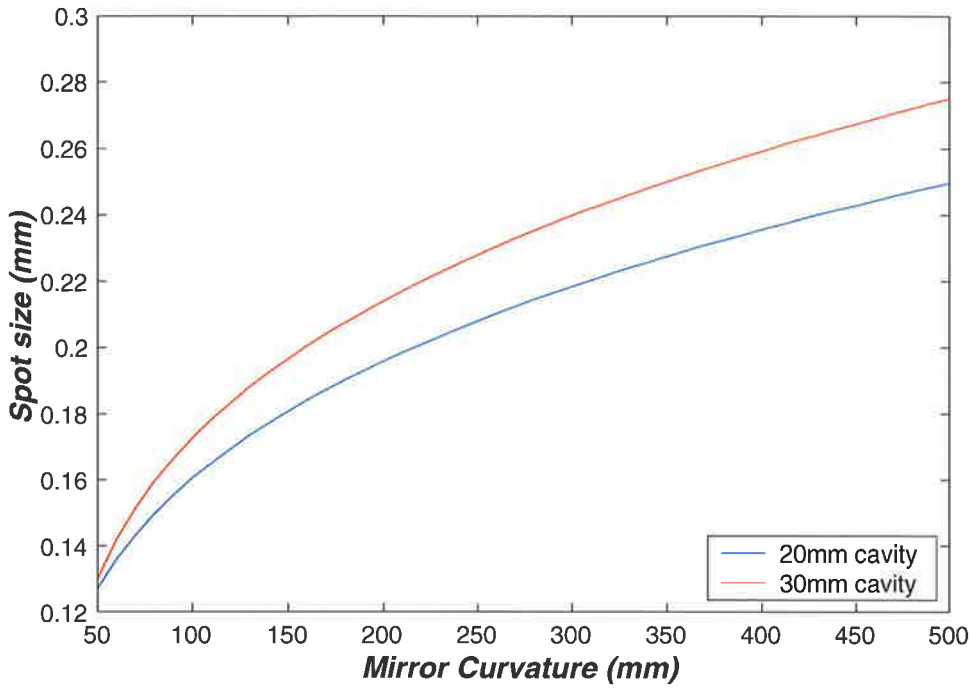

Figure 2.2: Variation of cavity waist size with back mirror radius of curvature.

2.2.2 Desired Sensitivity

Having determined the relevant parameters of the test cavity we can now estimate the expected thermoelastic noise. The thermoelastic noise of the two cavity surfaces can be considered uncorrelated for frequencies above the inverse diffusion time between the two faces. Evaluation of equation 2.5 for our cavity results in a correlation frequency of order 5 mHz.

$$f_d = \frac{1}{2\pi} \frac{\kappa}{\rho C l^2} \tag{2.5}$$

where l is the distance between the cavity faces.

In our frequency range of interest, we can therefore simply add the thermoelastic noise of the two surfaces in quadrature. Our estimate for test mass thermoelastic displacement noise is therefore:

$$S_{TE}(\omega) = C_{FTM}^2 \times \frac{8(1+\sigma)^2\alpha_l^2\kappa k_B T^2}{\sqrt{2\pi}C_v^2\rho^2}\frac{J(\omega)}{\omega_c^2} \times \left[\frac{1}{r_{flat}^3} + \frac{1}{r_{curved}^3}\right]$$

$$\text{or, } S_{TE}(f) = C_{FTM}^2 \times \frac{8(1+\sigma)^2\alpha_l^2\kappa k_B T^2}{(2\pi)^{5/2}C_v^2\rho^2}\frac{J(2\pi f)}{f_c^2} \times \left[\frac{1}{r_{flat}^3} + \frac{1}{r_{curved}^3}\right] \quad (2.6)$$

where $\begin{cases} r_{flat} = 172\mu m \\ r_{curved} = 185\mu m \end{cases}$ as described above.

For our choice of cavity geometry, extrapolation from the C_{FTM} graph in [31] leads to an estimate of C_{FTM} =1.1 and a critical frequency of 60 Hz for the change in the slope of the thermoelastic noise spectrum.

The spectrum of $S_{TE}^{1/2}$ predicted by equation 2.6 is plotted in figure 2.3, defining the sensitivity required for the experiment. Noise curves other than test cavity thermoelastic noise must not enter the pink region on the graph as this will prevent the measurement of thermoelastic noise in our frequency band of interest. We have set a desired performance region a factor of ten less than the expected thermoelastic noise level, as shown in green in the diagram. Noise sources in this region will not significantly degrade sensitivity. The yellow region indicates an adequate performance region where competing noise sources become a concern. Narrow-band noise sources, such as mains harmonics, are likely to violate the sensitivity specification but these are not expected to cause significant problems as they can be easily identified.

The target sensitivity can be converted to equivalent frequency noise by using equation 1.6.

$$S_f^{1/2}(f) = \frac{c}{\lambda L}S_{TE}^{1/2}(f) \quad (2.7)$$

$$= 8.03 \times 10^{15} S_{TE}^{1/2}(f) \quad Hz/\sqrt{Hz} \quad (2.8)$$

The equivalent frequency noise is also shown on figure 2.3.

The design specification may be numerically approximated as

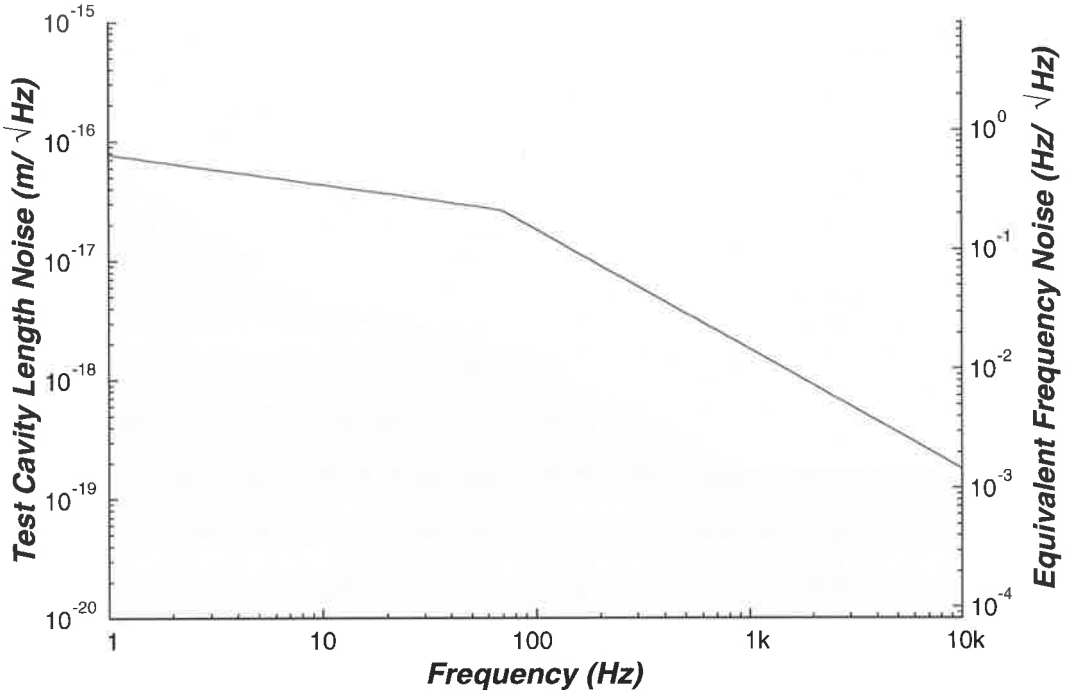

Figure 2.3: Target sensitivity for the experiment. The expected thermoelastic noise level is shown as the blue curve. The spectra of contaminating noise sources must not enter the pink region. The green region indicates the desired sensitivity performance, the yellow region indicates the region where the sensitivity performance is acceptable.

$$S_f^{1/2} \leq \begin{cases} 30 \times 10^{-3} Hz/\sqrt{Hz} \text{ at } 10\ Hz \\ 10 \times 10^{-3} Hz/\sqrt{Hz} \text{ at } 100\ Hz \\ 1 \times 10^{-3} Hz/\sqrt{Hz} \text{ at } 1\ kHz \end{cases}$$

2.3 Competing Noise Sources

2.3.1 Introduction

We will only be able to measure thermoelastic noise signal if it is larger than all other noise sources in the experiment. In the following sections we will therefore examine the noise sources expected to be present in the experiment. By estimating the magnitude and frequency distribution of each noise source we will be able to determine the frequency range over which we should expect to detect thermoelastic noise.

We begin this section with an examination of several noise sources that produce length changes in the test mass: thermal noise, thermorefractive noise and photothermal noise. We also consider the effect of radiation pressure on the test cavity. This analysis will show that thermoelastic noise is the dominant length noise in the test cavity.

As has been described in chapter 1, frequency noise in the laser used to probe the length of the test cavity is indistinguishable from cavity length fluctuations. In section 2.3.3 we will consider the free-running laser frequency noise and show that it is too high for our requirements. Consequentially the laser needs to be stabilised with the use of a reference cavity. Although a detailed description of the optical design of this cavity must wait until chapter 4, in section 2.3.4 we proceed with analysis of its noise contribution. This analysis will show that the use of our reference cavity will reduce the laser frequency noise sufficiently to meet our target sensitivity goal.

The noise sources examined in this chapter are restricted to those arising from physical effects in the test and reference cavities. The analysis is thus intended to show that an ideal system would be able to measure thermoelastic noise in sapphire. Additional noise sources such as detector and servo noise and intensity noise contamination of the readout scheme will be analysed in chapter 4.

2.3.2 Test Cavity Noise Sources

Apart from thermoelastic noise, there are four effects that cause fluctuations in the optical length of the test cavity: thermal noise, photothermal noise and radiation pressure noise, all of which cause fluctuations in the physical length of the cavity, and thermorefractive noise, which leads to an optical length fluctuation by changing the refractive index inside the cavity.

2.3.2.1 Test Cavity Thermal Noise

We know from section 1.2.1 that the spectrum of the displacement fluctuations of the surface of the test mass due to thermal noise with structural damping can be expressed as:

$$S_{x,Th}^{tm}(f) = 4k_B T \frac{1-\sigma^2}{\pi^3 E_o r_o} I\phi \times \frac{1}{f}$$

where
- k_B is Boltzmann's constant
- T is the temperature of the material
- σ is Poisson's ratio for the material
- E_o is Young's modulus of the material
- I is 1.873 22 for a Gaussian probe beam
- ϕ is the loss angle of the material
- r_0 is the spot size of the probe laser beam

Converting to equivalent frequency noise in a cavity we obtain;

$$S^{tm}_{Th,bulk}(f) = \left(\frac{c}{\lambda L}\right)^2 4k_B T \frac{1-\sigma^2}{\pi^3 E_o r_o} I\phi \times \frac{1}{f} \tag{2.9}$$

Since structural damping is the dominant loss mechanism in sapphire [33], the sapphire substrate exhibits a constant value of ϕ with frequency, with a value of $\phi = 3 \times 10^{-9}$ being cited in [12]. However, multi-layer dielectric coating of an optical material significantly degrades the achievable ϕ, because of the high local losses in the coating layers [34]. The analysis of coating thermal noise given in [35] shows that if the coating shares the mechanical properties of the bulk material, then the thermal noise spectrum is increased by a factor:

$$\frac{S^{tm}_{Th,coating}(f)}{S^{tm}_{Th,bulk}(f)} = 1 + \frac{2}{\sqrt{\pi}} \frac{1-2\sigma}{1-\sigma} \frac{\phi_c}{\phi_b} \frac{d}{r_0} \tag{2.10}$$

where d is the thickness of the coating and ϕ_b and ϕ_{coat} are the loss angles in the bulk and coating respectively. If the mechanical properties of the coating differ from those the substrate, then the equation can be extended:

$$\begin{aligned} \frac{S^{tm}_{Th,coating}(f)}{S^{tm}_{Th,bulk}(f)} &= 1 + \frac{1}{\sqrt{\pi}} \frac{1+\sigma_c}{(1-\sigma_b^2)(1-\sigma_c)} \frac{\phi_c}{\phi_b} \frac{d}{r_0} \frac{E_b}{E_c} \times M \\ \text{where } M &= (1-2\sigma_c) + (1-2\sigma_b)^2 \frac{(1+\sigma_b)^2}{(1+\sigma_c)^2} \end{aligned} \tag{2.11}$$

Here σ_b and σ_c are Poisson's ratio for the bulk and coating and E_b and E_c are

Young's modulus for the bulk and coating.

Experiments to measure coating thermal noise produced in SiO_2/Ta_2O_5 multilayer coatings have yielded values of ϕ between 1×10^{-4} [36] and 2.7×10^{-4} [16]. The thickness of the coating in our test cavity is unknown, but it is of the order 10 μm [16, 35]. Substitution into equation 2.11 shows that the coating leads to a considerable increase in thermal noise, as can be seen in figure 2.4. The increase in thermal noise is large because the coating thickness is a significant fraction of the laser spot size, so the d/r_0 term in equations 2.10 and 2.11 is much larger than it would be in a gravitational wave detector.

Substitution of the appropriate constants from Table 2.1 into equation 2.9 yields;

$$S^{tm}_{Th,bulk}(f) \approx 2.2 \times 10^{-6}\frac{1}{f} \quad Hz^2/Hz$$

and

$$S^{tm}_{Th,coating}(f) \approx 2.5 \times 10^{-3}\frac{1}{f} \quad Hz^2/Hz$$

As shown in figure 2.4, the total thermal noise level is below the thermoelastic noise level throughout our measurement frequency range.

2.3.2.2 Test Cavity Thermorefractive Noise

The random temperature variations present in all materials were discussed in section 2.2.1. At that time we were concerned only with coupling of these variations to length changes through a non-zero coefficient of thermal expansion as this is the cause of thermoelastic noise. However, in addition to the physical length changes, these random temperature fluctuations also produce variations in the local refractive index, which can be expressed as $\delta n = \beta \times \delta T$, where $\beta = \frac{\partial n}{\partial T}\big|_\rho$, the thermorefractive constant of the material.

The effect of refractive index variations has been studied extensively by the optical fibre sensor community [37, 38] where it is of particular significance due to the long optical paths inside the fibres. In the context of smaller optics the effect has been studied for reflective optics [31, 39] and some preliminary work has been conducted on transmissive optics [40]. To date no complete model for thermorefractive

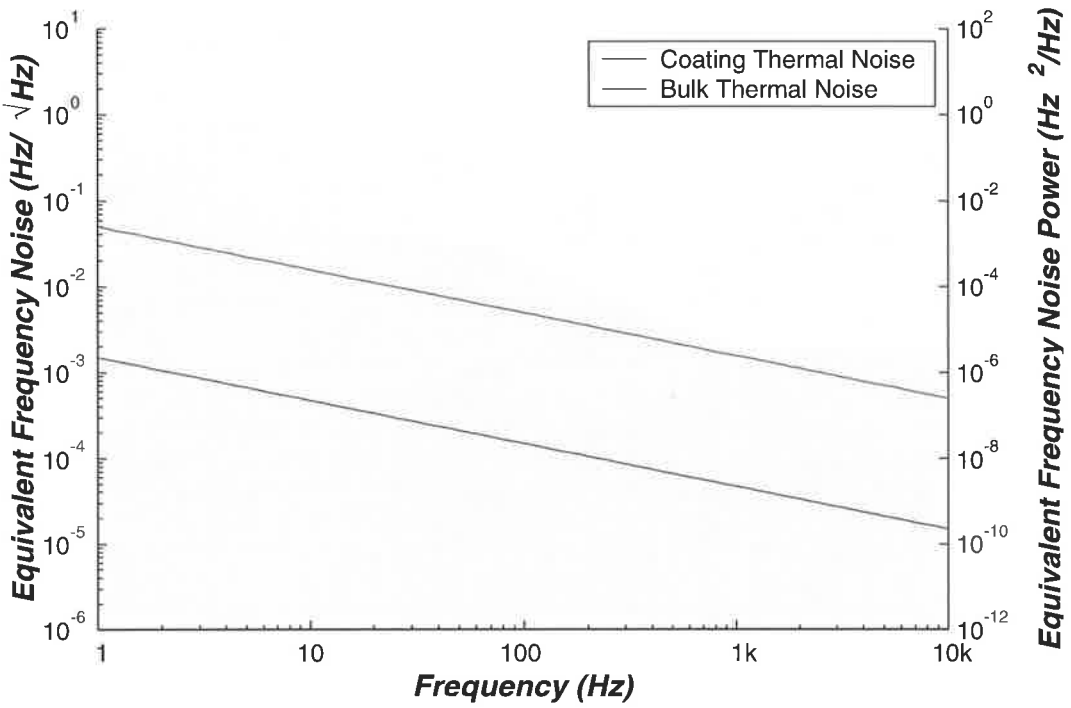

Figure 2.4: Estimated spectrum of test cavity thermal noise. The red curve shows the predicted thermal noise spectum for the coatings. The blue curve shows the noise predicted for the bulk material thermal noise. The background colours indicate the required sensitivity regions.

noise in short, large diameter optics has been published so here we use the fibre treatment described in [41], which gives a formula for the frequency noise produced by thermorefractive effects:

$$S_{TR}^{tm}(f) = \left(\frac{\alpha}{n}+\beta\right)^2\left(\frac{c}{\lambda}\right)^2 T_{rms}^2(f) \tag{2.12}$$

$$\text{where } T_{rms}^2(f) = \frac{k_B T^2}{2\pi\kappa L}\ln\left\{\frac{\left[\left(\frac{2}{\pi r_0}\right)^2+\left(\frac{Nf}{c}\right)^2\right]^2+\left[\frac{f}{2\pi D}\right]^2}{\left[\left(\frac{2.405}{\pi R}\right)^2+\left(\frac{Nf}{c}\right)^2\right]^2+\left[\frac{f}{2\pi D}\right]^2}\right\} \tag{2.13}$$

$$\Rightarrow S_{TR}^{tm}(f) \approx 1\times10^{-8}\ Hz^2/Hz \text{ at } 10\ Hz. \tag{2.14}$$

where R is the radius of the test mass and D is the thermal diffusivity of the material.

The spectrum of the test cavity thermorefractive noise is shown in figure 2.5. The expected noise is negligible for our purposes.

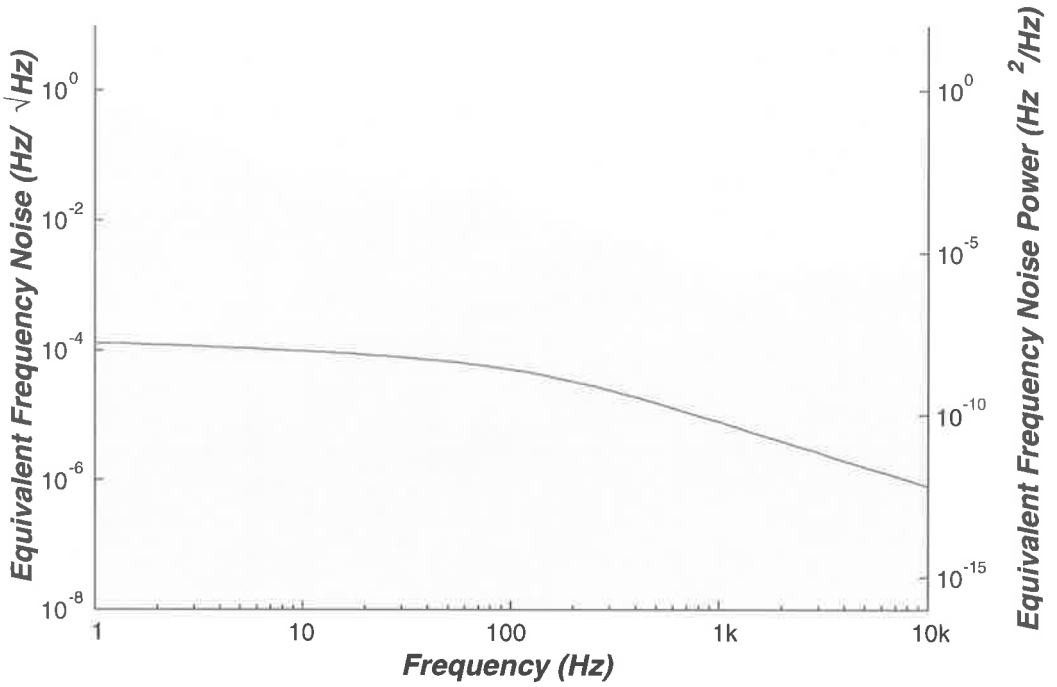

Figure 2.5: Estimated spectrum of test cavity thermorefractive noise. The background colours indicate the required sensitivity regions.

2.3.2.3 Test Cavity Photothermal Noise

Photons circulating inside the test cavity are randomly absorbed by the test mass material, resulting in local heating and a localised expansion of the material. This noise source is therefore sometimes referred to as photothermal *shot* noise to emphasise the statistical nature of the absorption process. Thus, the photothermal process acts as an additional driver for the thermoelastic and thermorefractive noise sources discussed above.

For realistic levels of laser power, the refractive index contribution to photothermal noise is negligible compared to the undriven thermorefractive noise [39] and can safely be disregarded. However, the noise contribution due to expansion is somewhat more significant and so we will examine it in further detail.

In [12] the photothermal noise due to absorption in mirror coatings is treated. The situation in our system is different as the laser remains within the optic and absorption can thus take place in the bulk of the material rather than just at reflections. The analysis assumes that the length scale over which the laser interacts

with the coating is much less than its radius. This assumption is no longer true in our case, so the analysis of [12] should be extended to use a two dimensional version of the thermal conductivity equation (cf. equation (3) of [12]). This is however well beyond the scope of this thesis. However, one would expect that absorption in the surface layers of the test cavity would be more likely to cause significant changes in cavity length due to the proximity to the surface of the resulting volume fluctuations. Absorption further into the bulk of the cavity bulk would need to strain a greater crystal volume to produce the same displacement of the cavity surface and would therefore need to have greater energy to have the same effect. We thus cautiously use the result presented in Braginsky's analysis and realise that it may underestimate the noise somewhat.

Length noise due to photothermal noise is given in [12] as

$$\begin{aligned} S^{tm}_{x,PT}(f) &= \frac{1}{2\pi^2}\alpha_l^2(1+\sigma)^2\frac{\hbar f_c P_c}{(\rho C r_o^2)^2}\times\frac{1}{f^2} \\ &= \frac{1}{2\pi^4}\alpha_l^2(1+\sigma)^2\frac{\hbar\kappa P_c}{(\rho C)^3 r_o^4}\times\frac{1}{f^2} \end{aligned}$$

where
- $\hbar$ is Plank's constant divided by 2π
- f_c is the critical frequency $= 1/2\pi \times \kappa/\rho C$
- P_c is the laser power inside the cavity
- All other symbols are are as defined previously

In our case we must double the predicted noise level to allow for noise in two independent mirrors. Converting to frequency noise and then substituting for the various parameters we find that

$$S^{tm}_{PT}(f) = 2\times\left(\frac{c}{\lambda L}\right)^2\times\frac{1}{2\pi^4}\alpha_l^2(1+\sigma)^2\frac{\hbar\kappa}{(\rho C)^3 r_o^4}\times\frac{P_c}{f^2} \tag{2.15}$$

$$\Rightarrow S^{tm}_{PT}(f) \approx 2.4\times10^{-18}\frac{P_c}{f^2}\quad Hz^2/Hz \tag{2.16}$$

We expect a test cavity circulating power of order 100 W, which gives an estimated photothermal noise level of $S^{tm}_{PT}(f) \approx 2.4\times10^{-16}/f^2\ Hz^2/Hz$. This noise level is insignificant for our purposes and thus it is apparent that more detailed

modelling of this noise source is not justified.

2.3.2.4 Test Cavity Radiation Pressure

Quantum mechanics tells us that every photon carries momentum. An outward pressure on the cavity is therefore produced when cavity photons reflect from the mirror surfaces. If the intensity of the radiation inside the cavity fluctuates then the radiation pressure changes, resulting in modulation of the cavity length in phase with the applied intensity fluctuation.

In a typical thermal noise experiment, the cavity is made with two independently suspended mirrors. Such suspensions inevitably have high mechanical compliance to provide good seismic isolation. Unfortunately the high compliance means that the cavity length is strongly influenced by radiation pressure forces. In a monolithic cavity (such as that used in this experiment) radiation pressure must strain the test mass crystal structure to produce length fluctuations. As the mechanical compliance of the crystal will be much lower than that of suspended systems, our system will be far less susceptible to radiation pressure.

The forces produced by classical and quantum radiation pressure act on the reflective end faces of the test cavity to produce deflection of the equilibrium cavity length. If the frequency of the variation in laser power is below the resonance of the crystal, then we can write the crystal strain as the applied stress divided by Young's modulus. Let us consider the deflection x caused by the reflection of the cavity power from one cavity end mirror.

$$\begin{aligned} strain &= \frac{1}{E_o} \times stress \\ \frac{x}{L} &= \frac{1}{E_o}\frac{F}{A} \qquad (2.17) \\ x &= \frac{L}{E_o A}F \\ &= \frac{L}{E_o \pi r_o^2}F \qquad (2.18) \end{aligned}$$

We can now convert to a power spectrum and double the result to coherently

add the noise from the second mirror:

$$S_{x,rp} = 2\left(\frac{L}{E_o \pi r_o^2}\right)^2 S_{F,rp} \tag{2.19}$$

2.3.2.4.1 Classical Radiation Pressure The momentum carried by a single photon is given by

$$\mathbf{p}_{photon} = \frac{h\nu}{c}.\hat{\mathbf{k}}$$

where
- $\mathbf{p}_{photon}$ is the momentum of the photon
- h is Plank's constant
- ν is the frequency of the photon
- c is the speed of light
- $\hat{\mathbf{k}}$ is a unit vector in the direction of inital photon propagation

We need to determine the force that will be produced when a photon is reflected from a mirror. After the reflection the direction of the photon is reversed. The momentum imparted to the mirror is therefore twice the initial photon momentum. If we have a total of n photons striking the mirror in unit time then the momentum transfer will increase by a factor of n. The force produced on the mirror is given by

$$\begin{aligned} \mathbf{F} &= \frac{\partial \mathbf{p}}{\partial t} \\ &= 2n\frac{h\nu}{c}.\hat{\mathbf{k}} \\ \mathbf{F} &= \frac{2}{c}P_c.\hat{\mathbf{k}} \end{aligned} \tag{2.20}$$

where P_c is the power in the laser beam. In our case this will be the power circulating inside the test cavity. Thus the spectrum has the form

$$S_{F,crp} = \frac{4}{c^2}S_{P_c}$$

It is generally useful to express the spectrum of the input power disturbance in

terms of relative intensity noise. Using

$$S_{P_c} = P_c^2 S_{RIN}$$

where S_{RIN} is the relative intensity noise spectrum of the laser, we can write

$$S_{F,crp} = \frac{4P_c^2}{c^2} S_{RIN}$$

2.3.2.4.2 Quantum Radiation Pressure In addition to any classical fluctuations in laser power there will be shot noise in the laser power. This is produced by quantum fluctuations in laser intensity caused by the quantisation of the cavity light. Even in the absence of excess laser intensity noise there will therefore be a fluctuating radiation pressure.

We know that a photon carries energy $E = h\nu$ and so we can express the number of cavity photons (in unit time) as:

$$n = \frac{P_c}{h\nu} \tag{2.21}$$

where $\begin{cases} P_c \text{ is the circulating power in the cavity} \\ h \text{ is Plank's constant} \\ \nu \text{ is the frequency of the radiation} \end{cases}$

If we approximate the photon statistics inside the cavity as Poissonian then we can write the variation in the number of photons in the cavity as $\sqrt{n}$.

The quantum noise power can thus be expressed:

$$\begin{aligned} P_q &= h\nu\sqrt{n} \\ P_q &= h\nu\sqrt{\frac{P_c}{h\nu}}, \end{aligned} \tag{2.22}$$

which, using equation 2.20, implies a force variation having magnitude

$$F_q = \frac{\sqrt{h\nu P_c}}{c} \tag{2.23}$$

Therefore, we can write

$$S_{F,qrp} = \frac{h\nu P_c}{c^2} \tag{2.24}$$

2.3.2.4.3 Effect of Radiation Pressure To convert the classical and quantum radiation pressures to their equivalent frequency noises we first substitute the appropriate expressions for $S_{F,rp}$ into equation 2.19 and then multiply by $\left(\frac{c}{\lambda L}\right)^2$to convert to frequency noise. For classical radiation pressure we obtain

$$\begin{aligned} S^{tm}_{crp}(f) &= 2\left(\frac{c}{\lambda L}\right)^2\left(\frac{L}{E_o\pi r_o^2}\right)^2\frac{4P_c^2}{c^2}S_{RIN}(f) \\ S^{tm}_{crp}(f) &= \frac{8}{\pi^2\lambda^2 r_o^4 E_o^2}P_c^2 S_{RIN}(f) \end{aligned} \tag{2.25}$$

Similarly, for quantum radiation pressure noise we obtain an equivalent frequency noise of:

$$\begin{aligned} S^{tm}_{qrp}(f) &= 2\left(\frac{c}{\lambda L}\right)^2\left(\frac{L}{E_o\pi r_o^2}\right)^2\frac{h\nu P_c}{c^2} \\ S^{tm}_{qrp}(f) &= \frac{2}{\pi^2\lambda^2 r_o^4 E_o^2}h\nu P_c \end{aligned} \tag{2.26}$$

This analysis assumed that the problem was one dimensional; that is we have assumed that the crystal was strained by radiation pressure having uniform intensity across the entire end face of the test cavity. In reality we know that the laser beam has a Gaussian profile with waist much smaller than the test cavity diameter. The non-uniform shape of the beam means that the radiation force must produce strain in both axial and transverse directions. This inevitably leads to a reduction in axial displacement when compared to the calculated results above.

The magnitude of this effect is not trivial to determine analytically so a finite element model was constructed using the Quickfield analysis package [42] and the

results analysed to yield an estimate of the effects of a finite beam. It should be noted that this analysis applies to both types of radiation pressure noise.

Figure 2.6 shows the response of a 50 mm diameter, 20 mm long sapphire test cavity to radiation pressure having a Gaussian profile and a waist of 200 μm. The non-uniform displacement of the test mass surface can be seen clearly. The displacement is localised around the position of the driving pressure but does not have the same shape; there is a "flattening" of the displacement when compared to the driving gaussian profile. The displacements indicated in the figure are of arbitrary units.

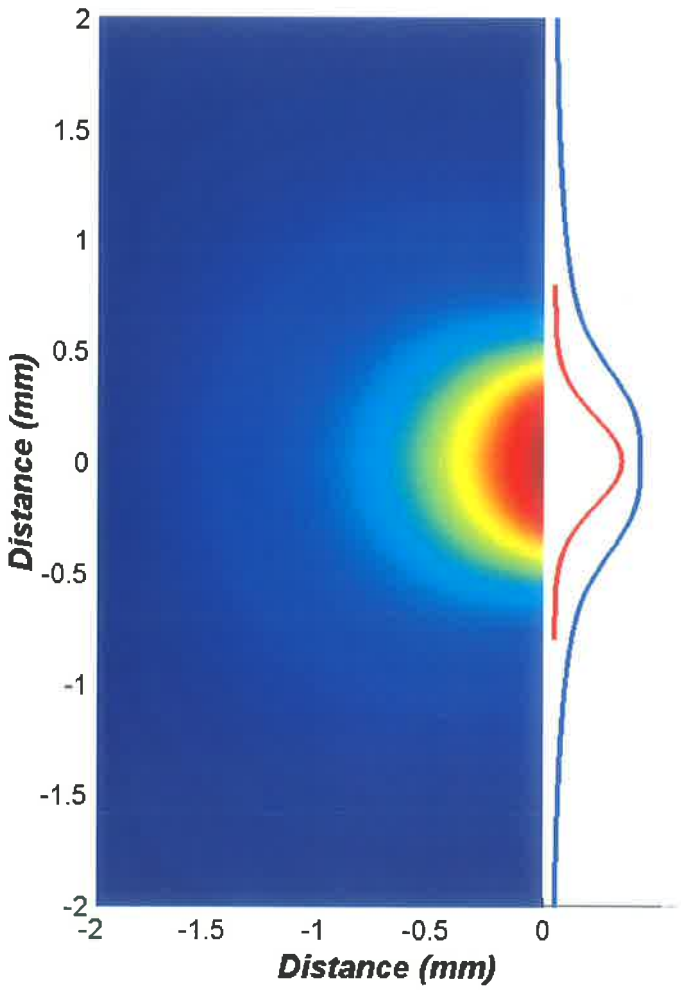

Figure 2.6: Displacement map of a sapphire cylinder with a 200 μm waist gaussian beam applied on axis. The forcing gaussian is shown on the right in red. The blue curve indicates the surface displacement.

To calculate the effective surface displacement seen by the laser beam we calculated the overlap integral of the deformed shape and the probe laser. This integral was then compared with the corresponding quantity obtained for a model having a uniform radiation pressure. This process yielded a factor of 0.48 reduction in surface deflection when compared to the infinite beam case. This factor would be different for different sized probe beams.

We modify equations 2.25 and 2.26 to include the coefficient, C_{FB}, which models the effect of the finite beam. This leads to new equations for the two radiation pressure effects;

$$S^{tm}_{crp}(f) = C^2_{FB}\frac{8}{\pi^2\lambda^2 r_o^4 E_o^2}P_c^2 S_{RIN}(f) \quad (2.27)$$

$$S^{tm}_{qrp}(f) = C^2_{FB}\frac{2}{\pi^2\lambda^2 r_o^4 E_o^2}h\nu P_c \quad (2.28)$$

where C^2_{FB} =0.48 and describes the effect of the finite beam.

Figure 2.7 shows the variation of the noise due to classical radiation pressure for a variety of intracavity powers. As knowledge of the power noise in the laser was required to evaluate equation 2.27, we anticipate results which will be presented in chapter 4. We approximate the relative intensity noise of our laser as $1\times10^{-4}/f$ $/\sqrt{Hz}$. The quantum radiation pressure noise spectrum is presented in figure 2.8.

The spectra for these two noise sources are found to be

$$S^{tm}_{crp}(f) \approx 1.4\times10^{-11}\frac{P_c^2}{f} \quad Hz^2/Hz$$

$$S^{tm}_{qrp}(f) \approx 6.3\times10^{-17}P_c \quad Hz^2/Hz$$

The graphs clearly show that the experiment is insensitive to radiation pressure noise, unless very high cavity power is used. Our system will employ an intracavity power of order 100 watts, so we would not expect radiation pressure to be a problem. This immunity to radiation pressure noise is one of the great advantages of using a monolithic test cavity.

Should the unstabilised classical radiation pressure effects be judged too large

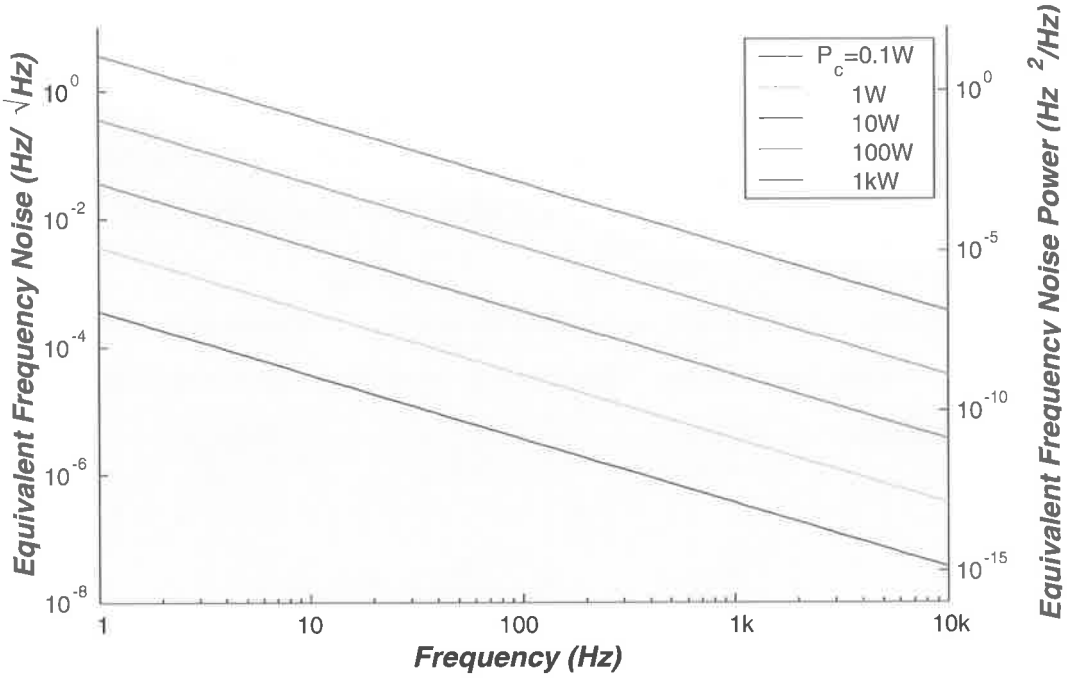

Figure 2.7: Classical radiation pressure noise in the test cavity. The noise spectra are shown for a number of different intracavity powers. The background colours indicate the required sensitivity regions.

then intensity stabilisation of the laser would enable this noise source to be considerably reduced. Given an ideal stabilisation scheme, the classical radiation pressure could be reduced to a level corresponding to the quantum radiation pressure level.

2.3.2.5 Summary of Test Cavity Noise Sources

Figure 2.9 shows the expected spectra for all of the noise sources in the test cavity assuming an intracavity power of 100 watts. As can be seen, thermoelastic noise is the dominant noise source in the test cavity over the entire frequency range, though coating thermal noise is a significant contributor at the high end of the frequency range.

In addition it should be noted that radiation pressure curves shown in the graph assume a laser without any intensity stabilisation. Because the laser is to be stabilised in the experiment, we will find that the real noise contribution of test cavity radiation pressure is significantly lower (see section 4.4.4).

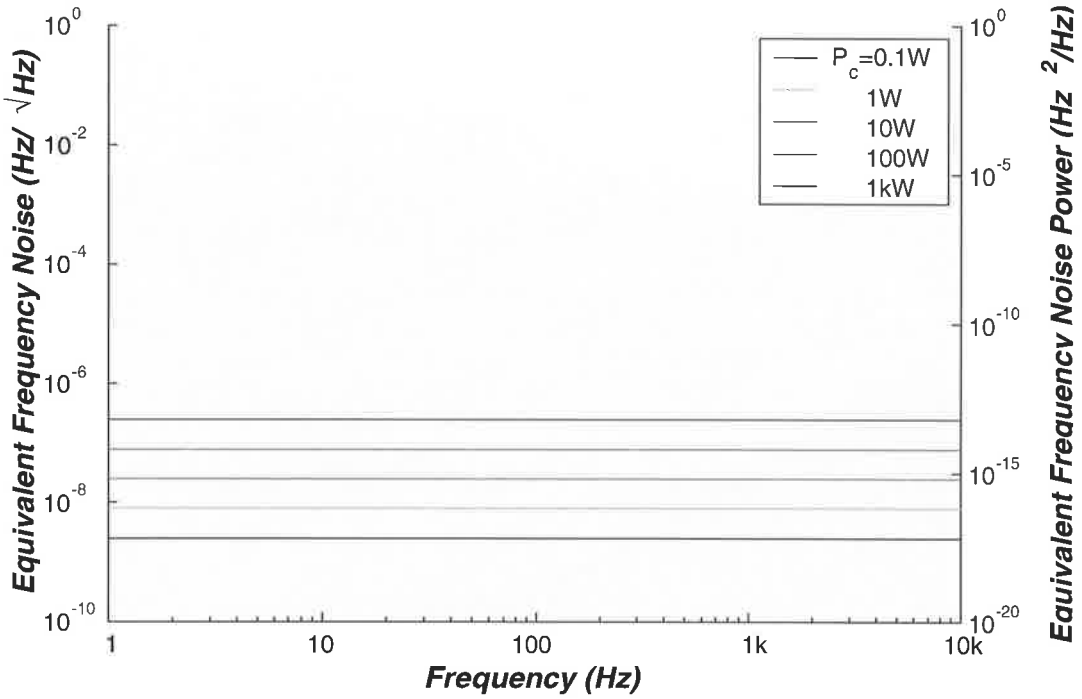

Figure 2.8: Quantum radiation pressure noise in the test cavity. The noise spectra are shown for a number of different intracavity powers. The background colours indicate the required sensitivity regions.

2.3.3 Laser Frequency Noise

The laser used in this project is a Nd:YAG nonplanar ring oscillator (NPRO), more details of which can be found in chapter 4. NPRO lasers are used in a wide range of high precision metrology experiments because their low frequency and intensity noise enables good sensitivity to be achieved. Figure 2.10 shows the typical frequency noise of an NPRO laser, $S_{npro} = 1 \times 10^4/f \quad /Hz$.

This frequency noise is much higher than that produced by the expected thermoelastic noise in our experiment and will thus need to be reduced if it is not to completely mask the test cavity thermoelastic noise. The laser frequency performance will be improved by locking the laser to a stable reference cavity.

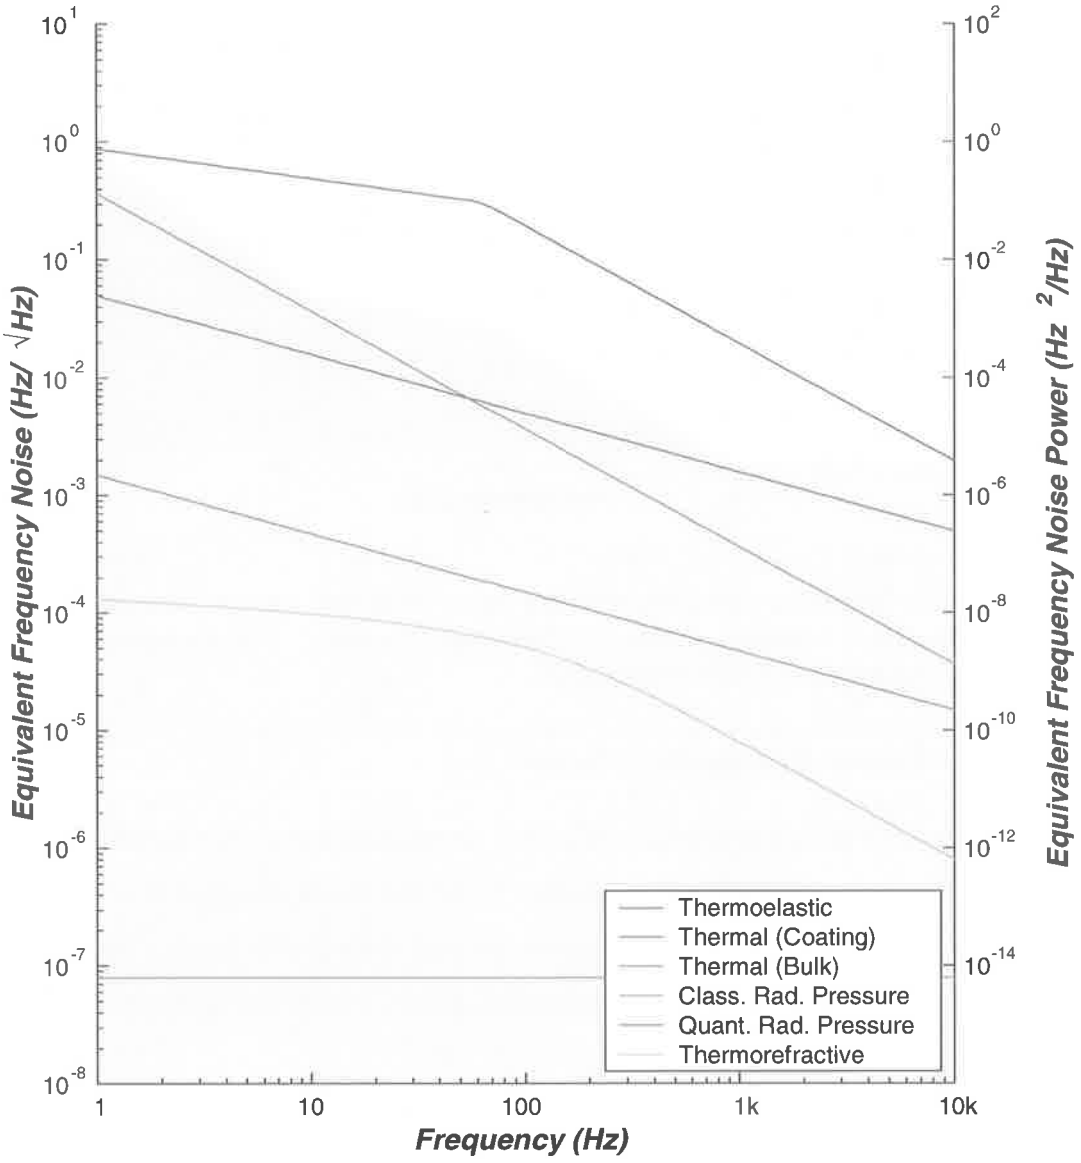

Figure 2.9: Composite plot showing the noise levels estimated for all noise sources present in the test cavity. The background colours indicate the required sensitivity regions.

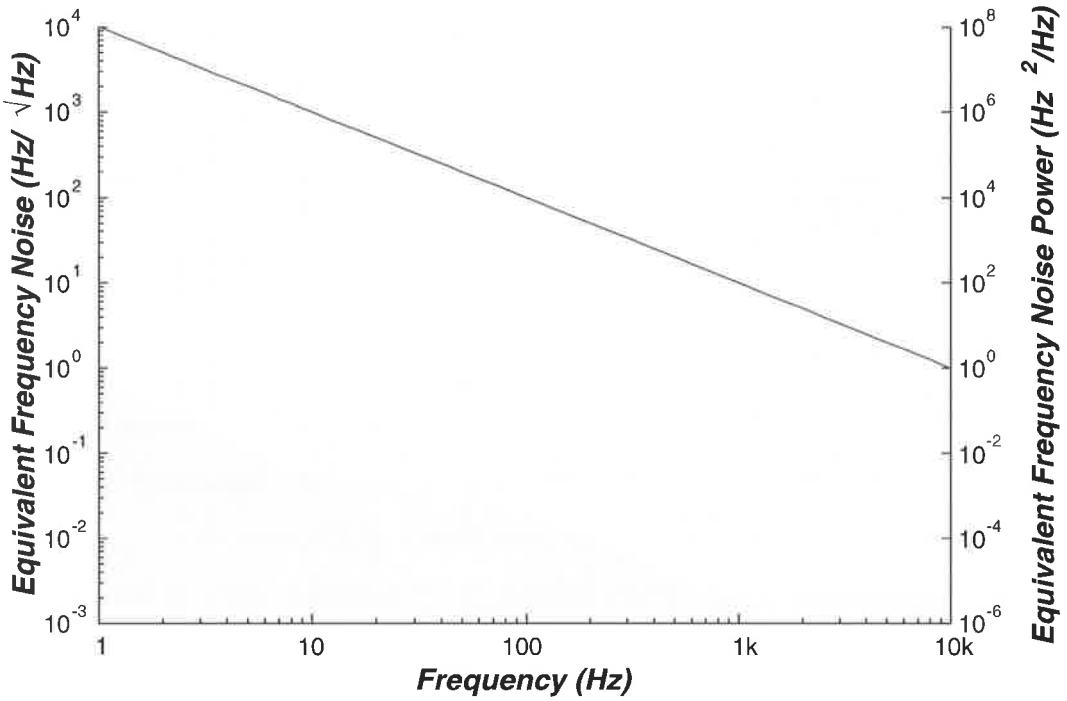

Figure 2.10: Frequency noise of a free running NPRO laser. The background colours indicate the required sensitivity regions.

2.3.4 Reference Cavity Noise Sources

In the previous section we established that the free running NPRO laser would have insufficient frequency stability to measure the thermoelastic noise signal in the test cavity. In chapter 4 we will describe in detail the techniques used to lock the laser to a reference cavity and the noise imposed by the servo control. For now we will simply assume that the laser can be configured so that the frequency of its output is determined by the resonant frequency of an external reference cavity, rather than by its free running value as presented above.

Ultimately, length noise in the reference cavity will limit the achievable frequency stability of a laser locked to the cavity. For this reason we must assure ourselves that the reference cavity length noise satisfies the experimental sensitivity specified above. As has previously been discussed, the reference cavity is made long relative to the test cavity so that the residual frequency noise of the laser beam is low compared to the test cavity thermoelastic noise, since for a given length noise a longer cavity

will experience a smaller frequency change.

Maintenance of the acoustic Q of the reference cavity is important to ensure that mechanical resonances of the cavity are contained in tight frequency bands. This ensures that reference cavity noise does not compromise the system sensitivity excessively. Analysis of the modal distribution in the reference cavity is difficult in this case because of its complex shape so details of the mechanical resonances will not be included in this document.

The reference cavity used in this project is described fully in chapter 4. In summary: the cavity is constructed of three fused silica supermirrors [43] bonded to a monolithic aluminium spacer in a ring configuration. The cavity has a finesse of approximately 1000 and an optical path length of 528 mm.

The catalogue of noise sources present in the reference cavity is much the same as that for the test cavity. We will begin with an examination of thermal noise and then discuss thermoelastic noise in the reference cavity. We will then turn to photothermal, thermorefractive and radiation pressure noise in turn. Finally, we will discuss the optical length noise due to refractive index noise in the gas inside the reference cavity, a noise source that was not present in the test cavity. We close this section with a brief discussion of the ways in which mechanical vibration could couple into the reference cavity.

2.3.4.1 Reference Cavity Thermal Noise

Thermal noise in the reference cavity occurs in both the mirrors and in the aluminium cavity spacer. Thermal noise in the mirrors results in movement of the mirror surfaces relative to their equilibrium position on the spacer. Spacer thermal noise results in changes in the distance between the mirror centres-of-mass.

Because of the reference cavity's complicated shape, a strictly correct formulation for thermal noise would require extensive finite element modelling to determine the mechanical modes. In a manner analogous to [44], the various modes would be summed over to determine the total length noise. As such an approach is well beyond the scope of this thesis we will make several simplifying assumptions in the course of this analysis.

We restate the thermal noise equation that was presented in section 2.3.2.1 above. As can be seen the level of thermal noise is inversely proportional to the radius over which the noise is integrated, r_0.

$$S_{x,Th}(f) = 4k_B T \frac{1-\sigma^2}{\pi^3 E_0 r_0} I\phi \times \frac{1}{f}$$

When calculating thermal noise in the mirrors we use the laser spot size, just as we did in the test cavity case. However, thermal noise in the spacer influences the cavity length by moving the mirrors, so we must consider thermal noise from the entire spacer and choose a value of r_0 that is characteristic of the spacer dimensions. The spacer radius is much larger than the spot size on the cavity mirrors, so despite the relatively high losses in the aluminium spacer, we expect that the spacer's contribution to the thermal noise will be negligible compared with the mirror component.

In the analysis that follows, we assume that the attachment of the mirrors to the spacer does not excessively degrade the thermal noise performance of the system. In reality the glue between the two bodies will be lossy, so the thermal noise will be somewhat worse than that predicted.

In chapter 4 we will see that the spot sizes on all three reference cavity mirrors are approximately 400 μm. We use the above equation to estimate the mirror thermal noise and then multiply by three to sum the uncorrelated length noises from the three cavity mirrors.

Table 2.2 shows the parameters for the mirror's fused silica substrate as obtained from [12] and [39]. The table also includes various additional material constants required for other noise sources to follow. The value for ϕ given in [12] (5×10^{-8}) is for fused silica bulk material. It has been shown that the presence of the mirror coatings significantly reduces the Q of the material [36], as expected from the analysis of [7]. Measurements of ϕ for coatings made from alternating layers of SiO_2 and Ta_2O_5 have yielded ϕ=1×10^{-4} [36] and 2.7×10^{-4} [?]. We use the more pessimistic value of 2.7×10^{-4} here.

Figure 2.11 shows the thermal noise level expected in the reference cavity. The

Symbol	Property	Value	Units
σ	Poisson's ratio	0.17	
E_0	Young's modulus	72×10^9	Pa
ϕ	Loss angle	2.7×10^{-4}	
κ	Thermal conductivity	14	$Wm^{-1}K^{-1}$
α_l	Thermal expansivity	5.5×10^{-7}	K^{-1}
C_v	Specific Heat	670	$Jkg^{-1}K^{-1}$
ρ	Density	2.2×10^3	kgm^{-3}
β	dn/dT	1.45×10^{-5}	K^{-1}

Table 2.2: Table Fused Silica material properties

expected noise level is approximately

$$S^{ref}_{Th,bulk} \approx 8.9\times10^{-9}\frac{1}{f} \quad Hz^2/Hz$$
$$and$$
$$S^{ref}_{Th,coating} \approx 2.0\times10^{-6}\frac{1}{f} \quad Hz^2/Hz$$

As can be seen these thermal noise levels lie considerably below the sensitivity requirement, even allowing for a modest increase in the noise level due to the bonding between the mirrors and spacer or because of a higher value of ϕ than assumed here.

2.3.4.2 Reference Cavity Thermoelastic Noise

Thermoelastic noise also depends critically on the radius over which the noise is averaged. In the case of thermoelastic noise, the dependance is even stronger, being proportional to $1/r_0^2$, so again we ignore the contribution of the spacer.

Figure 2.12 shows the thermoelastic noise level expected in the reference cavity. The reference cavity thermoelastic noise should not mask the test cavity thermoelastic noise. Above the critical frequency the expected thermoelastic noise level is;

$$S^{ref}_{te} \approx 1.1\times10^{-2}\frac{1}{f^2} \quad Hz^2/Hz$$

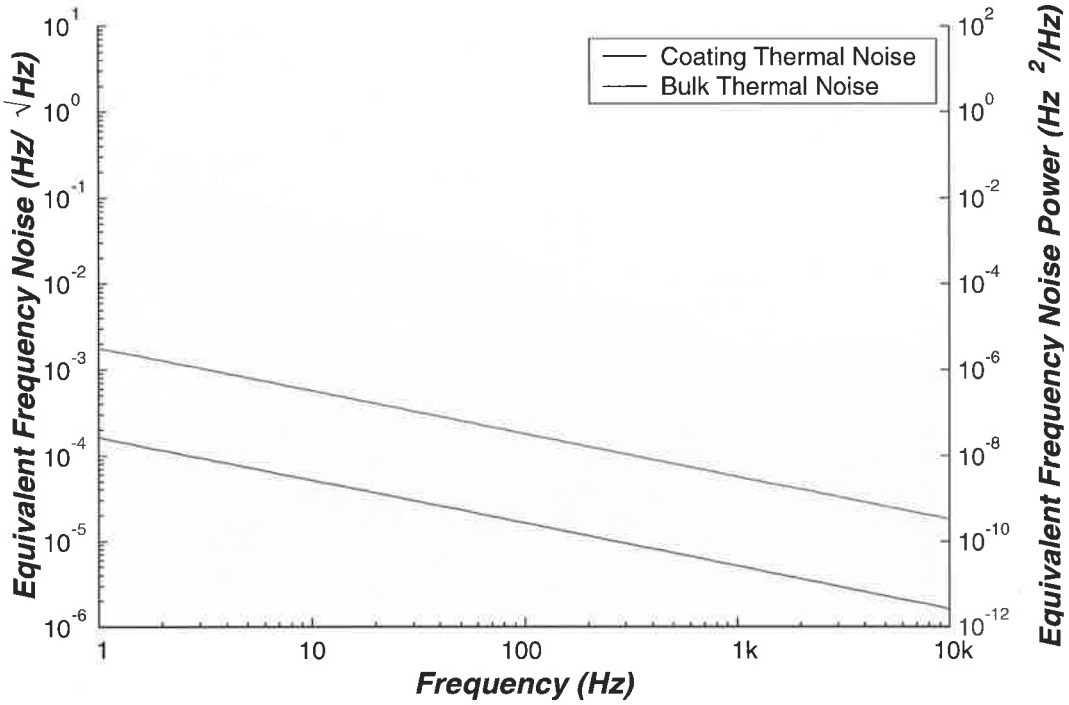

Figure 2.11: Thermal noise in the reference cavity mirrors. The red curve shows that thermal noise produced in the mirrors coatings, the blue that in the bulk of the fused silica. The background colours indicate the required sensitivity regions.

2.3.4.3 Reference Cavity Photothermal Noise

Photothermal noise in the reference cavity mirrors can be calculated using the formula presented in section 2.3.2.3. We see that reference cavity photothermal noise will be insignificant at the intracavity powers employed in this experiment.

$$S_{pt}^{ref} \approx 1.3 \times 10^{-22} \frac{P_c}{f^2} \quad Hz^2/Hz$$

2.3.4.4 Reference Cavity Thermorefractive Noise

Light spends relatively little time inside the mirrors of the reference cavity. Nevertheless, it does pass through the cavity's input and output couplers where it is subject to thermorefractive noise. Using the formula from section 2.3.4.4, we see that the noise spectrum can be expressed as

$$S_{tr}^{ref} \approx 2 \times 10^{-16} \frac{1}{\sqrt{f}} \quad Hz^2/Hz$$

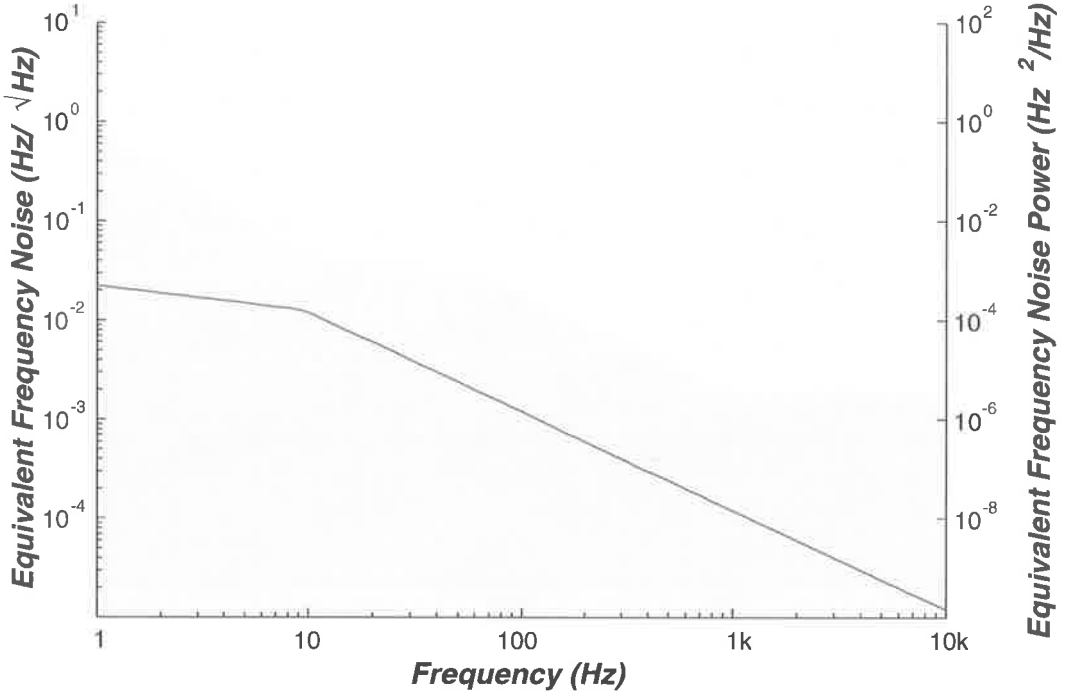

Figure 2.12: Thermoelastic noise level in the reference cavity mirrors. The background colours indicate the required sensitivity regions.

This noise level easily meets the sensitivity goal of the experiment.

2.3.4.5 Reference Cavity Radiation Pressure

Radiation pressure fluctuations produce a displacement of each mirror surface in the reference cavity. A finite element model of the mirror and spacer was constructed, the mirror being assumed to be perfectly bonded to the aluminium spacer in a ring 5 *mm* from its centre (as was the case in the reference cavity). A force having a Gaussian profile was applied to the centre of the mirror and the resulting static deflection calculated. Analysis of the resulting system showed that the displacement caused by radiation pressure is predominantly in the mirrors rather than the spacer material. The effects of the glue used to bond the mirrors to the spacer was neglected in the model. We assume that its effects are minor because the surface area over which it is applied is large compared to the laser spot size. Figure 2.13 shows the resulting displacement profile. Using the same procedure discussed in section 2.3.2.4, this model was used to determine the C_{FB}^2 coefficient required to adjust the

radiation pressure for the Gaussian force profile. A beam with a waist of 400 μm was found to have a C_{FB}^2 of 0.30 in fused silica.

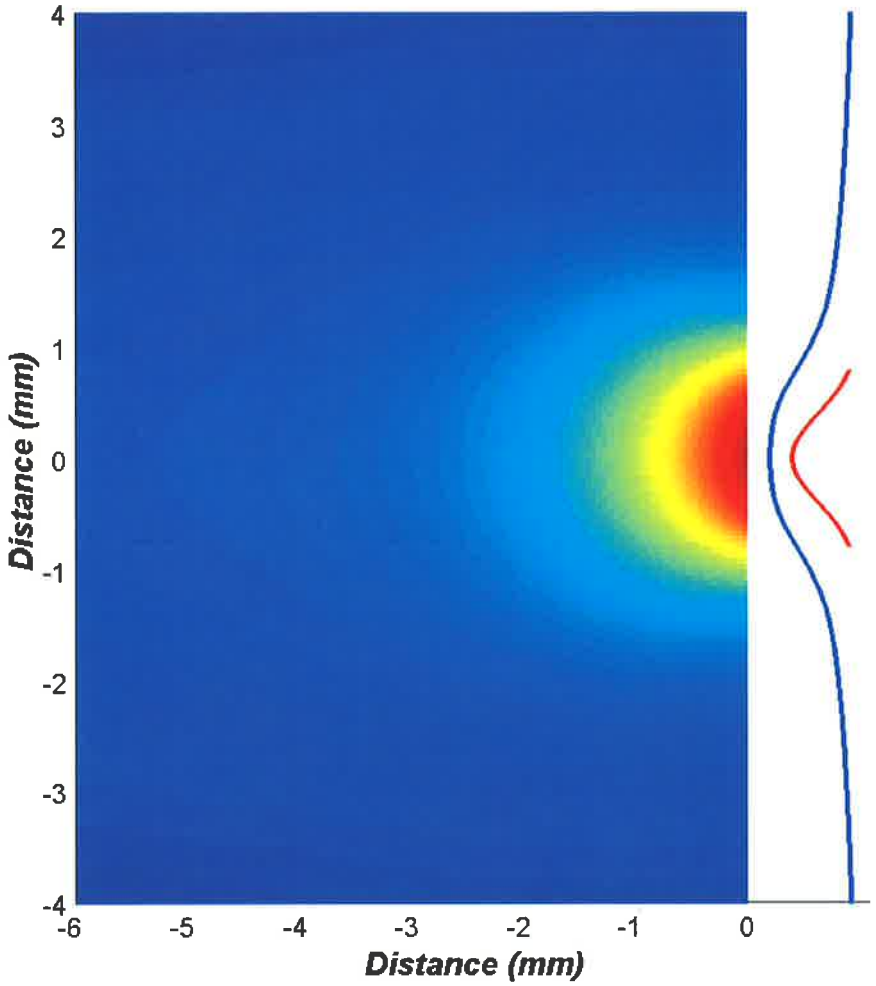

Figure 2.13: Displacement map of a fused silica cylinder with a 400 μm waist gaussian beam applied on axis. The forcing gaussian is shown on the right in red. The blue curve indicates the surface diaplacement.

Knowing C_{FB}^2 allows us to determine the displacement produced by three correlated radiation pressure effects at the three mirrors:

$$S_{x,rp}^{ref} = 3 \times C_{FB}^2 \left(\frac{X_{mirror}}{E_o \pi r_o^2} \right)^2 S_{F,rp} \tag{2.29}$$

where X_{mirror} is the thickness of each mirror substrate.

We now make the conversion from length to frequency noise;

$$\begin{aligned} S_{rp}^{ref} &= \left(\frac{c}{\lambda L}\right)^2 S_{x,rp} \\ &= 3C_{FB}^2 \left(\frac{c}{\lambda L}\right)^2 \left(\frac{X_{mirror}}{E_o \pi r_o^2}\right)^2 S_{F,rp} \end{aligned} \tag{2.30}$$

We can write equations for classical and quantum radiation pressure as follows:

$$\begin{aligned} S_{crp}^{ref}(f) &= 3C_{FB}^2 \left(\frac{c}{\lambda L}\right)^2 \left(\frac{X_{mirror}}{E_o \pi r_o^2}\right)^2 \frac{4P_c^2}{c^2} S_{RIN}(f) \\ &= 12C_{FB}^2 \frac{1}{(\lambda L)^2} \left(\frac{X_{mirror}}{E_o \pi r_o^2}\right)^2 P_c^2 S_{RIN}(f) \end{aligned} \tag{2.31}$$

and

$$\begin{aligned} S_{qrp}^{ref}(f) &= 3C_{FB}^2 \left(\frac{c}{\lambda L}\right)^2 \left(\frac{X_{mirror}}{E_o \pi r_o^2}\right)^2 \frac{h\nu P_c}{c^2} \\ &= 3C_{FB}^2 \frac{1}{(\lambda L)^2} \left(\frac{X_{mirror}}{E_o \pi r_o^2}\right)^2 h\nu P_c \end{aligned} \tag{2.32}$$

Substitution of the appropriate constants into equations 2.31 and 2.32 yields expected noise spectra of

$$\begin{aligned} S_{crp}^{ref} &\approx 1.3 \times 10^{-8} \frac{P_c^2}{f} \quad Hz^2/Hz \\ S_{qrp}^{ref} &\approx 6.0 \times 10^{-20} P_c \quad Hz^2/Hz \end{aligned}$$

These spectra are shown in figures 2.14 and 2.15 for a variety of cavity circulating powers. We see that radiation pressure noise will not be a problem at realistic power levels. As was discussed for the test cavity, classical radiation pressure can be further reduced by intensity stabilisation of the laser.

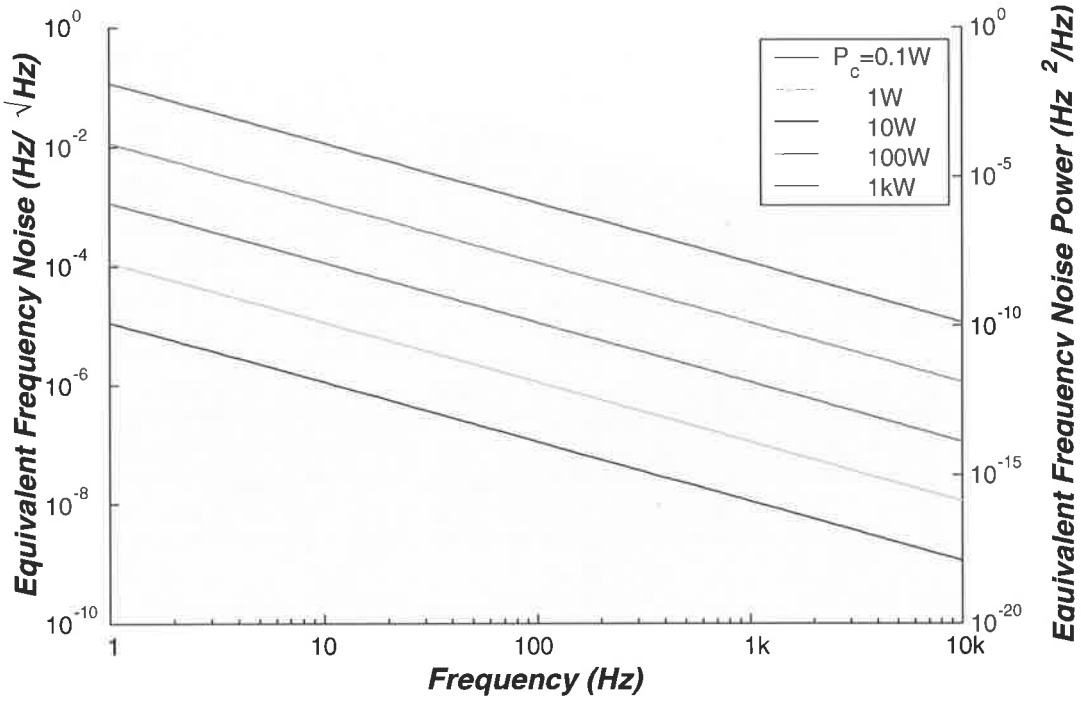

Figure 2.14: Classical radiation pressure noise in the reference cavity. The noise spectra are shown for a number of differnt intracavity powers. The background colours indicate the required sensitivity regions.

2.3.4.6 Reference Cavity Gas Refractive Index Variation

The presence of gas inside the test cavity spacer results in refractive index fluctuations. A change in refractive index produces a change in the optical path length inside the cavity and hence a shift in the resonant frequency. The refractive index changes will therefore produce frequency noise in the laser beam exiting the reference cavity [45].

It has been shown [46] that the equivalent strain spectrum due to pressure fluctuations in a cavity having optical length L can be approximated as

$$S_{h,gas}(f) \approx 2\sqrt{2}\frac{(n_0-1)^2}{\bar{v}N_0 L r_0}\left(\frac{P}{P_0}\right) \tag{2.33}$$

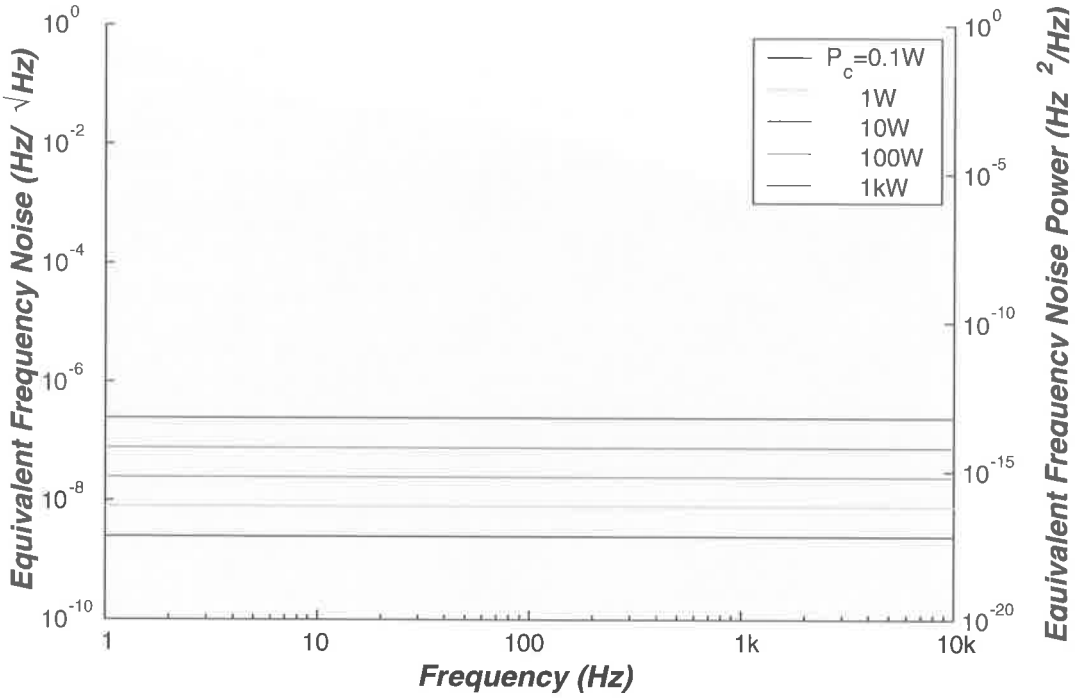

Figure 2.15: Quantum radiation pressure noise in the reference cavity. The noise spectra are shown for a number of different inracavity powers. The background colours indicate the required sensitivity regions.

where
- n_0 is the refractive index of air at 1atm
- N_0 is number of molecules/unit volume at 1 atm
- P is the pressure in the cavity
- P_0 is atmospheric pressure
- $\bar{v}_0$ is the mean molecular speed at 1atm
- r_0 is the beam radius
- L is the optical path length inside the cavity

We can convert equation 2.33 into length noise, $S_{x,gas}$, by multiplying by the cavity length. We then use the standard formula to convert to frequency noise, S_{gas}.

$$\begin{aligned} S_{x,gas} &= L^2 S_{h,gas} \\ &= 2\sqrt{2}\frac{L(n_0-1)^2}{\bar{v}N_o r_o}\left(\frac{P}{P_0}\right) \\ \Rightarrow S_{gas} &= \left(\frac{c}{\lambda}\right)^2 2\sqrt{2}\frac{(n_0-1)^2}{L\bar{v}N_0 r_0}\left(\frac{P}{P_0}\right) \end{aligned}$$

The analysis which leads to the above equation is applicable only in the regime where a gas molecule may be considered to pass through the laser beam in less than the observing time. The interaction between a given gas molecule and the cavity mode lasts for a mean time $t_i = \sqrt{2}r_o/v_m$,where v_m is the "most likely" molecular velocity and is equal to $\sqrt{2kT/m}$. The above analysis is therefore accurate only for frequencies less than a critical frequency f_c where,

$$\begin{aligned} f_c &= \frac{1}{2\pi t_i} \\ &= \frac{v_m}{2\sqrt{2}r_0} \end{aligned}$$

With our choice of r_o this critical frequency is of the order of a megahertz, so we may feel confident in using the above description.

Appropriate gas constants were obtained from [47] and [48] and are shown in table 2.3. Gas constants for nitrogen were used where a choice of gas was required giving

$$S_{gas}^{ref} \approx 1.4 \times 10^{-2}\frac{P}{P_o} \quad Hz^2/Hz$$

Symbol		Value
n_0		1.000 3
N_0		$2.7\times10^{25}\ m^{-3}$
P_0		$760\ Torr$
$\bar{v}_0$	$=\sqrt{\frac{3k_BT}{m}}$	$510\ m/s$

Table 2.3: Table Gas properties used to determine gas noise

Figure 2.16 shows the noise spectrum resulting from gas density fluctuations. At atmospheric pressure the gas induced noise would reduce the frequency range over

which thermoelastic noise can be seen, as it is higher than our design sensitivity for frequencies above 200 Hz. This limitation is acceptable for the phase of the project described in this thesis. The next phase of the experiment will involve a move into vacuum, which will reduce the gas noise to well below the sensitivity target.

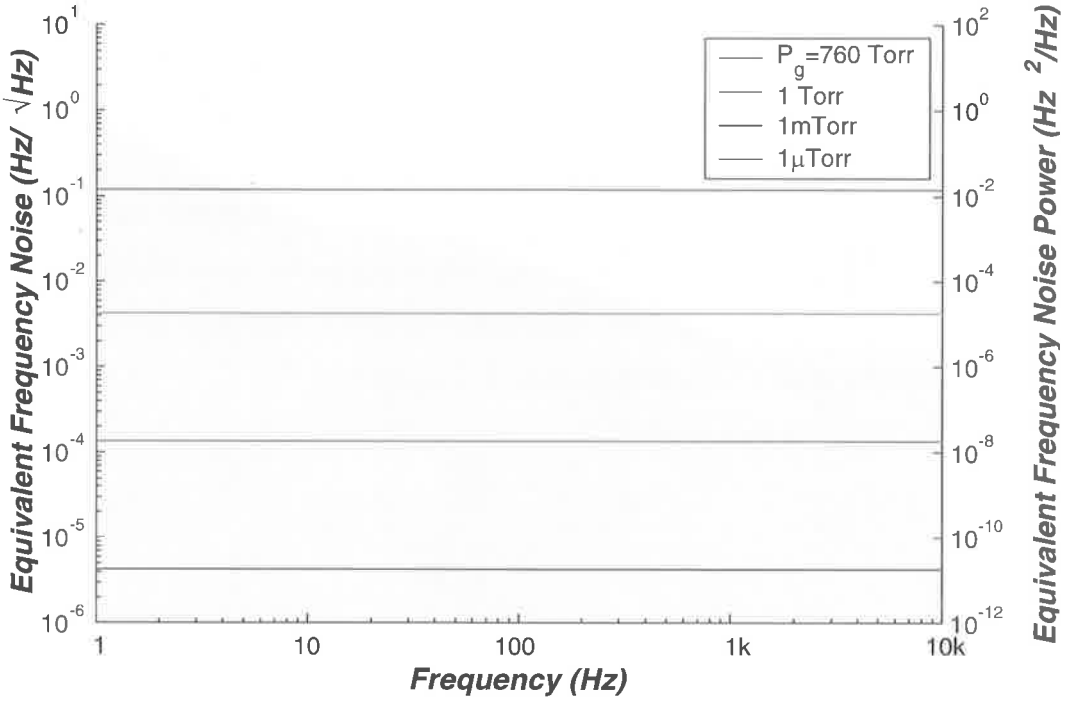

Figure 2.16: Frequency noise caused by gas density fluctuations inside the reference cavity. The noise spectrum is shown for a number of different gas pressures. The background colours indicate the required sensitivity regions.

2.3.4.7 Acoustic and Mechanical Noise

Acoustic noise can cause length noise in the reference cavity, either through direct excitation or by exciting surrounding structures which transmit noise to the cavity through its mount. The experiments described in this thesis were conducted during quiet times of the night to limit this effect. In future phases of the experiment the reference cavity will be in vacuum, so acoustic excitation will be negligible.

Mechanical noise is also coupled into the reference cavity from ground motion, either due to seismic excitation or other local mechanical excitation sources. The effect of this excitation is again rather hard to estimate, but its severity can be

mitigated by using an appropriate seismic isolation system. A description of the seismic isolation system used in this experiment can be found in chapter 3. While the work described in this thesis used the preisolator to minimise ground vibration, when the reference cavity is moved into vacuum it will be supported on a vibration isolation suspension to further reduce excitation. This isolation system will also prevent degradation of the reference cavity Q through mechanical coupling to the optical table.

2.3.4.8 Summary of Reference Cavity Noise Sources

Figure 2.17 shows a consolidated graph of all estimates of length noise in the reference cavity. As can be seen, refractive index noise caused by the gas inside the cavity will limit the experiment at atmospheric pressure. This means that we cannot expect to see thermoelastic noise in the measurements described in this thesis. Once the experiment is moved to a vacuum environment, we can see from the graph that the frequency stability provided by the reference cavity will be sufficient to allow detection of test cavity thermoelastic noise.

2.4 Summary

This chapter began with an analysis of the expected thermoelastic noise signal. We confirmed the result of Numata [14], the derivation of which was not reported. Thus, we defined a target sensitivity for the experiment, based on the goal of measuring the thermoelastic noise spectrum of sapphire between 10 Hz and 1 kHz. Though detailed design of the test cavity is described in chapter 4, in this chapter we described some of its mechanical properties. This test cavity is a short monolithic test cavity, that in contrast to a test cavity consisting of two individually suspended mirrors is largely insensitive to intensity noise. Analysis of all expected intrinsic length noise sources in this test cavity showed that thermoelastic noise will be the dominant source of length noise in the test cavity over the frequency range from 1 Hz up to beyond 10 kHz.

We also considered the effect of laser frequency and found that, as expected,

the free-running NPRO laser source will not be adequate for the measurement of thermoelastic noise in the sapphire test cavity and it will need to be stabilised using a long reference cavity.

Length noise in the reference cavity imposes a fundamental limit to the absolute frequency stability of the laser and thus we analysed this noise in detail in section 2.3.4. The calculations showed that thermoelastic noise in the reference cavity mirrors will be the limiting factor in reference cavity length noise for all frequencies of interest. However, the level of the thermoelastic noise is sufficiently small that it will not mask test cavity thermoelastic noise.

Our model of reference cavity thermal noise neglected the effects of the glue used to bond the mirrors to the spacer. It is possible that the glue could increase the reference cavity thermal noise level beyond the thermoelastic noise level. If this is found to be the case in the next phase of the experiment, then an alternate reference cavity made with a ULE spacer with optically bonded mirrors will be employed.

The reference cavity is at atmospheric pressure for the first phase of the experiment. Refractive index variations in the gas inside the reference cavity will limit the achievable length noise performance of the reference cavity and prevent the measurement of test cavity thermoelastic noise. Placing the reference cavity in vacuum in the next stage of the experiment will, however, make this source of noise negligible, as well as significantly reducing acoustic excitation of the reference cavity. The lack of gas damping will also improve thermal noise in the reference cavity spacer as will mounting the reference cavity on a suspension system. In addition, seismic and other ground borne excitation of the reference cavity will be reduced by the suspension.

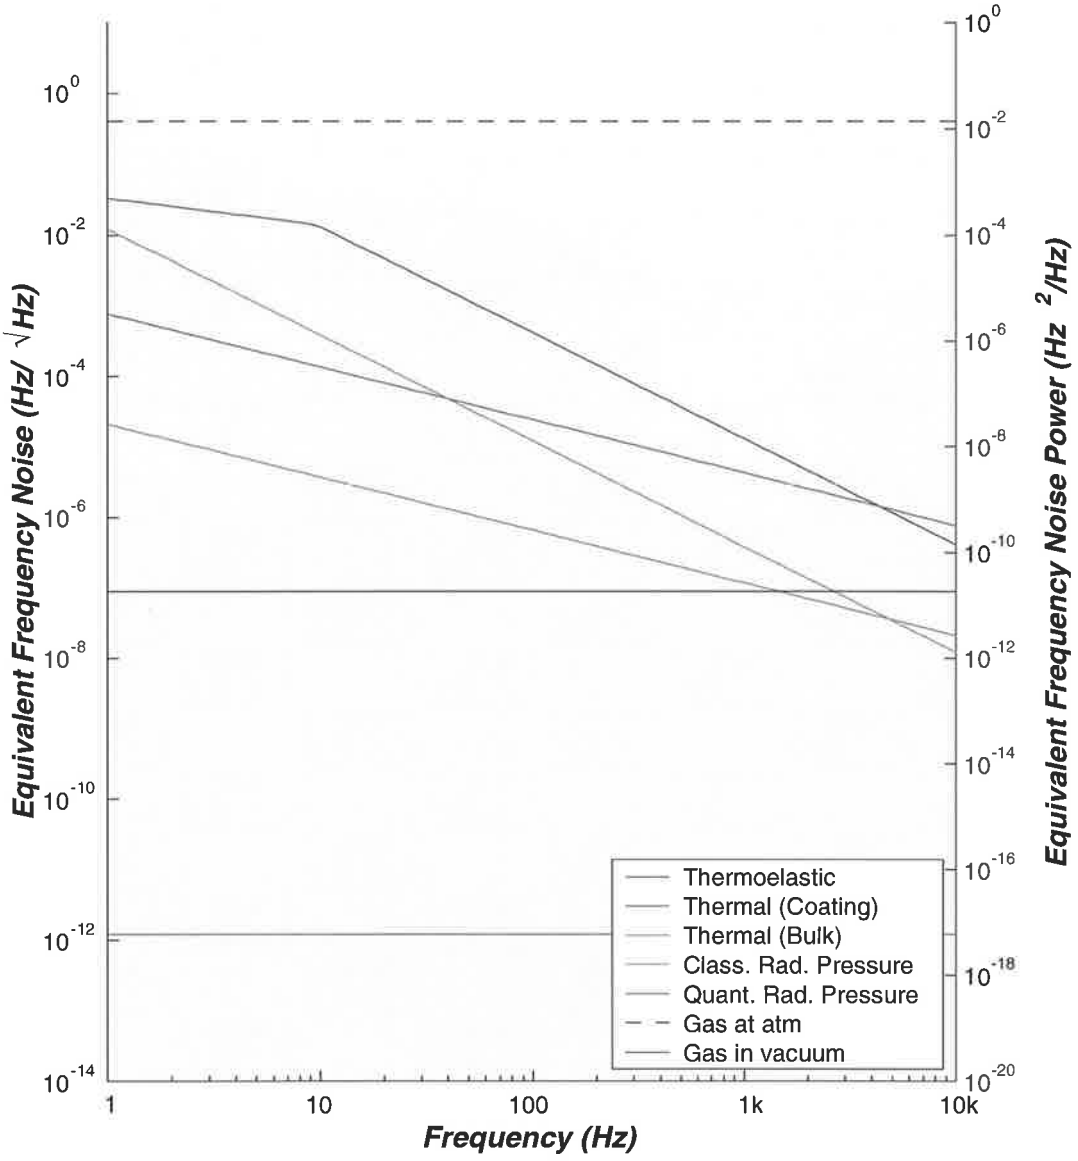

Figure 2.17: All estimated noise sources in the reference cavity. The background colours indicate the required sensitivity regions.

Chapter 3

Mechanical System

3.1 Introduction

The use of a monolithic test cavity results in an experiment which is relatively insensitive to seismic noise. Alternative topologies require significant effort to isolate the cavity from seismic disturbances and although the complex isolation systems used in other experiments were not required here, a simple isolation system was needed to decouple the test cavity from mechanical vibrations that might excite the internal vibrational modes of the test cavity.

This chapter will describe the design of the mechanical isolator used to support the test cavity. This isolator consists of a two stage scheme; a double sling of steel wire supports an intermediate mass from which a monolithic niobium flexure mount is used to hold the test cavity. Oscillations in the suspension are damped using eddy current damping.

3.2 Pre-Isolator

The laboratory in which this project was conducted was once home to a large UV monochromator system [49]. Part of the legacy of this experiment was a large seismic pre-isolator, consisting of a 30 ton concrete block mounted on nine critically damped air springs. Each air spring is made from a pair of pressurised air chambers connected by a narrow bore tube to provide damping.

The pre-isolator provides low frequency isolation between seismic noise and the experiment. A second advantage of the pre-isolator slab is that it provides a rigid structure on which the whole experiment is placed. On such a structure, low frequency seismic noise becomes common mode to the entire experiment, leading to an additional reduction in seismic noise at low frequency [50]. For our experiment this low frequency breakpoint will occur at a frequency in the few tens of hertz.

In additional to natural seismic noise, the pre-isolator also shields the experiment from man-made impulsive shocks transmitted through the ground. The experiment was located in the central business district of Adelaide, South Australia, a noisy environment with large mechanical shocks being transmitted through the earth. Such impulsive shocks are troublesome for the experiment as they can lead to unlocking of the various control loops and hence reduce the robustness of the system. This lack of reliability is particularly troublesome in a low frequency experiment because it makes it difficult to conduct the necessary long data collection runs. Unlocking events were reduced significantly when the pre-isolator was operational.

The vacuum chamber that will be used in the next phase of this project is positioned above the pre-isolator slab and mounted on large steel posts which are bolted directly to the laboratory floor. Two vacuum bellows separate the pre-isolator slab from the vacuum chamber, ensuring that any vibrations present in the chamber walls can not be coupled through to the enclosed work space. Figure 3.1 shows the mechanical layout of the pre-isolator and its connection to the vacuum system. The vacuum chamber and the design of its pumping system is described in detail in appendix C.

Figure 3.2 is a photograph that shows a section of the pre-isolator slab. The white concrete slab runs the entire length of the laboratory (about $\frac{1}{3}$ of the slab is in the foreground of the picture and is not visible in the picture) and extends for approximately 600 mm under the false floor. Seen in the distance is the vacuum chamber with an optical table directly in front of it. The upright grey posts along each side of the slab are supports for the original monochromator vacuum chamber that ran the full length of the photograph; the current vacuum chamber is one of four sections of the original chamber. As shown in figure 3.1, the support posts are

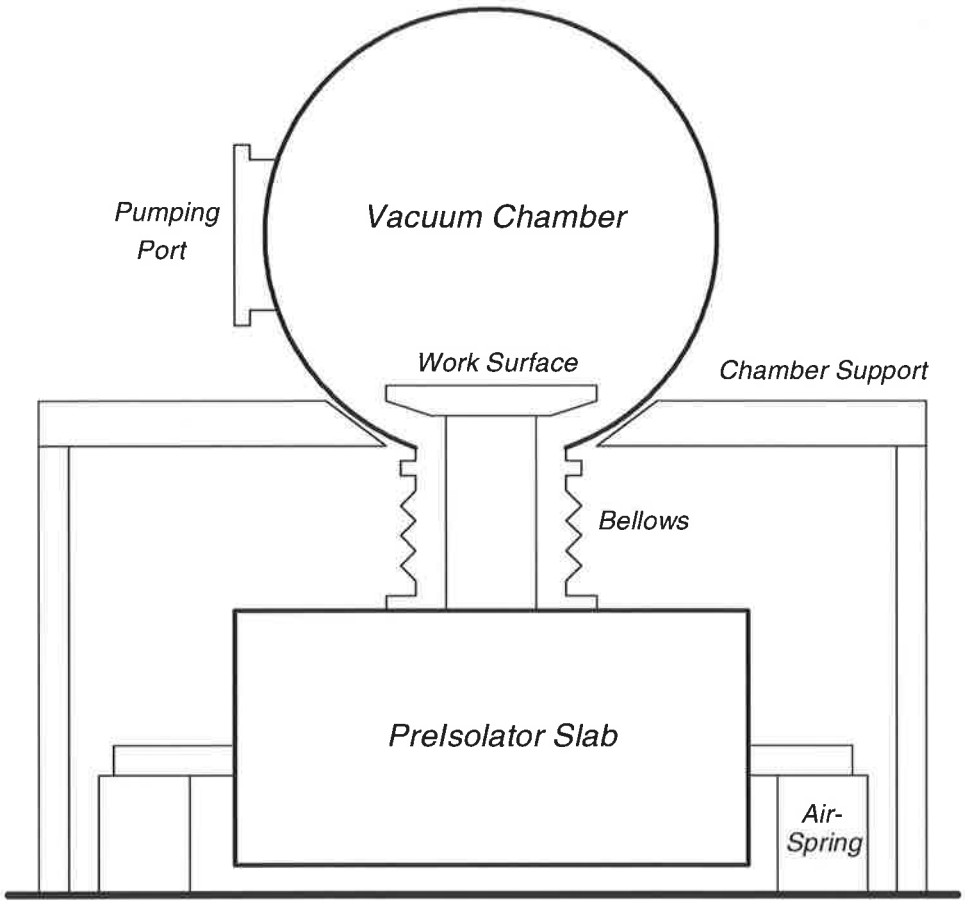

Figure 3.1: Cross-section of the vacuum system structure.

bolted directly to the floor.

After reconditioning of the air-springs, the slab was refloated and the seismic noise levels were measured with a seismometer. The results of this measurement may be seen in figure 3.3. The graph shows both vertical motion and horizontal motion in the north-south direction (the direction corresponding to the optical axis of the test cavity). These curves were taken during the day and thus reflect worst-case noise levels. As can be seen, the pre-isolator reduces the seismic noise level, although the effect is rather more modest than was hoped. The pre-isolator has a vertical resonance at approximately 10 Hz, beyond which it starts to reduce the seismic noise level. In the frequency band beyond 30 Hz the noise is reduced by a factor of approximately 4. Horizontal noise attenuation is better, having a

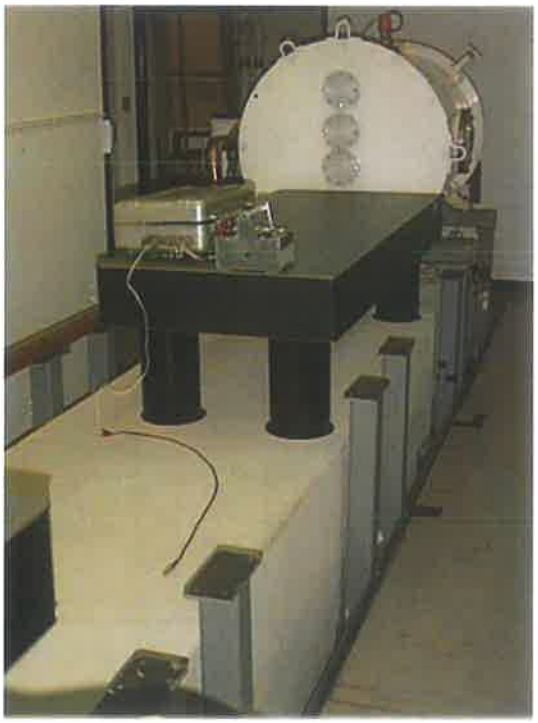

Figure 3.2: Photograph of the pre-isolator slab. The phot also includes the vaccum chamber and the optical table used in the experiment. Equiment seen on the table is seismic monitoring aparatus.

resonance at 6 Hz with the noise level dropping off to a factor of 10 below the ground noise. Although the calibration of the seismometer was unknown, a visiting TAMA delegation measured the background seismic noise level at 1 Hz at 6×10^{-6} $m/\sqrt{Hz}$ during the day. One would expect a better noise level during the night due to the different activity levels in the vicinity of the laboratory.

3.3 Test Cavity Suspension System

3.3.1 Introduction

The extreme precision required by the gravitational wave community has led to an explosion in seismic isolation system design in recent years. State of the art suspension systems are now complex electro-mechanical devices, using springs and pendula to provide passive isolation, high dynamic range piezoelectric transducers

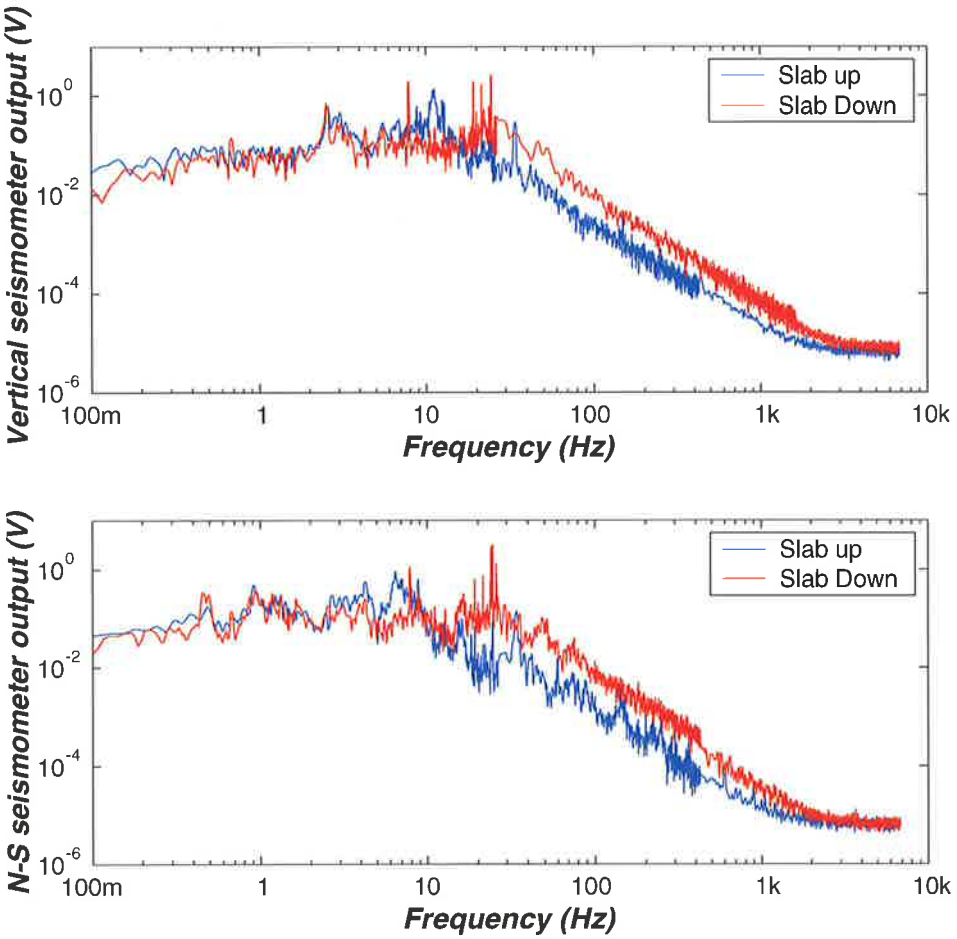

Figure 3.3: Effect of the preisolator on seismic noise levels. The top graph shows the vertical motion of the slab, the bottom graph shows horizontal motion along the north-south axis. In each case the blue curve is when the slab is raised on its air springs and therefore the pre-isolator is operational.

to actively reduce motion and voice coil actuators to damp resonances [51–57].

One of the advantages of using a compact monolithic test cavity in this experiment is the remarkable degree of insensitivity to seismic noise that it provides. Because of their common suspension, the two faces of the test cavity experience only common mode seismic excitation for all frequencies lower than the first mechanical mode of the test mass. The small size of the test cavity ensures that this common mode behaviour extends to high frequencies, where the seismic excitation is minimal in any case. Because the seismic noise is common to both faces of the test mass it does not produce changes in the cavity length. This leads to a very significant relax-

ation of the requirements for the vibration isolation required in the experiment, and indeed makes the experiment possible in the relatively small vacuum environment available.

The main task of the suspension system in our experiment is thus decoupling the test mass from its environment rather than isolating the system from seismic noise. Mechanical coupling through the test mass mounting inevitably spoils the test mass acoustic Q, leading to spreading of thermal noise energy away from the resonance frequencies and an increase in the rms thermal noise level.

The long-term aim of this experiment is to measure the noise levels of many test masses, thereby allowing the length noise to be characterised for a large number of materials and cavity geometries. Consequently, we require that various test masses be able to be swapped in and out of the system with relative ease. The design also needs to be robust to endure repeated handling. These requirements led to the selection of a monolithic flexure mount, a relatively new suspension type which will be described more fully in section 3.3.4.

Figure 3.4 shows a diagram of the overall system and indicates the main subsystems as described in the figure caption. Each of these subsystems is described in detail in the sections to follow.

Basic features of the suspension system are:

- the test mass is suspended from an intermediate mass with a small flexure mount (see section 3.3.4);
- the intermediate mass is suspended from a rigid tower using a double loop of stainless steel wire (see section 3.3.3) and
- the pendulum modes of the intermediate mass are damped using passive eddy current damping (See section 3.3.3.1).

3.3.2 Tower

The rigid frame of the suspension tower was made from aluminium. The frame consists of solid top and bottom plates with thick right-angle section posts at the

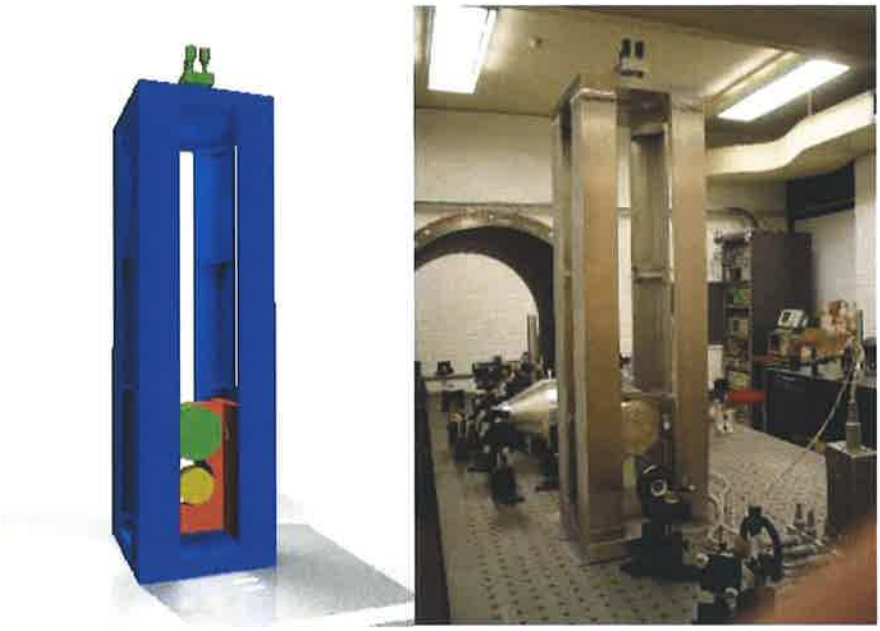

Figure 3.4: Schematic and photograph of the full suspension system. The tower is shown as blue, the reaction assembly red, the intermediate mass and its suspension system green and the test mass yellow.

four corners. Rigidity of the tower was paramount to ensure that it provided a stable base for the remainder of the system. The top and bottom plates were made from thick pieces of aluminium in order to keep internal resonance frequencies high. The corner posts were each fitted with two gussets to further increase their rigidity and each corner post was bolted securely to the top and bottom plates using two large bolts. The tower base provides mounting points for the damping subsystem as well as tie down points with which it was bolted to the optical table. The top plate houses the wire suspension system as described in section 3.3.3. Its design allows for the suspension of two independent masses, which allows an upgrade path to a suspension system utilising a reaction mass. The use of a reaction mass would allow the suspension damping to be isolated from ground and suspension tower resonances, a feature which may prove useful in future upgrades of the system.

3.3.3 Intermediate Mass Suspension

The intermediate mass was hung from the tower using a conventional two loop pendulum suspension. Each of the loops was made from 0.002 inch diameter stainless steel wire and had a length of 475 mm, yielding a pendulum frequency of 0.72 Hz. While it might have been desirable to provide more isolation by using longer suspension wire, this was impossible due to the size of the available vacuum tank.

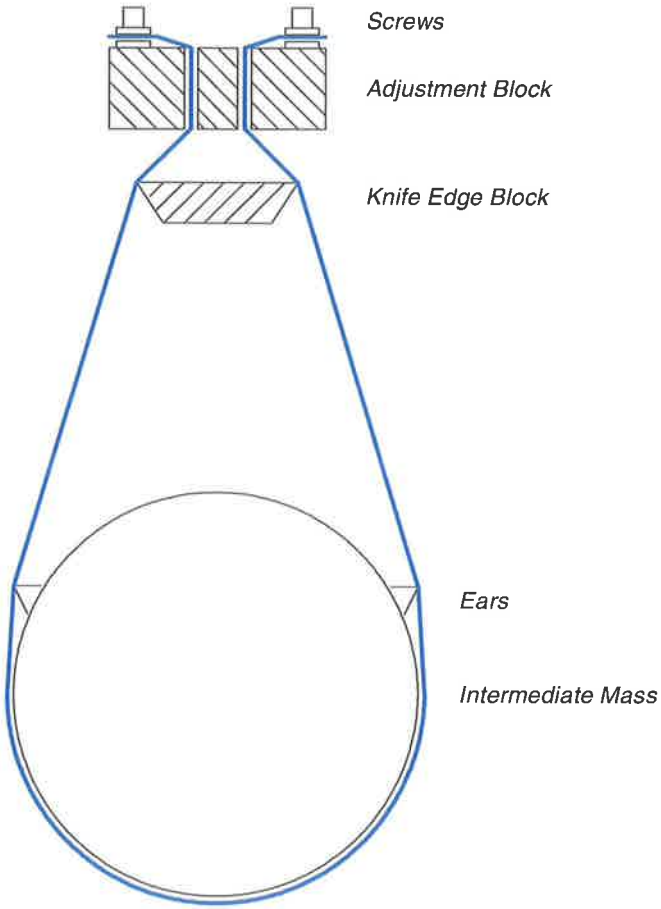

Figure 3.5: Diagram of the configuration of one of the two loops of wire used in the intermediate mass suspension system. The wire is shown in blue. Note that the distance between the suspension hardware and the intermediate mass has been reduced for the purposes of this diagram.

Figure 3.5 shows an overview of intermediate mass suspension design. The ends of each wire loop were attached to an adjustment block that sat on the top plate of the tower. The wire was then passed over knife edges and to the intermediate

mass. Contact with the intermediate mass was made using knife edge "ears" which precisely defined the contact point between the wire and its mounting. The contact point was sharp to prevent frictional losses caused by rubbing of the wires during movement of the intermediate mass. Any such losses lead to degradation of the acoustic Q of the intermediate mass and a resultant loss in performance of the suspension system [58, 59].

Figure 3.6 is a photograph of the assembly used to connect the wires to the tower. The assembly consists of two main parts; the black anodised knife edge block and the raw aluminium adjustment block. The adjustment block is fitted with three high precision lead screws with which the yaw, pitch and height of the intermediate mass can be adjusted. As can be seen in the photograph, the lead screws are not mounted in a right-angled triangle arrangement, which would provide independent adjustment of yaw and pitch. Stability requirements forced the back lead screw to be mounted centrally, leading to coupling of the tilt and pan adjustments. This coupling is relatively small and has not been found to cause a problem in practice. The adjustment screws are located in appropriately machined mounting points on the knife edge block: a cone, a vee and a flat were used to constrain the mounting in all degrees of freedom.

The suspension wires are firmly clamped to the top of the adjustment block by a screw at each end of the wire, allowing easy access during the wire mounting process. The wires then pass through a small hole in the adjustment block and then to the knife edge block. Figure 3.7 shows a close up of suspension wires crossing the knife edge block.

The knife edge block was made of anodised aluminium, ensuring that the knife edges are hard to prevent damage caused by the high pressures that are exerted by the wires when under load. If a soft knife edge were used then the extremely thin suspension wires can easily cut slots in the knife edge, resulting in spreading of the contact region between mount and wire [59]. This increase in contact area would lead to an increase in loss and a degradation of system performance.

The knife edges in our system were not damaged by loading; microscopic examination of the knife edge block after repeated and extended loading showed no

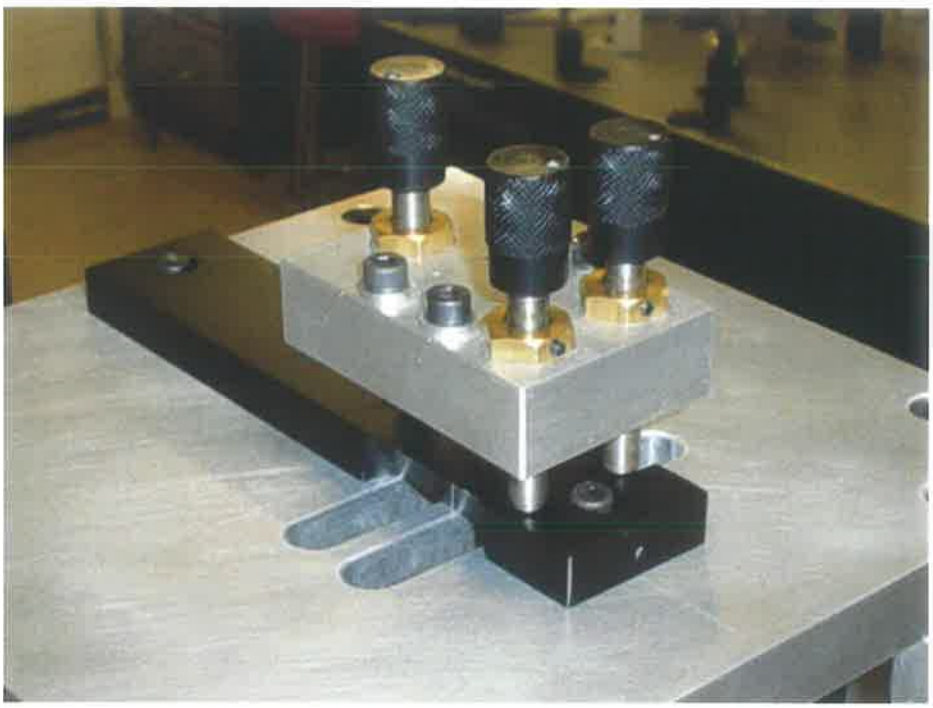

Figure 3.6: The adjustment block mounted on the knife edge block with the three high precision lead screws.

discernible degradation in surface quality. It was found, however, that the suspension wires were badly deformed at the high pressure point where they crossed the knife edge, leading to the production of a noticeable kink in the wire. Adjustment of the lead screws also led to wire damage as the wires caught and slipped on the knife edge while under tension. It was very important to ensure that any damaged wire sections not be used in the pendulum as it could degrade performance. Damaged wire was prevented from effecting the suspension by ensuring that the lead screws were only ever adjusted upwards. This upward adjustment pulled the damaged wire sections above the knife edge where their effect would be minimal.

At the intermediate mass, the wires again passed over knife edges as shown in figure 3.8. These knife edges were formed as part of small "ears" and were again machined from aluminium and anodised to increase their surface hardness. The knife edge ears were glued to the intermediate mass using Torr-Seal, ensuring mechanical rigidity and vacuum compatibility. Attaching the ears to the intermediate mass inevitably provides a source of mechanical loss for intermediate mass vibration and

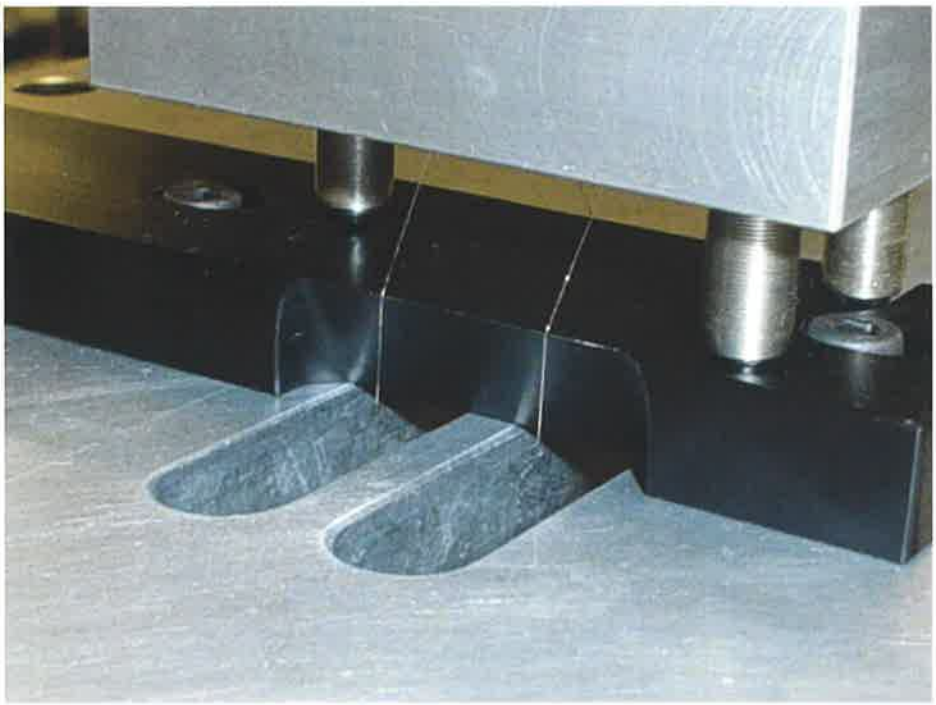

Figure 3.7: The suspension wires crossing the knife edge block. The wires can be seen passing through slots in the top plate of the suspension system on their way to the intermediate mass.

hence reduces the Q of the intermediate mass. However, it has been shown that gluing the ears helps to minimize this loss in Q, presumably by reducing sliding friction between the mass and its supporting ear [58]. Degradation of the acoustic Q was not critical, however, because of the common mode rejection of the test cavity and because of the presence of an additional isolation stage between the intermediate and test masses.

3.3.3.1 Damping Assembly

Low loss oscillators can be excited easily and have long decay times. These strong resonance effects are unacceptable in a suspension system as they lead to large uncontrolled motion of the suspended optics at the resonance frequencies. Such oscillations causes difficulties in alignment, lock acquisition and stability and in low noise measurement. Damping of suspension resonances is therefore required in all seismic isolation systems. Eddy current damping was chosen to damp pendulum res-

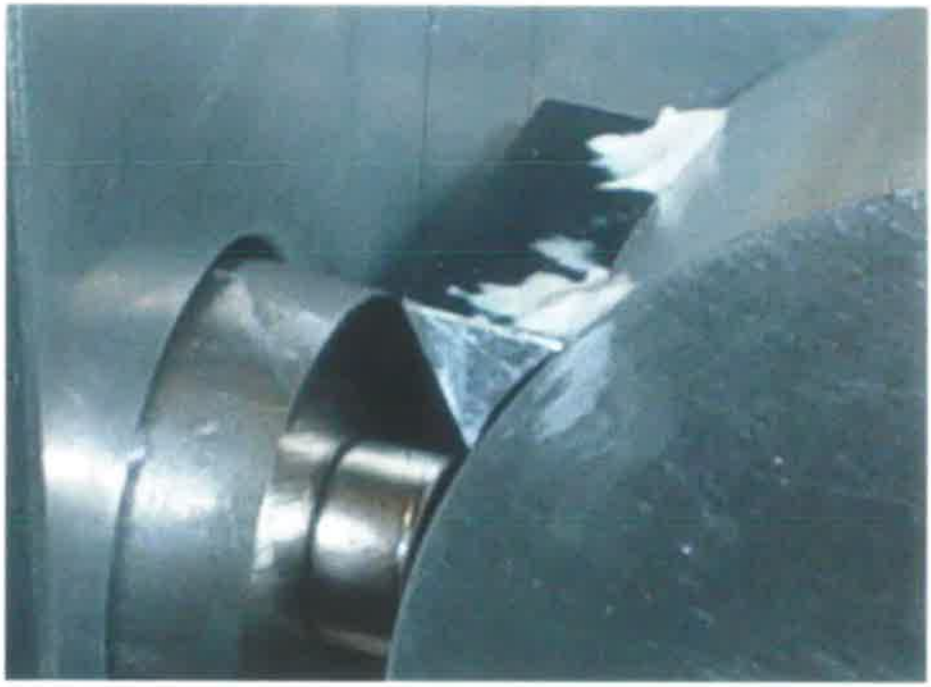

Figure 3.8: Knife edge "ears" mounted on the cylindircal intermediate mass (lower right of picture). White Torr-seal adhesive is clearly visible in the photograph. There is an additional single spot of adhesive on the underside of the ear.

onances in this experiment as it is simple, cheap and robust. Eddy current damping has long been used in a diverse range of instruments, such as seismometers and electron microscopes. The simplicity and effectiveness of the technique made it a natural choice for early gravitational wave detectors [60]. Unfortunately, motion of the damping magnets is coupled into suspended mass motion by the eddy current damping mechanism [60],which allows mechanical noise to "short-circuit" the isolation system. In the current state-of-the-art damping systems eddy current damping is confined to early stages of suspension systems, where its large possible dynamic range is particularly valuable and the coupling is not so troublesome [61]. Research on how best to utilise eddy current damping in gravitational wave experiments is ongoing [61,62].

In the later stages of large vibration isolation systems, eddy current damping has now been widely replaced by voice-coil actuators [45,51,63] and a new "active" eddy current scheme [64], both of which offer tunability of the damping. In tunable

systems the damping only acts at those frequencies which correspond to mechanical resonances of the structure being damped and hence the tendency to couple extraneous motion into the suspended mass is reduced. As the requirements for our suspension system were relatively relaxed it was deemed sufficient to use the simpler eddy current damping system. Provision was made, however, for the provision of an active system at a later date, should it become necessary.

One scenario leading to the production of eddy currents is the movement of a conductive body through a magnetic field. As the material moves it carries along the unbound electrons present in its conduction band. These conduction charges experience the magnetic component of the familiar Lorentz force ($\mathbf{F} = q\mathbf{v} \times \mathbf{B}$) due to their motion with respect to the field. This force in turn leads to a bulk flow of electrons in the reference frame of the conductor. In a conductor having finite conductance this eddy current dissipates energy by ohmic heating. The eddy current therefore withdraws energy from motion of the conductive body and dissipates it as heat within the conductor, leading to a braking force on the conductor's motion. As can be seen in the Lorentz force equation, the dissipation will be proportional to the velocity of the conductor and the force produced will therefore act as a viscous force on the conductor.

Four strong zinc coated Nd:FeB magnets were used to realise our damping scheme. Each of the cylindrical magnets had a diameter of 10 *mm* and a length of 5 *mm*. The arrangement of the four magnets is shown in figure 3.9. As can be seen, two magnets are used to damp both yaw and axial pendulum motion of the intermediate mass. One magnet was used to damp the sideways motion of the mass and a fourth used to damp the pitch degree of freedom. These magnets are denoted as **L**eft and **R**ight, **S**ide and **T**op respectively. It was found that the damping efficiency of the pitch mode was increased if the intermediate mass was lowered slightly so that the L and R magnets did not act through the centre of mass of the intermediate mass; this alteration did not effect the efficiency of either yaw or pendulum damping. A photograph of the damping arrangement is shown in figure 3.10. The Left, Right and Top damping magnet assemblies can be seen in the photograph. The side magnet is also visible in figure 3.8.

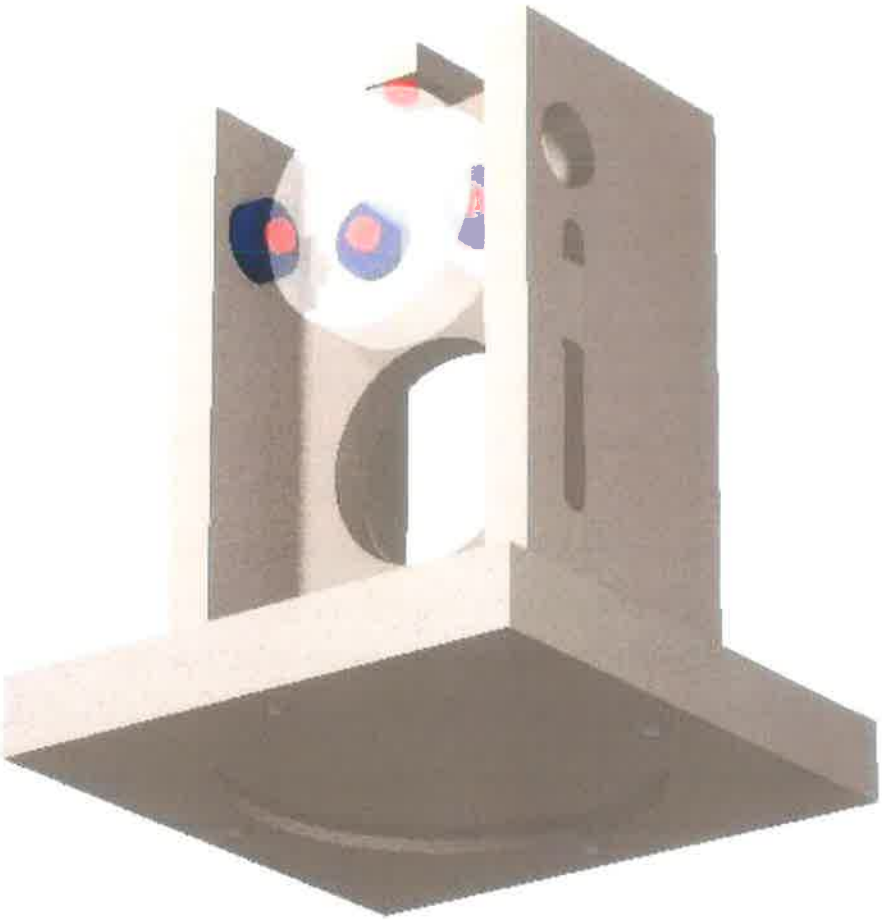

Figure 3.9: The reaction assembly with the location of the damping magnets highlighted. The four magnets are indicated in red. Steel spacer pieces are coloured blue. The position of the intermediate mass is indicated by the transparent disk.

The magnets were held in a solid reaction assembly as shown in figure 3.9. This assembly was made from 15 *mm* thick aluminium plate and securely bolted to the tower base. Mounting the damping magnets on an independent base rather than attaching them to the tower ensures that tower resonances do not move the damping magnets and couple into intermediate mass motion.

The mounting holes in the reaction assembly were made considerably larger than the size of the magnets so that an active system could be retrofitted if needed. The holes were sized to suit the sensor/actuator modules used by LIGO, a set of which was kindly provided for our use by the LIGO team. These damping devices consist of an optical sensor and a voice coil mounted on a common body. These devices

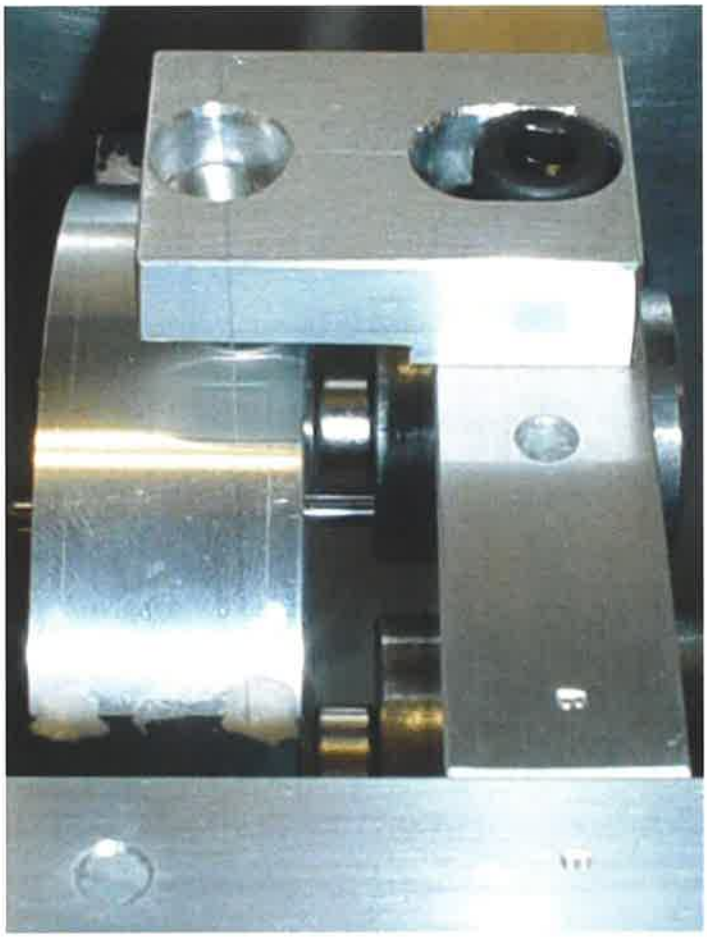

Figure 3.10: Arrangement of damping magnets around the intermediate mass. The Left and Right magnets are visible as the two silver cylinders in the centre of the frame. The top magnet sits underneath the projecting finger at the top of the photograph.

are known as **O**ptical **S**ensor and **E**lectro **M**agnetic actuators or simply "OSEMs" (see [65] for a general description of OSEM design). Because the OSEMs defined the size of the holes in the reaction assembly, OSEM sized pieces of steel were used to hold the magnets for the eddy current damping in use during the experiments described in this thesis. Details of the OSEMs used may be found in a LIGO document [66] and references thereon.

As shown in figure 3.9, there is a large hole in the reaction assembly at the height of the test mass. This aperture allows the mounting of long test cavities that need more space than can be provided in front of the reaction plate. A slot cut into

one side of the reaction assembly allows access to facilitate installation of the test masses.

Two sets of safety pins were built into the reaction plate. These pins were designed to catch the intermediate and test masses should the suspension fail. The pins were located so that a failure in the suspension wires would lead to the intermediate mass being caught by its safety pins. Only in the unlikely event of a flexure mount failure would the test mass safety pins be required.

3.3.3.2 Damping Performance

Five LIGO OSEMs were mounted in a test jig so that motion of the intermediate mass could be measured. The tests were conducted both with and without the damping magnets present. The OSEMs were configured as simple light gates as shown in figure 3.12 below; the actuator part of the OSEMs were left unconnected. Five non-reflective metal flags were temporarily attached to the intermediate mass and an OSEM was used to monitor each flag's position. The location of the sensors around the mass is shown in figure 3.11. The four OSEMs on the front face were each placed 20 *mm* from the centre line of the intermediate mass, the side OSEM was placed in the centre of the side. This arrangement of sensors was sufficient to extract information about the axial pendulum, sideways pendulum, pitch and yaw modes of the intermediate mass. The OSEMs can be considered in three groups, each of which would ideally be sensitive to different mechanical modes of intermediate mass suspension as shown in table 3.1. In practice there was found to be some cross coupling between the modes.

OSEM	Mode
L,R	Axial (Z), Yaw (Y)
T,B	Axial (Z), Pitch (P)
S	Side (S)

Table 3.1: OSEM/Mode correspondence

As the flags moved they interrupted a varying amount of the light beam produced by the infrared LED in the OSEM. This modulated the amount of light falling on

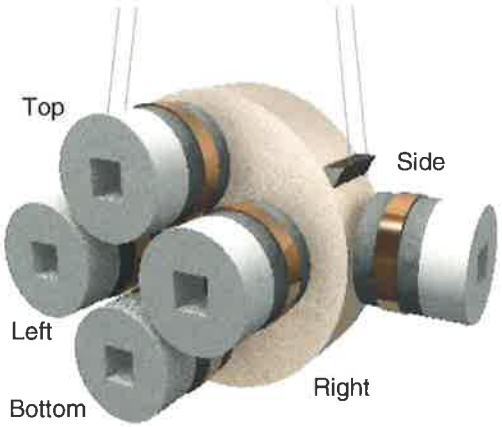

Figure 3.11: Positions of the five OSEMs used for testing of the suspension system damping

the phototransistor, leading to changes in photocurrent. A simple transimpedance amplifier was then used to convert the photocurrent to a voltage output which was then used for analysis. The OSEM drive circuit also included constant current drivers for the OSEM LEDs and is shown in figure 3.13.

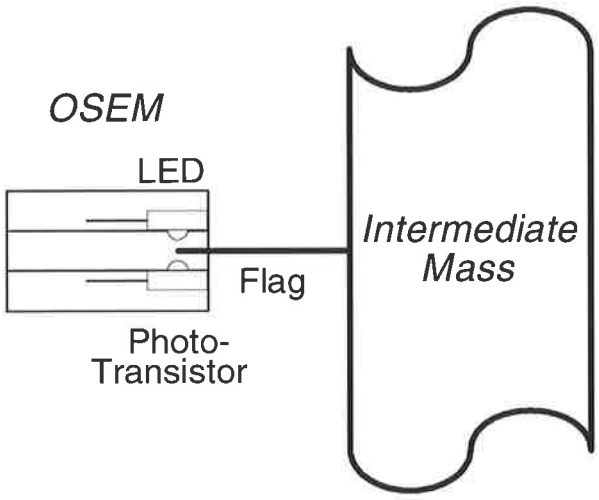

Figure 3.12: Operation of the OSEM as a simple light gate.

The manufacture of both LEDs and phototransistors leads to large process spread, often resulting in the parameters of such components varying quite markedly between devices. Each OSEM therefore required individual calibration. The OSEMs were mounted in a jig and a micrometer mounted flag was moved through the active

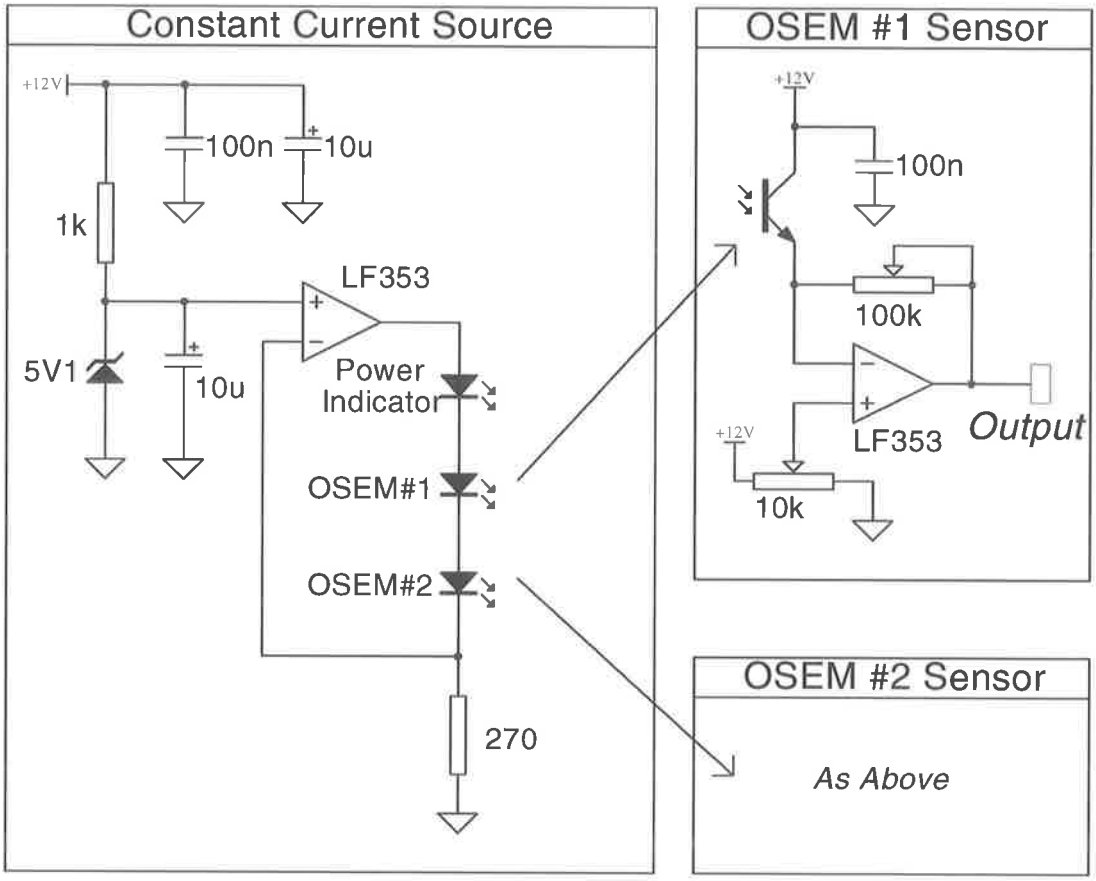

Figure 3.13: Schematic of the circuit used to drive the optical sensor component of the LIGO OSEMs for use in the damping measurement experiments.

region of each OSEM. Output voltages and micrometer positions were recorded at each position and are plotted in figure 3.14.

The OSEM responses are all of similar shape although the details of the curves are slightly different. Note that the absolute values on the x-axis of the graph are meaningless; they simply reflect the location of the various OSEMs within the testing jig. The Bottom OSEM (represented by circles rather than crosses on the graph) shows slightly different behavior from the other four OSEMs. Inspection of the Bottom OSEM showed that it was of a different model to the others and so the different curve shape is perhaps not surprising; a different LED or detector device may well have been used in this OSEM.

All five OSEMs displayed linear operation over a significant part of their response. Visual inspection suggested that linear range extends from approximately -5 V to 5 V for the Bottom OSEM and from approximately -7 V to 7 V for the other sensors. Linear least square fits were carried out over the linear regions to determine the slopes of the lines, the results of which are also shown in figure 3.14 and in table 3.2.

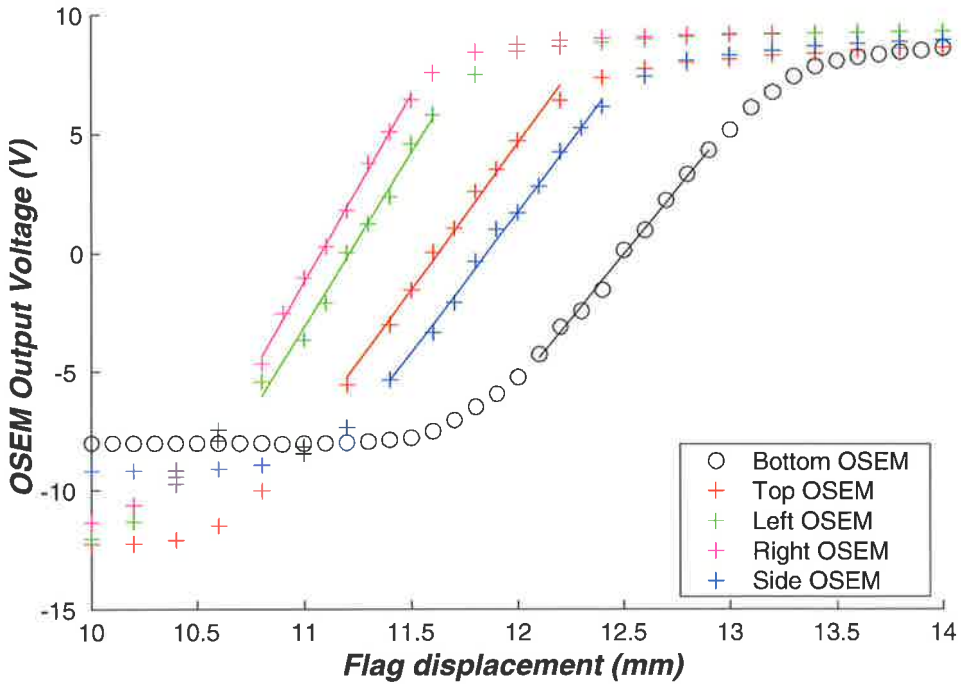

Figure 3.14: Calibration curves for the five OSEMs. Raw data is shown as well as linear least square fits to the linear region of each OSEM. The absolute value of the x-axis is arbitrary as discussed in the text.

OSEM	Response (V/mm)
Bottom	10.9
Top	12.3
Left	14.6
Right	15.7
Side	11.8

Table 3.2: OSEM calibration constants

Spectra of the five OSEMs output voltages were recorded and converted to displacements using these calibration factors. Figure 3.15 shows the spectra for the Top, Left and Side OSEMs. This plot in conjunction with table 3.1 above, allows identification of the frequencies that correspond to the different modes of the intermediate mass. The results of this procedure allow the construction of table 3.3.

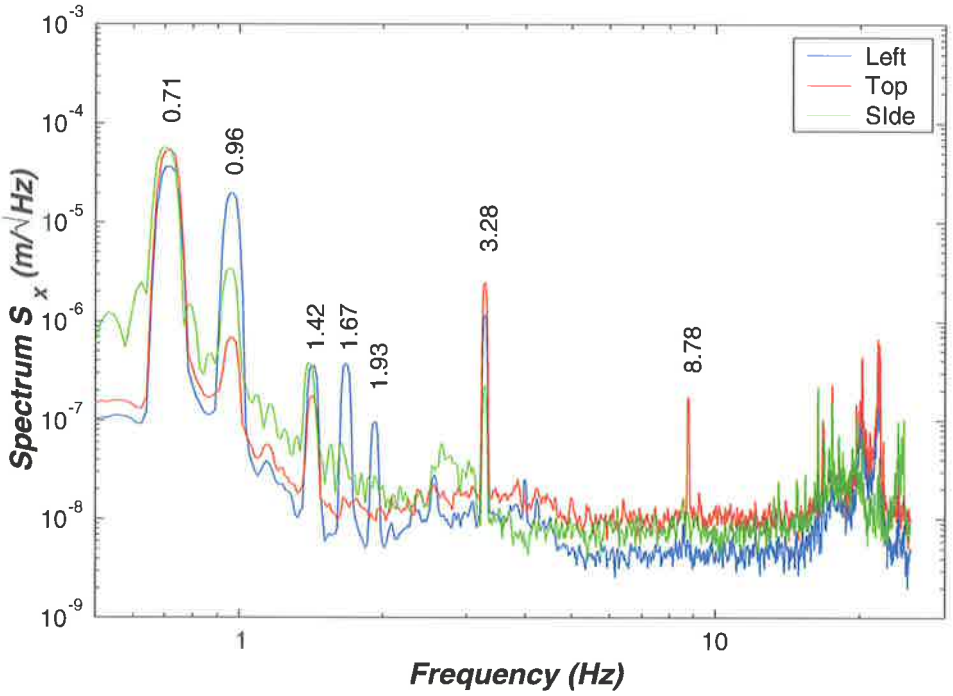

Figure 3.15: Calibrated open loop spectra of three OSEM output spectra allowing identification of modes.

Frequency (Hz)	Dominant OSEM	Mode	Harmonic
0.71	All	z,S	1^{st}
0.96	L	yaw	1^{st}
1.42	L,S	z,S	2^{nd} (2×0.71)
1.67	L	yaw	1^{st}
1.93	L	yaw	2^{nd} (2×0.96)
3.28	T	pitch	1^{st}
8.78	T	pitch	1^{st}

Table 3.3: Position and nature of the intermediate mass resonances

Closed loop spectra were taken for each of the five sensor with Eddy current damping in operation and the results were compared with the open loop spectra. Comparison between the L and R sensors allows identification of the yaw mode because motion in the yaw mode motion should be correlated in the two sensors. Similarly pitch should be correlated in the T and B sensors. Figure 3.16 shows the results from the L and R sensors, figure 3.17 the results from the T and B OSEMs. These graphs also include coherence plots which indicate when the two OSEM outputs are in agreement; a coherence near one indicates that the signal is due to a real spectral feature that is present in both channels, a lower coherence indicates that the measured spectral level is more likely due to noise. The included cross spectral plots are a convenient method for showing those features that are common to both OSEMs as features in the OSEMs output that didn't have a high coherence would not appear in the cross spectrum.. The sensor noise floor was measured with the OSEM connected to its drive circuit but no flag present to interrupt the optical sensor. The resultant noise floor measurement is also shown in the figures.

Open and closed loop curves for the Side OSEM are shown in figure 3.18.

Each of the graphs show a number of similar features; the open loop curves show a broad low frequency noise component with a number of narrow, poorly damped resonances. At high frequencies there is an additional complex of lines. Once damping is applied the suspension resonances are markedly reduced, in most cases disappearing from the traces. Even with damping present, the broad low and high frequency noise features remain, suggesting that they were **not** due to resonances in the suspension pendulum. These noise features were probably due to mechanical resonances in the comparatively flimsy test jig frame or electronic noise.

The noise spectra shown in the figures correspond to the suspension system operating in a quiet laboratory environment. Experiments showed that air currents were the primary means of excitation of the suspension modes, although acoustic excitation was also significant. Both of these forms of excitation would be absent in the vacuum environment, so a better noise level is be expected once the experiment has been moved into the vacuum chamber. The responses presented here should not therefore be regarded as accurate estimates of the final motion of the intermediate

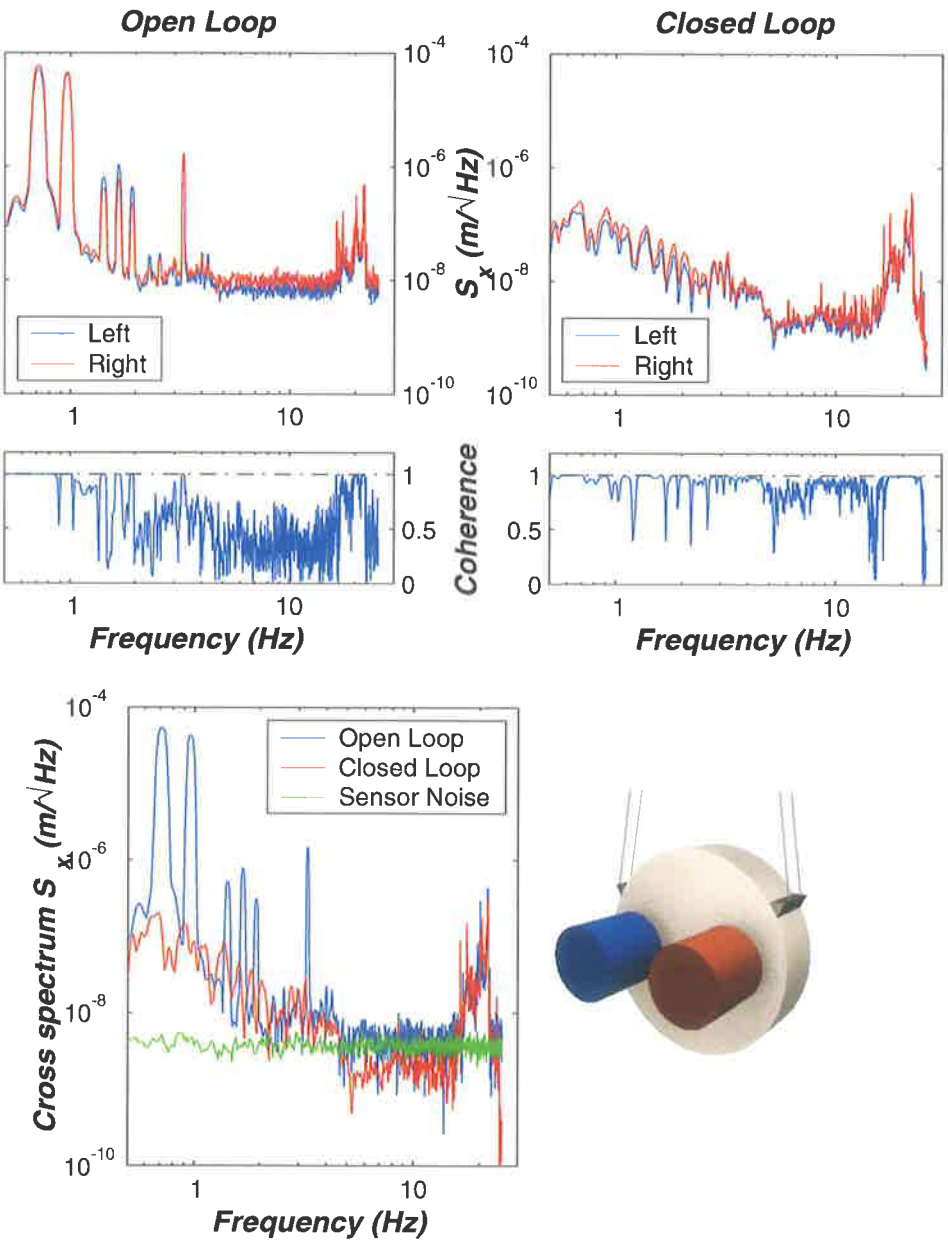

Figure 3.16: Comparison of open and closed loop curves for the L and R sensors.

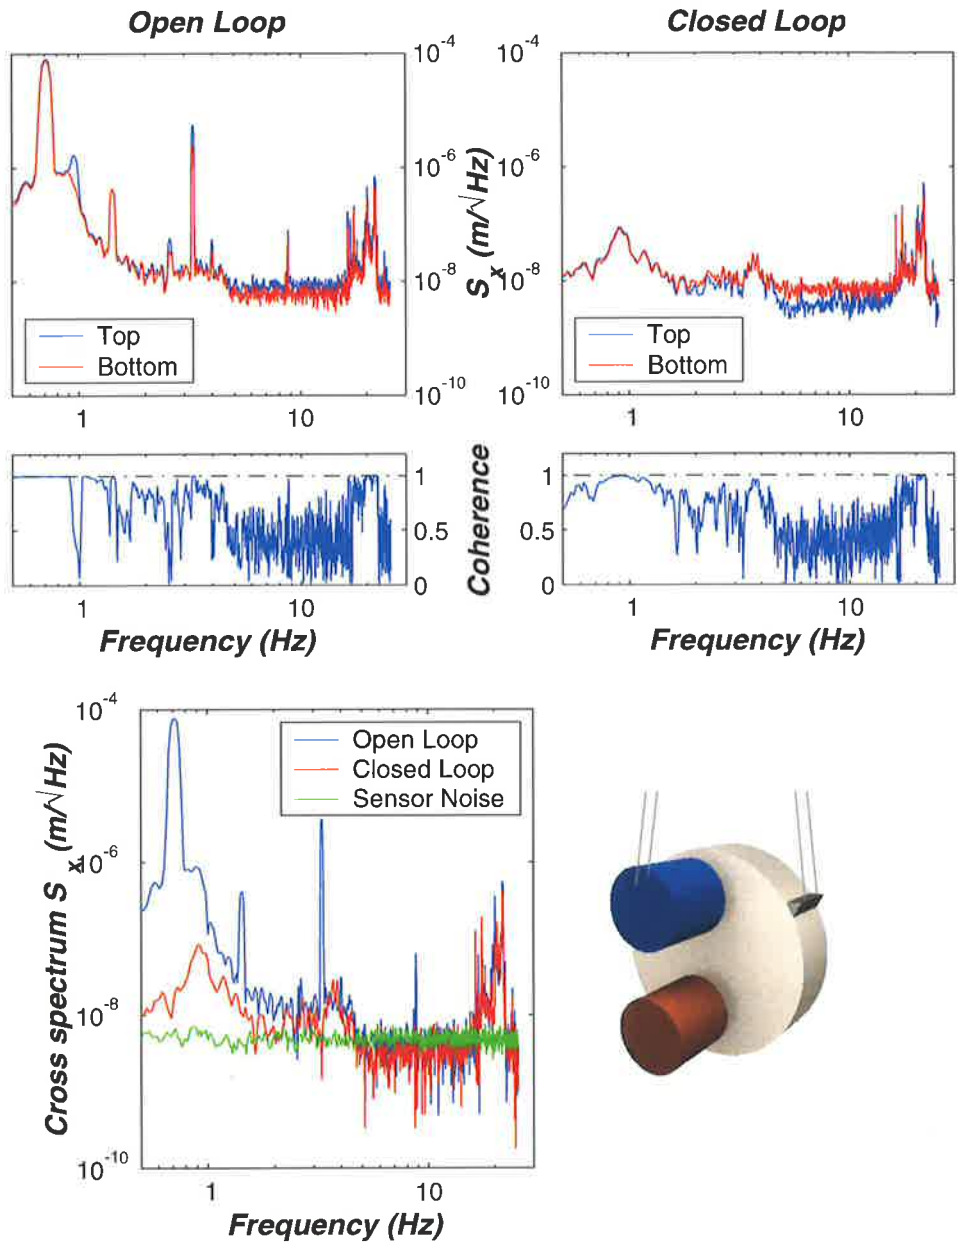

Figure 3.17: Comparison of open and closed loop curves for the T and B sensors.

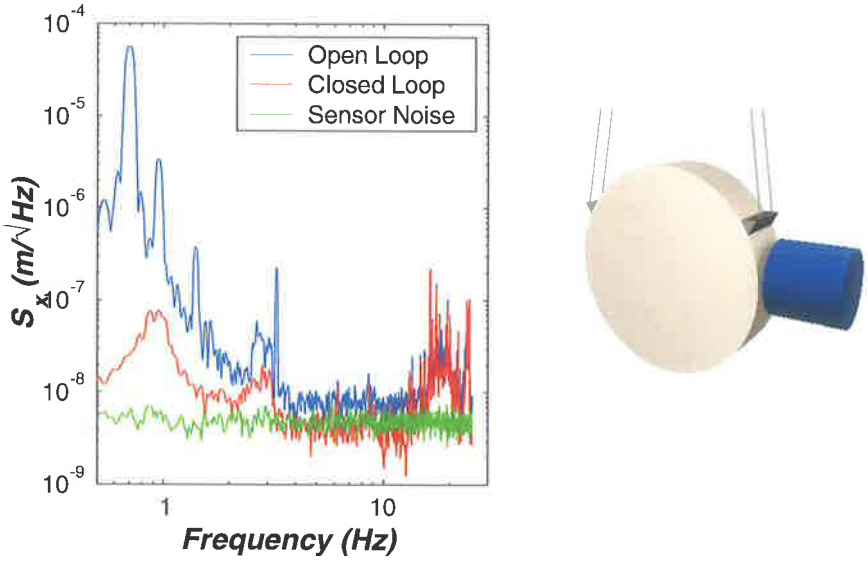

Figure 3.18: Comparison of open and closed loop curves for the S sensor.

mass.

A second measurement of the damping effectiveness was obtained by measuring several transient responses. This allowed the ringdown of various modes to be observed. The pendulum mode of the intermediate mass was extremely lightly damped in the open loop case. This was expected, as the low damping indicates low dissipation in the pendulum and hence good localisation of pendulum thermal noise around the resonance frequency. The pendulum was set in motion and the amplitude of the resultant oscillation was measured by recording the amplitude with the top OSEM every 30 seconds for 14 minutes. Measuring a single OSEM output was sufficient for this test because the pendulum mode decayed much more slowly than other modes. The system was therefore allowed to settle for a few minutes after excitation so that the undesired modes had time to decay. The ringdown curve obtained from this experiment is shown in figure 3.19 and yields a time constant of 650 s. The damping ratio ζ can be determined from the time constant and the pendulum decay equation $A(t) = A_o \frac{1}{\sqrt{1-\zeta^2}} e^{-\zeta \omega_n t}$. Fitting of this exponential curve to the data yields an open loop damping ratio estimate of 350×10^{-6}. This corresponds to a Q of approximately 2800.

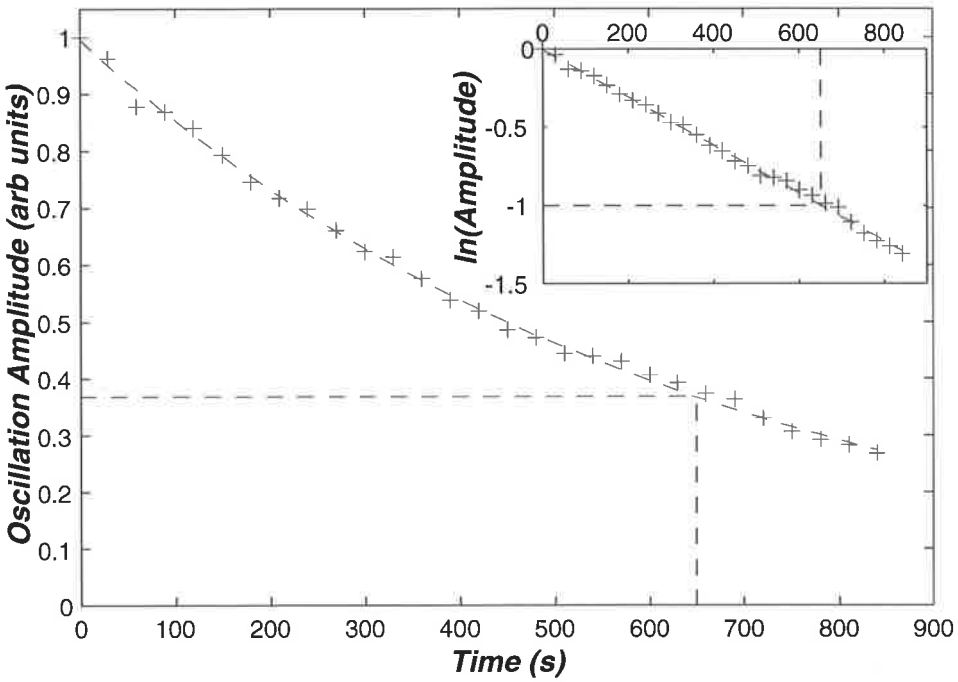

Figure 3.19: Decay of the pendulum mode of the intermediate mass suspension. Insert shows the natural logarithm of the amplitude to confirm the exponential nature of the decay. Dashed black lines indicate the 1/e point of the decay and the correspnding time constant. The dashed blue line denotes the fitted theoretical curve.

Figure 3.20 shows the raw data for this same resonance when damping was added. In this experiment the intermediate mass was again impulsively excited. Although attempts were made to drive only the pendulum mode, other modes were also inevitably excited by this procedure. Data was collected using the top and right OSEMs as the positions of these sensors ensured that they sampled different contaminating modes. The analysis of the two data sets can therefore be expected to yield different estimates of the damping.

It is immediately apparent from a comparison of this data with that shown in figure 3.19 that there had been a marked increase in damping. Because the quantisation noise present in the data, the data was preprocessed by filtering, removing the dc component and taking the absolute value of the data. The results of this process are shown in figure 3.21 for the right OSEM and figure 3.22 for the top OSEM. The time constants obtained for the two cases are now $0.58s$ and $0.47s$ respectively as

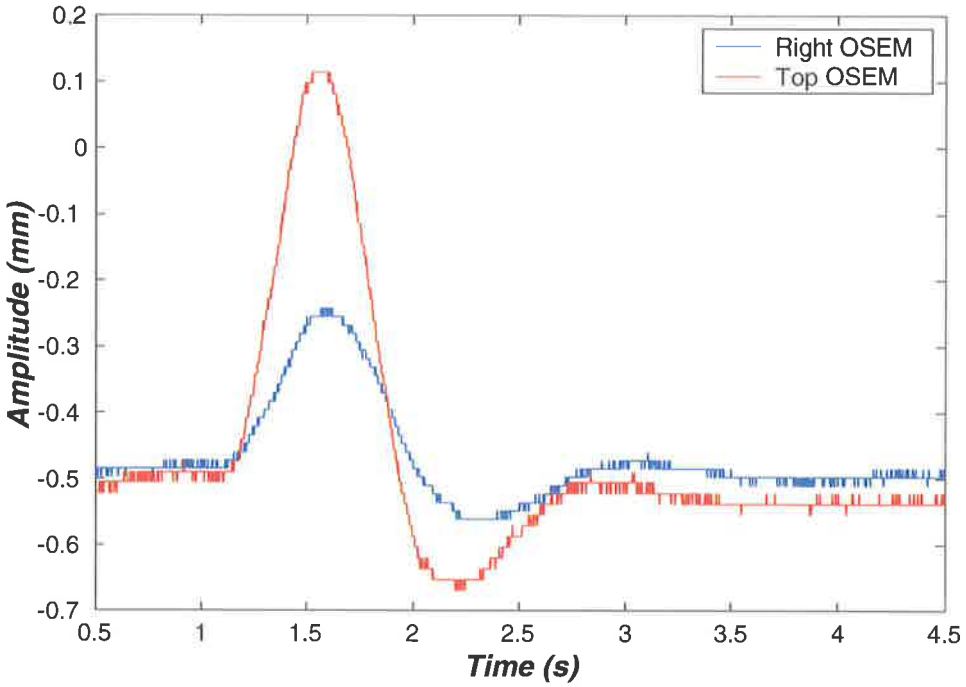

Figure 3.20: Raw data obtained for decay of the closed loop pendulum mode. Blue curve shows the motion of the right OSEM, red curve is that of the top OSEM.

compared to the previous 650 s. This reduction of the time constant corresponds to an increase in ζ from 350×10^{-6} to 0.42 and 0.58 for the two estimates.

Although the pendulum modes were most effectively damped by the damping system, the yaw and pitch modes were damped too. For example, figure 3.23 shows the response of the system to a pitch impulse in both open and closed loop configurations. As can be seen, excitation of the pitch mode also led to injection of energy into the low frequency pendulum mode. In the open loop case both modes oscillate with a steady amplitude on the time scale of the graph. In the closed loop configuration the low frequency pendulum mode is damped very quickly in accordance with the analysis above, while the high frequency pitch mode continues for rather longer.

The poorer damping of the yaw and pitch modes is mitigated by the fact that they occur at higher frequencies than the pendulum modes and should therefore be exposed to reduced levels of excitation. It is also expected that the geometry of the suspension support should ensure that coupling of seismic noise into the yaw mode

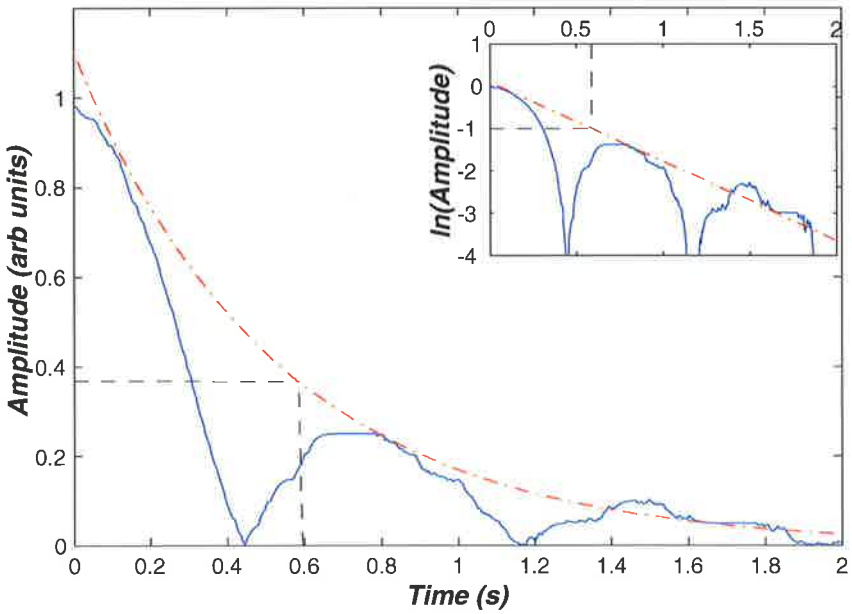

Figure 3.21: Decay of the pendulum mode as seen at the right OSEM. Fitted curve shown as the red dashed line. Black dashed lines show the point where the envelope has decayed to 1/e of the inital value. The inset shows the same data but uses the logarithm of the envelope amplitude.

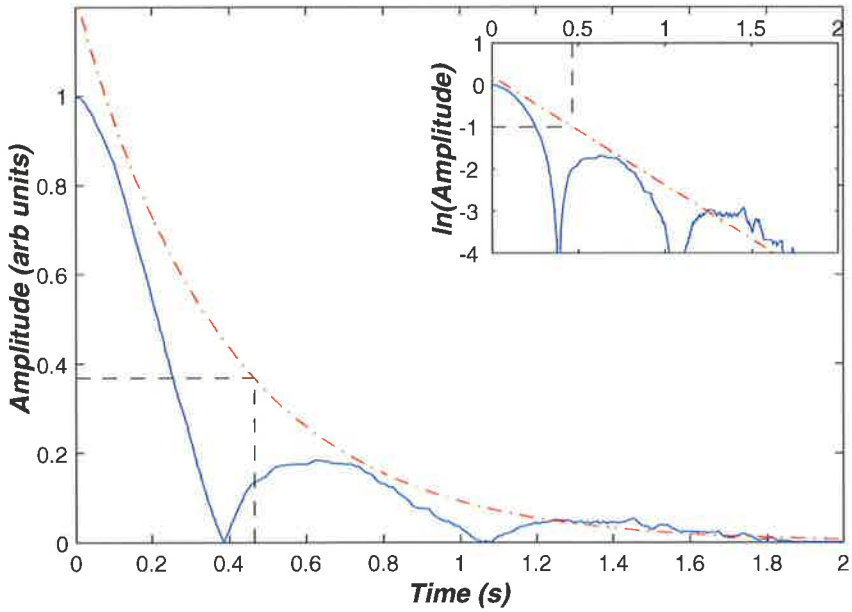

Figure 3.22: Decay of the pendulum mode as seen at the top OSEM. Fitted curve shown as the red dashed line. Black dashed lines show the point where the envelope has decayed to 1/e of the inital value. The inset shows the same data but uses the logarithm of the envelope amplitude.

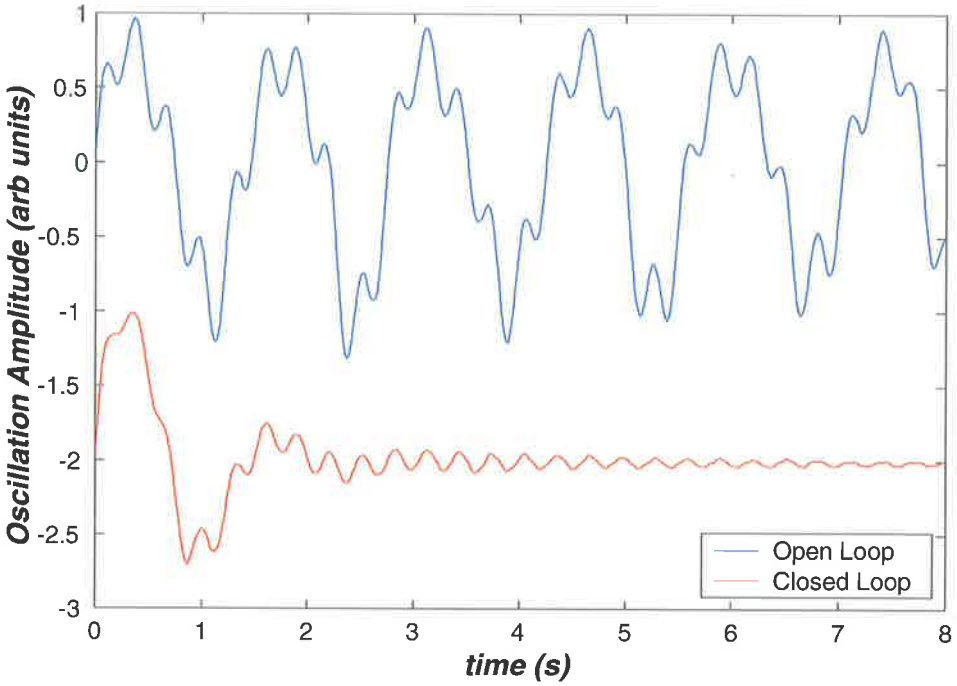

Figure 3.23: Open and Closed loop ringdown for the top OSEM after deliberate excitation of the pitch mode.

and particularly the pitch mode would be relatively inefficient. The combination of reduced drive and low coupling efficiency for these two modes should lead to small resonance amplitudes.

3.3.3.3 Conclusion

Damping of the intermediate mass motion by eddy current damping has been shown to effectively damp suspension resonances. The tests were conducted in air and as such were subject to relatively high level of modal excitation *via* air-currents and acoustic noise. The jig used to hold the position sensors used in analysis of the motion was not particularly rigid and was likely to have introduced spurious features into the spectra of the motion of the various suspension modes. The combination of these effects is sufficient to require retesting of the suspension performance once the suspension has been moved into the vacuum environment.

3.3.4 Test Mass Suspension

Most current suspension systems use a double pendulum linkage between the intermediate mass and the test mass [58, 60, 67]. The pendulum wire has traditionally been made of steel [68, 69] although the use of fused silica fibres to reduce degradation of the test mass acoustic Q is now being investigated [70–72]. The point of attachment between the test mass and its suspension is inevitably a source of mechanical loss and hence degrades the performance of the system. As it is crucial that losses in the test mass be minimised, the design of optimal mounting systems remains a topic of vigorous research.

The gravitational wave group at the University of Western Australia (UWA) has developed flexure mounts to replace the more traditional pendula [73]. The flexure mount is made from a very thin membrane of niobium (approximately 70 μm thick) that is rigid in one dimension but flexible in the other. Each end of the membrane is expanded into a mounting wedge that is held fast by small grooves cut into the intermediate mass or the test mass. These grooves are termed "dovetails" by the group at UWA, although they are not shaped as dovetails, but as two angled cuts (see figure 3.24 below); we follow the UWA nomenclature here.

Flexure mounts have several significant advantages when compared to pendulum alternatives, most notably their robustness, small size and good thermal conductivity [74]. Niobium is used for the flexure material because it exhibits the lowest acoustic loss of any metal, which ensures good Q performance. The flexure mounts have been found to possess excellent pendulum mode Q [75,76] and maintain excellent test mass Q [74]. Flexure mounts have the additional advantage of allowing test masses to be changed easily by lifting a test mass off the flexure mount and replacing it with another. This simplicity should be contrasted with the time consuming balancing and realignment procedures required when rehanging a pendulum.

A small flexure mount was generously provided by the group from UWA and was used throughout the experiment. It should be noted that the flexure was not specifically designed for our experiment; in particular the flexure mount was designed to hold a significantly larger test mass (4 kg as opposed to the 630 g test mass used in this work). We would expect that the different loading on the flexure mount

would lead to some change in the reported characteristics of the mount, although the variations should not be significant for our purposes. The different loading conditions would lead to some minor shifting of the flexure mount internal resonances, but as these resonance frequencies are very high (100s of kHz) they do not affect this experiment.

A diagram of the flexure mount is shown in figure 3.24. Each flexure is fabricated from a single piece of niobium using electron discharge machining. The flexure is then treated to relieve mechanical stress and restore its low loss properties.

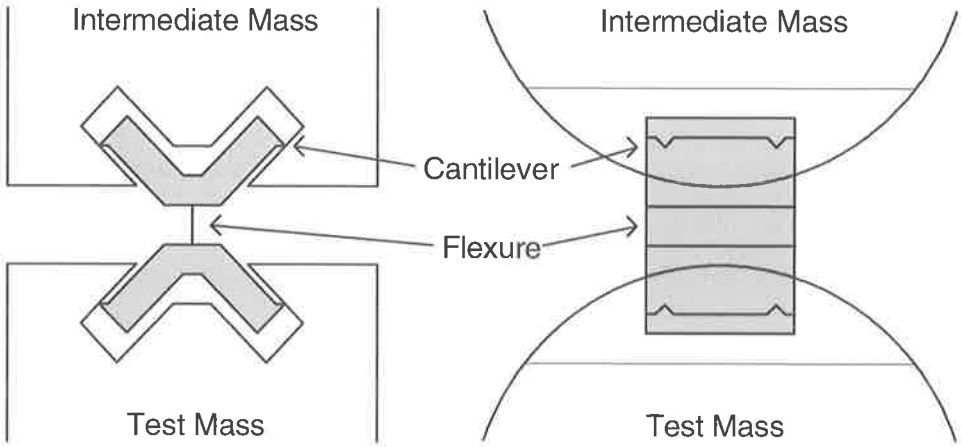

Figure 3.24: Diagram of the flexure showing its mounting into the intermediate and test massses.

The microcantilevers that form the contact points between the flexure and the two suspended masses are indicated in figure 3.24. These cantilevers are crucial to maintenance of high test mass acoustic Q as they produce high pressure boundary regions that minimise slippage and hence reduce losses [73]. Care was taken to ensure that the load was borne by all four cantilevers. Should either set of contacts be made with only three of the cantilevers then significant stick-slip losses would occur at the unused cantilever, resulting in marked degradation of the test mass Q.

The figure also shows the dovetails that need to be cut into the intermediate and test masses. Tests at UWA have shown that cutting of small dovetails produces only minor reduction in test mass acoustic Q provided the test mass is annealed after the cut [73, 74]. Our test mass dovetails were cut by the manufacturer before

final annealing so good performance was expected. The dovetail in the intermediate mass was not as critical and was done within the department. The photograph in figure 3.25 shows the flexure mount positioned between the intermediate and test masses. As can be seen the direction of the flex is in the direction of the optical axis in order to best utilise the additional vibration isolation provided by the flexure.

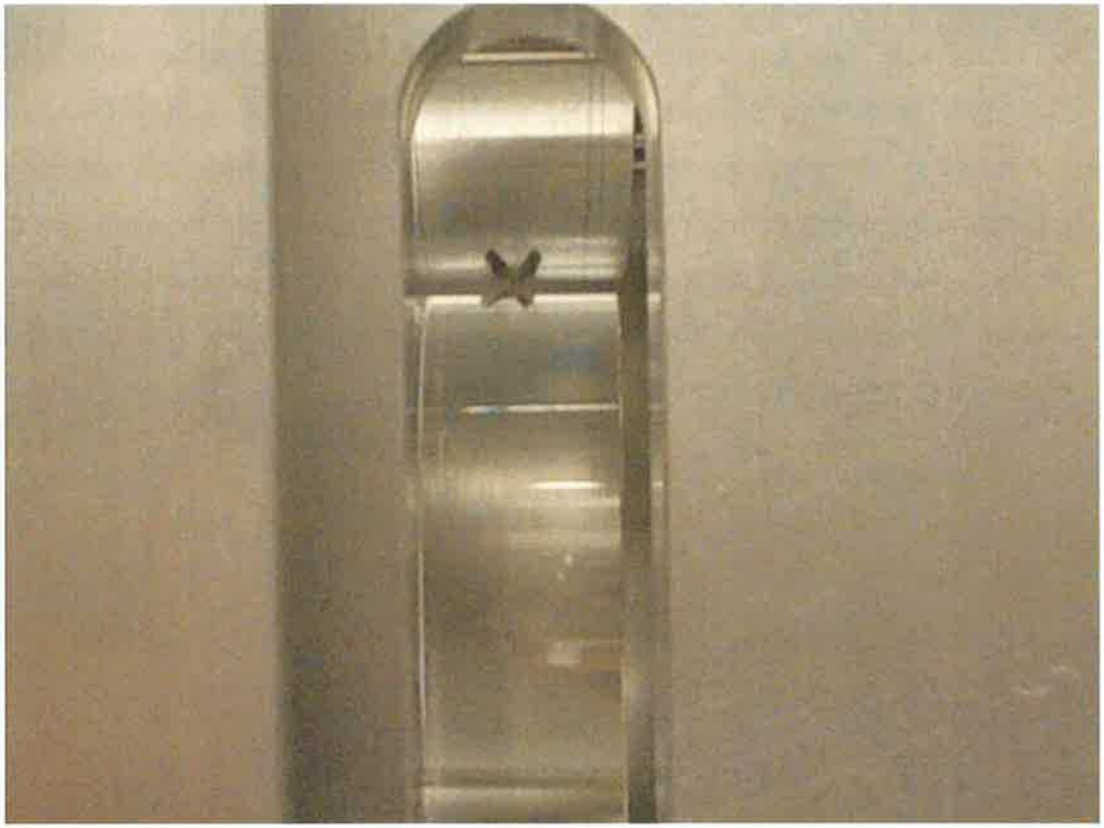

Figure 3.25: The small "X" shaped flexure mount positioned betwen the intermediate mass (top) and test mass (bottom). The view is through the access slot in the side of the reaction mass. Note that the flexure is slightly canted over in this photograph.

The flexure mount was not specifically characterised for this experiment, its behaviour being assumed to be typical of the others built at UWA. Typically these flexure mounts have allowed acoustic Q's of the order of 1×10^7.

3.4 Conclusion

In this chapter we have examined the construction of a seismic isolation system used to mount the test cavity. This suspension was built to fit into an existing vacuum chamber and made use of a large seismic pre-isolator. As the use of a monolithic test cavity renders the experiment relatively insensitive to seismic noise the isolation system was relatively simple, consisting of a wire suspended intermediate mass from which the test mass was suspended using a monolithic flexure mount built at the University of Western Australia. Eddy current damping was employed to damp movement in the suspension system, although the damping tests will need to be repeated when the system is placed into vacuum. While the suspension system was designed with regard to minimising the acoustic losses in the intermediate and test masses, the acoustic Q's of the test and intermediate masses will also need to be measured in the vacuum environment.

Chapter 4

Optoelectronic Design

4.1 Introduction

In this chapter we will describe the optical layout of the experiment and the associated electronics. The optical layout consists of two main subsystems: the laser is locked to a stable reference cavity using the Pound-Drever-Hall technique [77], and the reference cavity is locked to the test cavity at low frequencies using a fringe side locking technique.

An overview of the experiment's optical topology is shown in figure 4.1. As was discussed in section 1.4, the experiment uses two cavities: a long reference cavity is used to impart the required frequency stability onto the NPRO laser, which is then used to measure the length of the short test cavity. A Pound-Drever-Hall discriminator is used to produce an error signal that indicates the frequency difference between the laser and the reference resonance. This error is used to lock the laser to the reference cavity, thereby imparting the cavity's stability onto the laser frequency.

A fringe side detector (D4) is used to measure transmission of the probe light through the test cavity and hence to determine its length changes. In addition, D1 and D2 allow measurement of the laser intensity and its noise before each of the two cavities.

Figure 4.1 also shows that much of the optics is intended to be placed inside a vacuum chamber. Vacuum incompatible components will be kept outside the chamber as will some of the reference cavity mode matching optics due to space

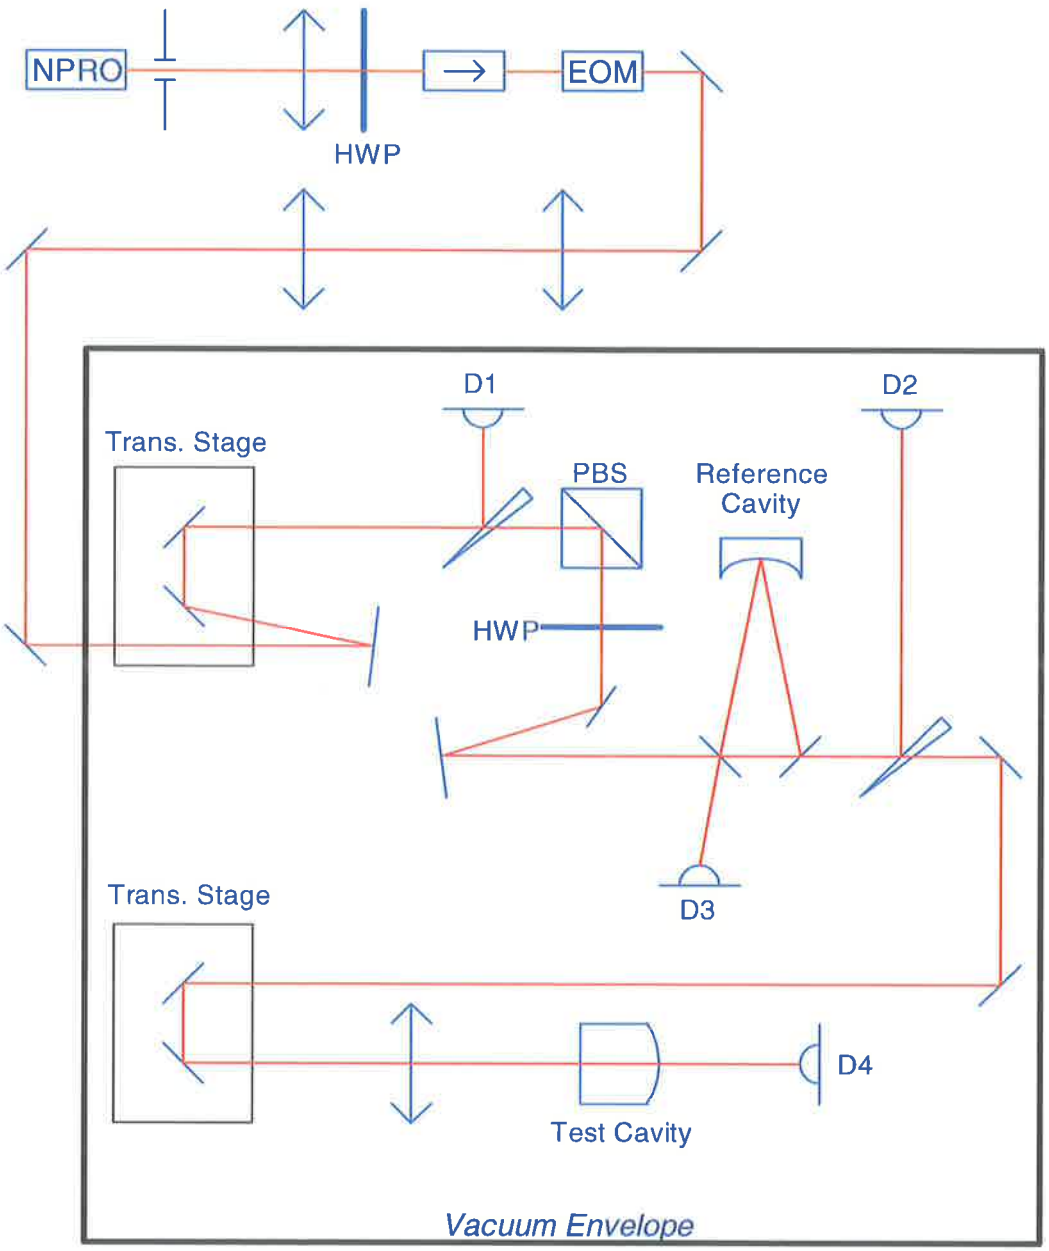

Figure 4.1: Layout of the optical components used in the experiment.

restrictions. Two steering mirrors allow alignment of the beam from the external optical bench to the optics inside the vacuum chamber. The vacuum chamber was not ready for use during the phase of the experiments described in this thesis, so all measurements were carried out on an external optical table. So that any problems due to space restrictions could be anticipated before the move into vacuum, the breadboard was laid out in the same way that the in-vacuo table will be in the next phase of the project. The photograph in figure 4.2 shows the entire experiment with the area to be placed inside the vacuum chamber highlighted.

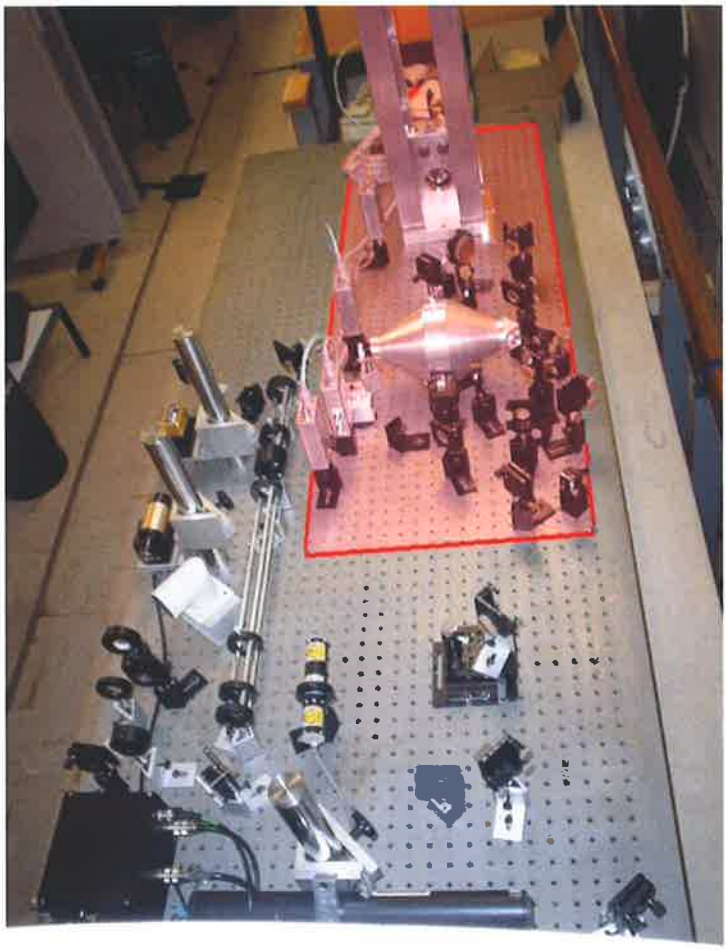

Figure 4.2: Photograph of the optical setup. The section to be moved into the vacuum chamber is highlighted.

Section 4.2 examines the system used to stabilise the laser to a reference cavity. Details of the reference cavity are provided in section 4.2.2, followed by measure-

ments of cavity properties. This section continues with a description of the Pound-Drever-Hall sensor [77] used to measure frequency differences between the laser and the reference cavity (section 4.2.3). The actuators used to provide feedback to the laser are then characterised in section 4.2.4, followed by a description of the compensator design (section 4.2.5). Finally, the results of the frequency stabilisation scheme are discussed in section 4.2.6.

Section 4.3 describes the scheme used to lock the reference cavity to the test cavity and measure length noise in the test cavity. Again, we begin with a description of the design and construction of the test cavity and then discuss the frequency difference sensor (section 4.3.2), the actuators (section 4.3.3) and the compensator (section 4.3.4).

In the course of this work we discover that free-running laser intensity noise could contaminate the test cavity readout. Section 4.4 thus describes the intensity stabilisation of the laser. Section 4.4.1 describes the sensor used to measure the intensity fluctuations. In the now familiar sequence, we will then discuss the actuator (section 4.4.2), compensator (4.4.3) and performance (section 4.4.4) of the intensity stabilisation. Finally, in section 4.4.5, we discuss the coupling from intensity actuation to frequency noise (and *vice versa*).

4.2 Laser Frequency Stabilisation

4.2.1 Introduction

As was discussed in chapter 2, a free running NPRO laser does not have sufficient frequency stability to be used in the experiment. To remedy this situation the laser frequency stability is improved by locking it to the resonant frequency of a reference cavity. In section 1.3.2 we saw that the reference cavity was required to be very long compared to the test cavity so that length noise in the reference cavity would not mask the length noise signal in the test cavity. Design of the reference cavity centered on the need to make it as stable as possible, so that its resonant frequency would serve as a good reference. This required a design that would be mechanically rigid while being long. The reference cavity was also required to be tunable, so tuning elements were included.

A Pound-Drever-Hall (PDH) sensor is used to measure the frequency difference between the laser and the cavity. This sensor provides a measure of the frequency difference between the laser and the cavity resonance while the two are reasonably close together. However, when the laser and cavity frequencies are more than approximately one cavity linewidth apart the PDH technique does not provide a useful error signal. If a transient frequency disturbance is sufficiently large to drive the system outside this range then the servo system will be unable to correct for the error. When this occurs the servo system is said to lose "lock", a situation that can only be remedied by external intervention to drive the two resonance frequencies back into coincidence.

Two actuators are used to control the laser frequency so that it follows that of the reference cavity. Fast laser frequency modulation is produced using a piezoelectric transducer (PZT) mounted on the laser crystal. Voltages applied to the PZT strain the laser crystal and thereby change the laser frequency. The PZT actuator has a relatively high bandwidth (about 100 kHz) but only a small dynamic range of about 200 MHz.

A slower, high dynamic actuator uses temperature tuning of the laser crystal to change its optical path length. The thermal actuator has a very large dynamic

range but is very slow, having a bandwidth of significantly less than one hertz.

4.2.2 Reference Cavity Design

A three-mirror, ring topology was chosen for the reference cavity. This ensures that optical feedback to the laser is minimised, as light reflected from the input coupler of the reference cavity is directed away from the laser, rather than retracing its input path. Optical feedback has been shown to affect laser frequency stability [78, 79] and although an optical isolator is used to protect the laser from this effect, the additional isolation provided by the ring topology is desirable.

As the reference cavity is required to have a stable resonance frequency, mechanical rigidity of the cavity was of great importance. This requirement led to the choice of a monolithic-like construction for the reference cavity. The main cavity spacer was machined from a single piece of aluminium and shaped into a double-tapered "football" shape. This shape maximises the mechanical rigidity of the spacer and keeps its thermal mass low. Aluminium was chosen for the spacer material because of its high thermal conductivity and low acoustic losses. The cavity's high conductivity allows the cavity length and hence resonance frequency to be tuned thermally if required.

As was discussed in section 1.3.2, we require the reference cavity to be long compared to the test cavity. The size of the reference cavity was, however, limited by the available space inside the vacuum chamber. This limited the length of our reference cavity to about half a meter.

The cavity length was further constrained by the need to choose a suitable Free Spectral Range (FSR) for the reference cavity. The feasible cavity size range corresponds to accessible FSRs from about 300 MHz to 750 MHz. We chose to have an FSR greater than 500 MHz as efficient resonant electro-optic phase modulators were available at those frequencies. The availability of such modulators was deemed useful as it allowed for the possible transmission of modulation sidebands through the reference cavity, should that prove necessary. The requirement for a resonant frequency greater than 500 MHz implies an optical path length less than 600 mm.

The mechanical design of the cavity is shown in figure 4.3. It has an overall length

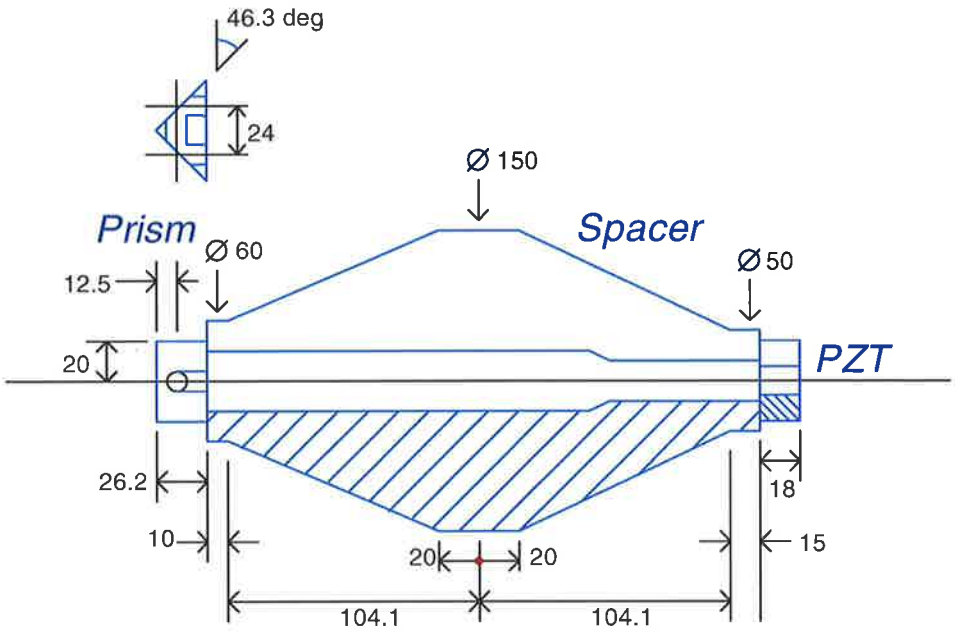

Figure 4.3: Mechanical design of the reference cavity.

of 258 *mm* and a diameter of 150 *mm*. It was firmly attached to a standard optical kinematic base so that it could be easily removed from the system for diagnostics as required. While this mounting was suitable for this phase of the project, the final version of the experiment will include a custom made, low-loss mechanical mount for the cavity to improve mechanical isolation of the spacer and to reduce acoustic losses.

A small breather hole was drilled through the spacer to allow evacuation of the cavity. This breather hole was plugged to prevent dust ingress during benchtop testing.

The reference cavity was fitted with three fused-silica supermirrors having reflectivities in excess of 99.95% [43]. Two planar mirrors were glued to the small aluminium prism using a vacuum compatible epoxy. This prism was then glued to the spacer. The curved mirror was mounted on a high dynamic range piezoelectric transducer (PZT) to allow the cavity resonance frequency to be tuned, and the assembly was then glued to the spacer. The high reflectivity, curved mirror has a one meter radius of curvature.

Figure 4.3 shows the location of the cavity piezoelectric transducer (PZT), used

to affect rapid variation of the cavity resonance frequency. In addition, a small ledge was machined into the spacer near the PZT mounting point. This ledge is intended to allow resistive wire to be wound onto the spacer thus enabling temperature tuning of the cavity if it is required; however, this would result in increased acoustic losses.

The cavity was assembled by gluing the various parts together with Torr-Seal, a commercial vacuum-compatible epoxy based glue produced by Varian incorporated. It has been reported that cavity mirrors glued with Torr-Seal can lead to degradation of cavity finesse [private communication, Malcolm Gray 2001], presumably due to fuming of the epoxy during curing. The finesse of our cavity was not monitored during the curing process so it is unknown whether such a degradation took place in this case. However, as will be discussed below, the resulting finesse of the cavity indicates that this degradation did not occur in our case.

Figure 4.4 shows a photograph of the reference cavity.

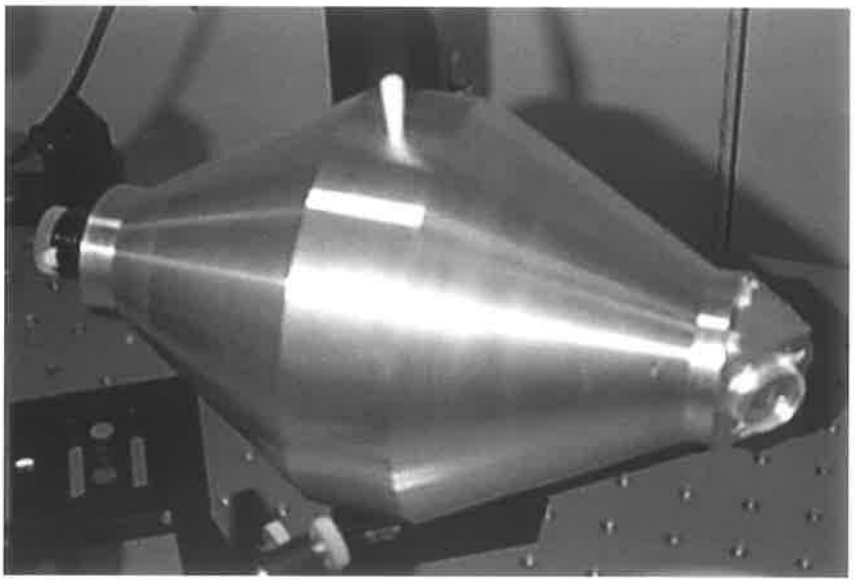

Figure 4.4: Photograph of the reference cavity. The input coupler is on the near side of the prism on the right of the spacer. The curved mirror may be seen mounted on its black PZT at the left end of the spacer. The white cylinder on top of the cavity is the plug blocking the breather hole.

4.2.2.1 Optical Properties of the Cavity

The calculated optical properties of the cavity are shown in figure 4.5. The cavity has an FSR of 561 MHz with a transverse mode spacing of 96 MHz. The spot radius on the curved mirror is 452 μm, while those on the flat entry and exit mirrors are 387 μm. The optical length of the cavity is 0.53 m.

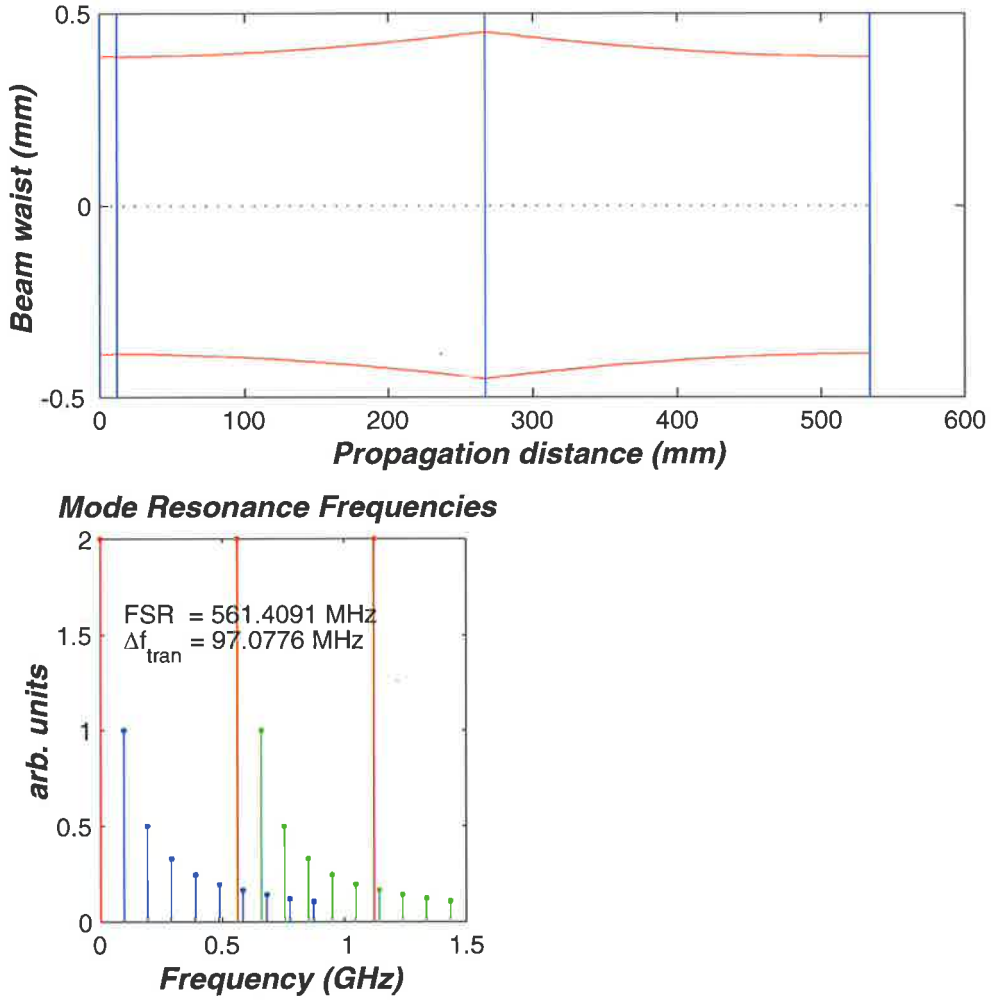

Figure 4.5: Calculated optical properties of the reference cavity. The top graph shows the beam size of the laser mode inside the cavity. The lower graph shows the frequencies of the cavity modes. The red lines indicate longitudional modes, each separated by one cavity FSR. The blue and green modes are transverse modes corresponding to different longitutional modes.

The finesse of the cavity describes the full-width half-maximum (FWHM) linewidth of the cavity resonance normalised to the free spectral range, where the FWHM

linewidth is the frequency range over which at least half the maximum possible amount of power is transmitted through the cavity:

$$F = \frac{FSR}{FWHM}. \tag{4.1}$$

Measurement of the cavity finesse was performed by applying a ramp voltage to the PZT element and monitoring the transmitted power. As the cavity moved into resonance with the incoming light, a cavity mode would build up and some of the light would be transmitted though the cavity. The transmitted light was detected using a photodetector, details of which can be found in appendix E. Figure 4.6 shows a typical scan from this measurement. The figure shows a sweep through a cavity resonance and two modulation sidebands spaced 17 MHz from the main resonance. These sidebands are discussed in section 4.2.3. The cavity was found to have a full width half maximum (FWHM) linewidth of 136±18 kHz, yielding a cavity finesse of 4100±500. This finesse corresponds to a mirror reflectivity, R, between 99.97% and 99.98%, a range in accord with the manufacturers specification of R>99.97%. We therefore conclude that the mirrors were not significantly effected by the curing epoxy.

4.2.2.2 Reference Cavity Mode Matching

Laser light incident on the reference cavity needed to be well mode matched to ensure efficient transfer of laser power into the cavity mode and thus maximum transmission through the locked cavity. Maximising the power going into the cavity TEM_{00} mode also improves the performance of the locking system and reduces the chance of locking the laser to higher-order cavity modes.

Mode matching requires that the incident laser beam properties be modified so that the laser mode overlaps the resonant cavity TEM_{00} mode. Overlap of the two modes requires three conditions be satisfied:

1. The laser beam and the cavity TEM_{00} mode must be colinear.

2. The laser beam waist must be in the same position as the waist of the cavity TEM_{00} mode.

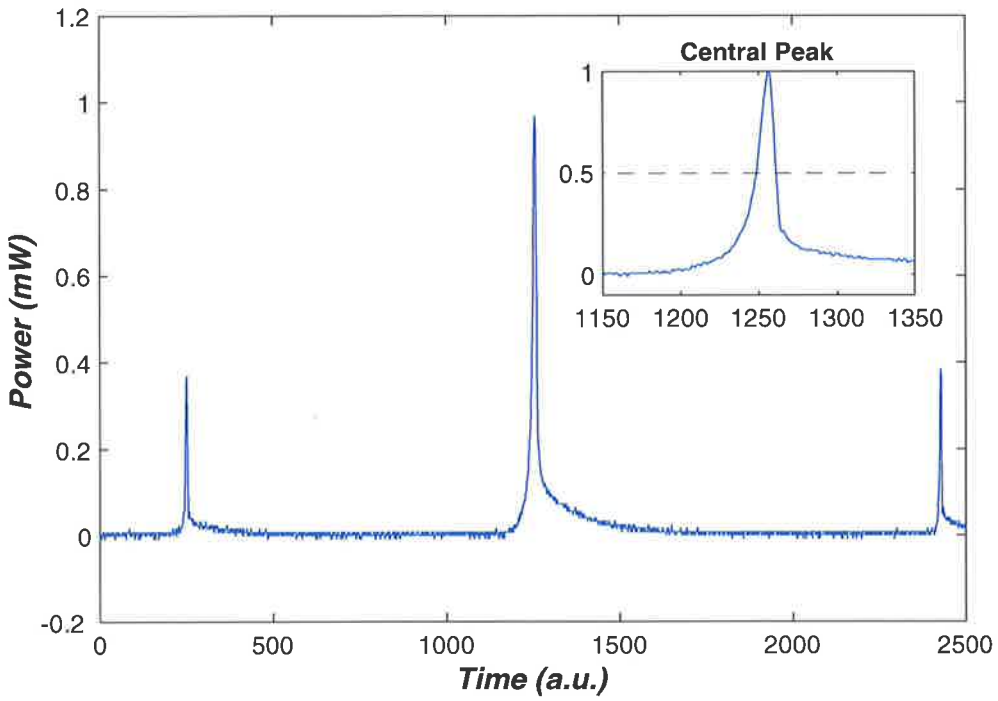

Figure 4.6: Plot of transmitted power as the resonance frequency of the cavity is swept through the laser frequency. The figure shows the main resonance with its two 17 MHz sidebands. The inset shows a close up of the central peak.

3. The laser beam waist must be the same size as the cavity TEM_{00} mode waist.

The properties of the incident laser beam were determined by measuring the beam radius with a CCD camera at several points and then fitting to the equation of a propagating gaussian beam. The laser beam was found to have a radius of 400 μm and a radius of curvature of 170 mm at the output window of the laser.

Using beam propagation code written in Matlab, a mode-matching system was designed to match the laser mode into the known cavity mode. The optical components of this system are shown in figure 4.7. The two one-meter focal length lenses, L1 and L2, are located 450 mm and 810 mm from the laser and produce a waist of the correct size. The beam is aligned into the cavity using the steering mirrors M3 and M4. Two additional mirrors, M1 and M2, are mounted on a translation stage, which allows the path length between the laser and the cavity to be adjusted to the required 2.86 m without significantly effecting the alignment. The beam predicted by the software is shown in figure 4.8.

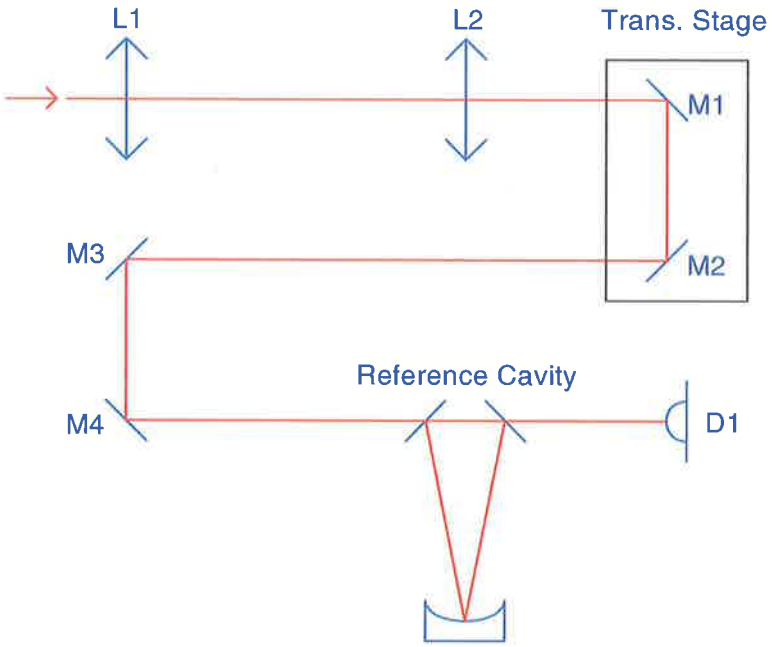

Figure 4.7: Optical components used to modematch the laser to the reference cavity. The two lenses, L1 and L2, each have a focal length of 1 m and are located 450 and 810 mm from the laser respectively. M1-4 are plane mirrors allow alignment of the beam into the reference cavity.

Optimisation of the mode matching was performed by using the reference cavity as an optical spectrum analyser. The length of the cavity was varied by applying a triangular voltage to the cavity PZT. The cavity resonances were thus swept across the laser frequency and different cavity transverse modes became resonant at different times [80]. The different resonances were detected using photodetector D1 and examination of the relative magnitude of the different modes allowed easy identification of any mismatch. Odd order transverse modes are caused by misalignment of the laser beam with the cavity TEM_{00} mode while even order modes arise from waist size and position errors [80]. The odd order modes were easily removed by adjustment of the steering mirrors M3 and M4, shown in figure 4.7. Removal of the even order modes was somewhat more difficult, relying on the simultaneous adjustment of the positions of the lenses L1 and L2 and the variable length stage M1/M2.

The mechanical stability of the mode matching optics was found to be excellent;

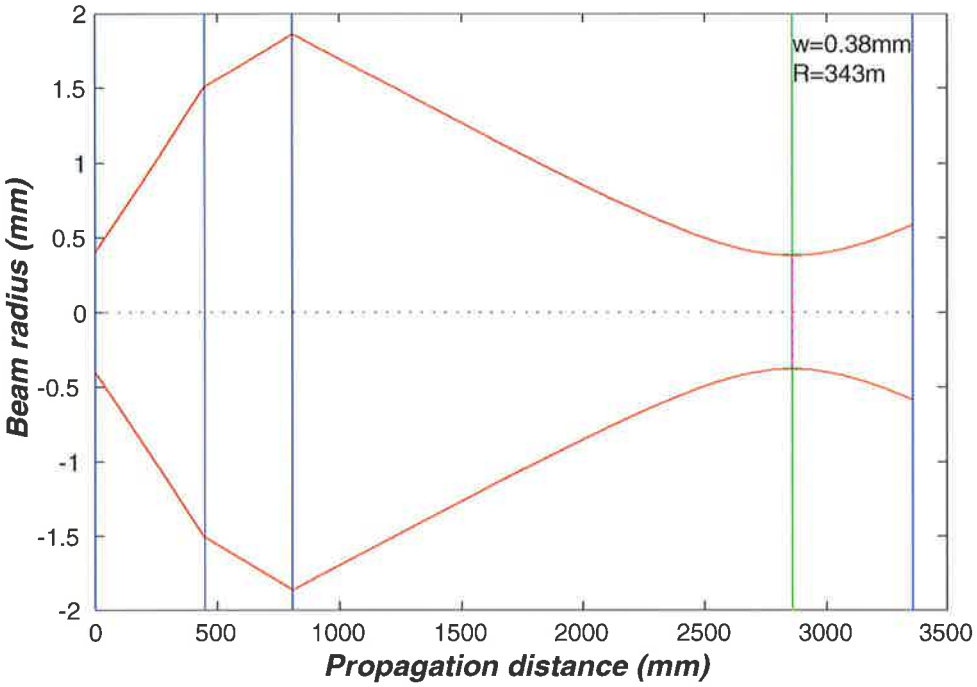

Figure 4.8: Predicted beam propagation through the reference cavity modematching optics. The two modematching lenses are located at the two vertical blue lines, 450mm and 810mm from the laser. The resultant waist occurs at the magenta line and is coincident with the cavity waist, the position of which is indicated by the green line.

realignment was not generally required for periods of several months.

4.2.3 Frequency Sensor

4.2.3.1 The Pound-Drever-Hall Technique

Reflection locking techniques provide a number of advantages over transmissive methods. Transmissive techniques are sensitive to laser intensity noise because there is no way to determine whether a transmitted power variation is due to genuine frequency noise or contamination by intensity noise. In contrast, in a reflection locking scheme there is ideally no detected power when the laser and cavity are locked. Reflection locking is thus a form of nulled lock-in detection, a class of techniques that are known to show good performance in the presence of intensity noise [81].

Simply looking at the reflected light from a cavity is not sufficient to extract a

useful error signal. Such a procedure would simply produce a null when the laser and cavity frequencies match; when they were not matched there would be no information about whether the laser frequency was lower or higher than the cavity resonance. To extract this information, we frequency modulate the input light. Examination of the phase of the reflected light at the modulation frequency then allows determination of the sign of any frequency mismatch. The modulation frequency must be higher than the cavity bandwidth so that the modulation sidebands are not transmitted into the cavity. Optical locking using this technique was first demonstrated in [77] and was an extension of the microwave technique of Pound. This resulted in the technique being named Pound-Drever-Hall locking, often abbreviated to PDH locking.

Since the Pound-Drever-Hall technique was first reported it has become the standard method for locking optical cavities. The success of the method has been, in part, due to its insensitivity to laser intensity noise. A second advantage of the PDH frequency discriminator is that it has a better shot noise limit than fringe side locking (see [79], [82] and appendix D).

Examples of the use of PDH detection may be seen in [83], [78], [84], [85] and [86] among many others. The technique has been extensively analysed in the literature, both in the quasi-static case [85, 87, 88] and more recently in dynamic situations [89,90]. Use of the technique has grown beyond passive cavity locking to applications such as injection locking [91, 92].

4.2.3.2 Operation of the PDH technique

Figure 4.9 shows the components of a Pound-Drever-Hall detector. An electro-optic modulator (EOM) is used to impart frequency modulations sidebands onto the laser beam. The modulated light is then incident on the laser cavity, where the sidebands are totally reflected from the input coupler. The reflected sidebands and the cavity leakage field are detected with the photodiode. Finally, the photodiode signal is mixed with a sample of the EOM drive signal and low pass filtered (LPF) to yield the PDH error signal.

If the input carrier is sufficiently close to the cavity resonance frequency then

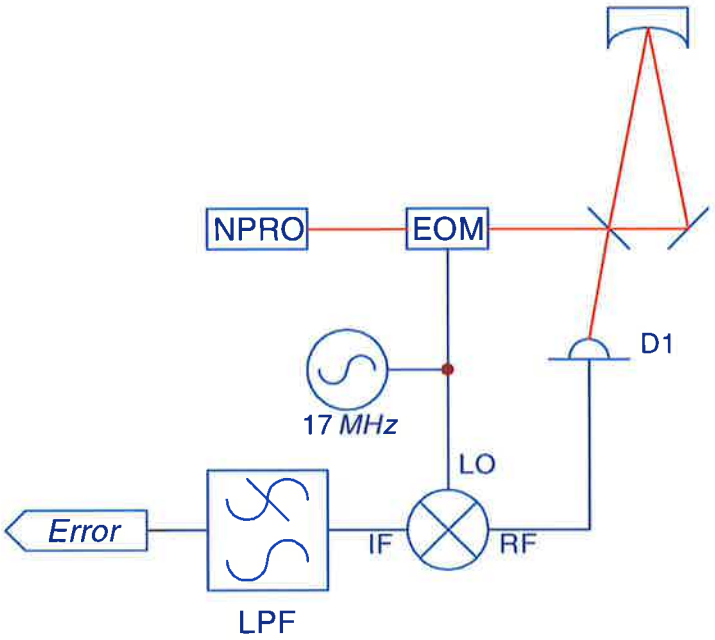

Figure 4.9: Major components of the Pound-Drever-Hall frequency discriminator.

some of the light will be coupled into the cavity, leading to the development of a cavity field. Some of the cavity field then leaks back through the input coupler and if the input field were perfectly in resonance with the cavity then this leakage field would be 180° out of phase with the input. However, if there is a small frequency difference between the incoming light and the cavity resonance then the phase shift picked up by the cavity field will be close to, but not quite equal to 180°. When the cavity leakage field interacts with the reflected input carrier they interfere to produce a signal having a phase shift that is strongly dependant on the frequency difference between the fields.

Let us consider the quasi-static case, where the laser frequency varies sufficiently slowly for the cavity field to remain in equilibrium with the input field. This case corresponds to frequency variation on timescales slower than the cavity storage time. In this case the signal component that results from interference of the reflected input and the cavity leakage fields will have the same frequency as the input carrier, but a frequency dependant phase shift.

Each of the two reflected sidebands also interact with the carrier. These modulation sidebands are placed well outside the cavity linewidth so that they are totally

reflected from the cavity when the carrier is at (or near) resonance. If the laser is on resonance then the two resulting signals will have equal amplitude but opposite phase. Cancellation will occur and the output will be zero. However, if the laser is not on resonance then the two signals produced by the sidebands do not cancel. This results in an amplitude modulated signal at the sideband frequency that can be demodulated to baseband with an RF mixer. This signal then results in the PDH error signal.

It should be noted that the PDH discriminator is non-linear; its output is proportional to the frequency difference between the cavity and laser only when the deviations between the two are small. The PDH discriminator is not unusual in this regard, as all sensors that measure frequency deviation from a cavity resonance are affected by the inherently narrow band nature of the resonance. A consequence of the non-linearity is that simple classical control theory cannot be used to design an adequate control loop to cover all frequency differences. The additional required control system complexity will be discussed in section 4.2.5.

As is discussed in appendix D, if we consider the region close to resonance then the sensor has a response of

$$H_{PDH} = H_o\sqrt{P_c P_s}FL \tag{4.2}$$

where P_c is the power in the carrier, P_s is that in the sidebands and H_o is a constant that describes the electronic gain of the detector system used to implement the PDH system.

As is also shown in the appendix, the power in the carrier and the sidebands of the modulated laser beam are given by

$$\begin{aligned} P_c &= J_o(\beta)^2 P_{inc} \\ \text{and } P_s &= J_1(\beta)^2 P_{inc}, \end{aligned} \tag{4.3}$$

where, $\begin{cases} J_o \text{ and } J_1 \text{ are Bessel functions of the first kind with order 0 and 1.} \\ \beta \text{ is the moduluation depth} \\ P_{inc} \text{ is the incident laser power} \end{cases}$

The response of the PDH error signal is maximised by judicious choice of the modulation depth β,with a value of 1.08 giving the best response [88]. Detailed derivations of the PDH error signal may be found in appendix D and also in [88] and [85].

4.2.3.3 The Shot Noise Limit for PDH

As the front end of a PDH system operates at high frequency, it is largely immune to the low frequency excess technical noise that occurs in all electronic circuits. The modulation frequency can also be made sufficiently high that there is no excess intensity noise in the laser. The limiting sensitivity of an ideal Pound-Drever-Hall system will thus be due to shot noise in the detected beam. The shot noise limit is derived in appendix D and results in a sensitivity limit of

$$S_{sn} = \frac{1}{FL}\sqrt{\frac{hc^3}{8\lambda P_c}}. \tag{4.4}$$

Figure 4.10 shows that a shot noise limited PDH sensor should satisfy the sensitivity requirement of the experiment when illuminated with a beam having a carrier power of more than 1 mW.

4.2.3.4 Sensor Design

An overview of the circuit used to implement the Pound-Drever-Hall sensor is shown in figure 4.11. A $17MHz$ signal is first produced by a voltage controlled oscillator (VCO). This VCO was tuned to match the resonant frequency of the EOM by monitoring the size of the error signal. A high power amplifier boosts the VCO output to the power required to drive the modulator. A portion of this amplified power is then sampled using a directional coupler and used as the local oscillator

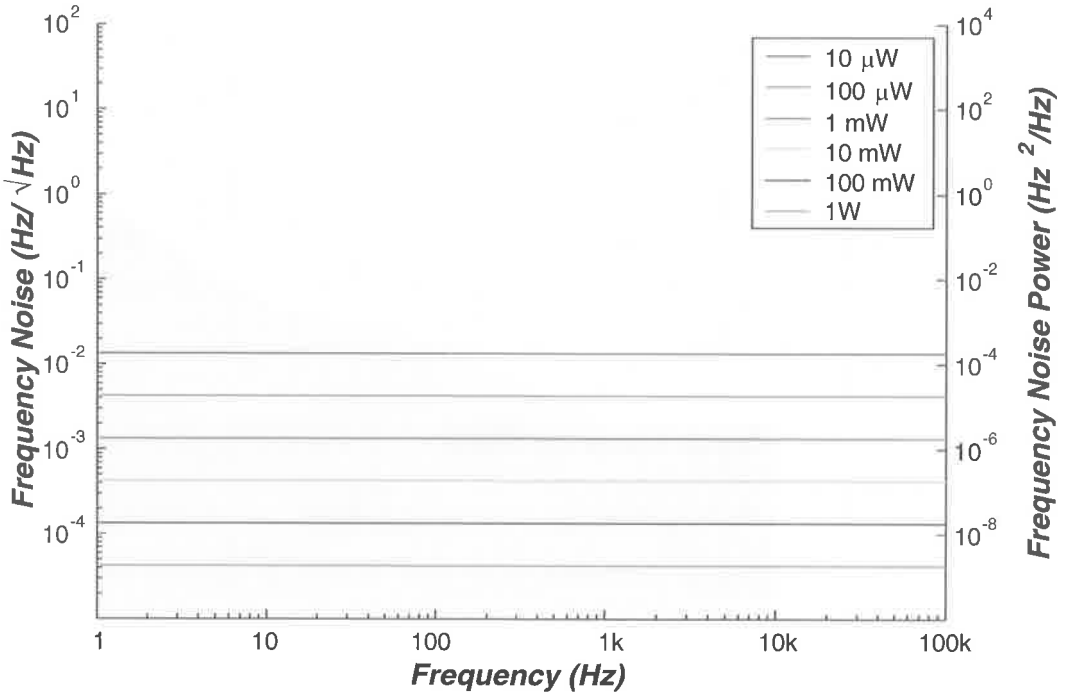

Figure 4.10: Dependence of shot noise limit with detected laser power using a Pound-Drever-Hall discriminator.

(LO). The remainder of the power is used to drive a model 4003 resonant electro-optic modulator made by New Focus, which was factory tuned to 17.0 MHz.

The signal reflected from the reference cavity is measured using a photodetector based on a high power InGaAs PIN diode. Details of the photodetector design can be found in appendix E. The photodetector is placed approximately 10 m from the PDH circuitry, so its output signal is amplified to minimise interference. The photodetector signal is passed to a phase delay unit (not shown in the schematic) that consists of a variety of lengths of coaxial cables that can be switched in and out of the signal path. The delay unit enables the phase of the detected signal to be matched to that of the local oscillator so that the error signal's shape could be optimised. The delayed signal is then low pass filtered with a filter having a bandwidth of 21.4 MHz to remove the second harmonic of the EOM drive (at 34 MHz). This harmonic arises when the sensor is operating near cavity resonance, but as it is not useful in forming an error signal it can be removed to prevent overloading the mixer. As the photodetector signal is large to prevent interference, it must now

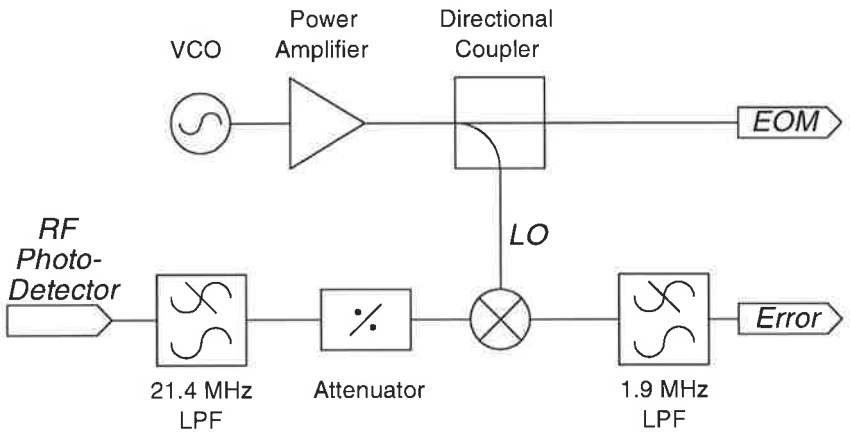

Figure 4.11: Pound-Drever-Hall detector schematic.

attenuated to prevent overloading the mixer. The mixer then demodulates the signal to baseband. A low pass filter having a bandwidth of 1.9 MHz is used to remove the high frequency mixer products at 17 MHz and 34 MHz and produce the final error signal. Further details on the design and construction of the PDH circuit can be found in appendix E.

As optimisation of the modulation depth is important to achieving good sensor performance, the modulation depth was measured directly. To do this the laser frequency was swept through cavity resonance for different EOM drive powers, which were produced by adding a variety of fixed 50 Ω attenuators to the EOM drive signal. For each value of attenuation, the transmitted power was monitored during the sweep and the amplitudes of the various peaks were recorded.

Values of β were determined for each level of EOM drive attenuation using equations 4.3 and 4.3. These data are plotted in figure 4.12. Because the power difference between the points in the graph is known, we can determine the EOM responsivity by fitting a curve to the experimental data. The results of this procedure are shown in the figure and results in an estimate of the EOM responsivity of 6 rad/W. As the EOM is a 50 Ω device this corresponds to 0.12 rad/V, a result which is at the lower end of the manufacturer's quoted range of 0.1-0.3 rad/V [93].

As can be seen in the figure the points corresponding to EOM drive attenuation of 3 dB and 6 dB straddle the "optimal" β value of 1.08 and hence either would be

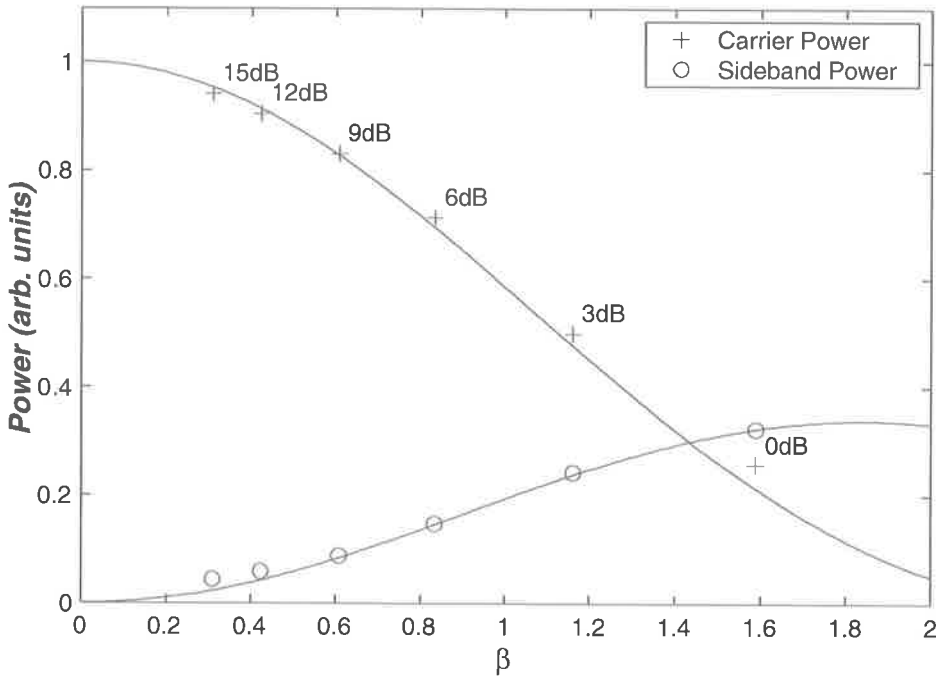

Figure 4.12: Variation of carrier and sideband powers with modulation depth. The red points are measured data, the blue curves are fitted curves yielding an estimate of 0.12 rad/V for the EOM sensitivity. The figure beside each carrier power gives the corresponding RF attenuation.

an adequate operating point. The 6 dB attenuator leaves more power at the laser carrier frequency and consequently allows more power through the reference cavity when locked. This has the advantage of providing more laser power for measurement of the test cavity. Thus, the 6 dB attenuator setting was used for the remainder of the experiment.

Figure 4.13 shows the PDH error signal obtained when the reference cavity was swept through resonance. The characteristic PDH error signal curves are seen as the cavity came into resonance with the carrier and its sidebands. The bottom curve in the figure shows the transmitted power during the sweep. As expected the graph shows that the PDH error signal has a null whenever the cavity and laser beam are in resonance. The figure also shows that the PDH nulls corresponding to the 17 MHz sidebands have opposite slope to that of the carrier wave; there is thus no danger of the system locking to a modulation sideband once the polarity of the compensator has been correctly set.

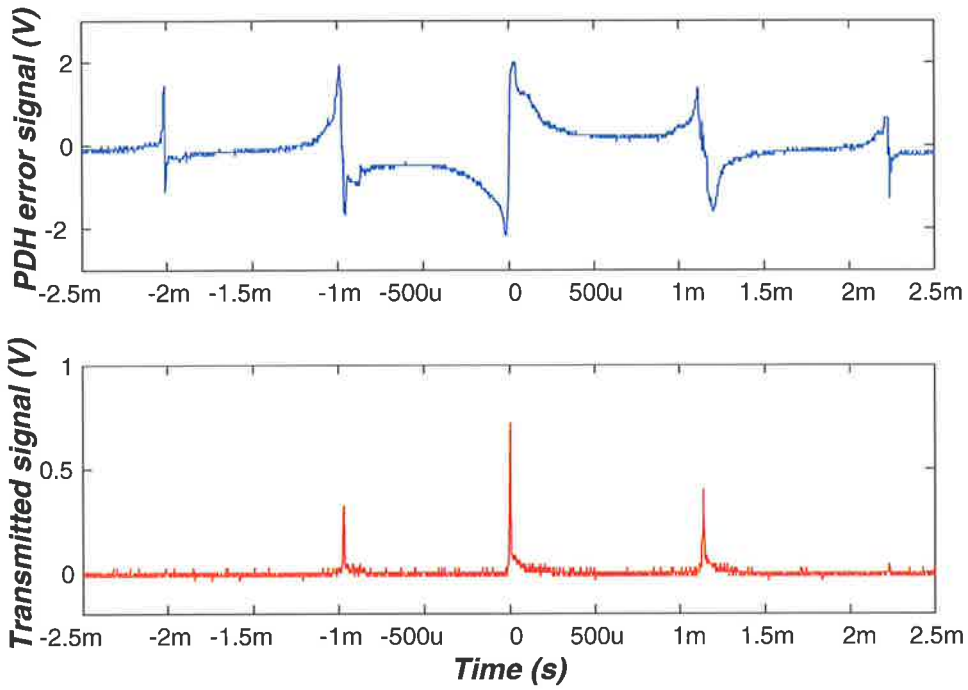

Figure 4.13: PDH error signal and transmitted signal as the reference cavity was swept through resonance with the incoming laser beam.

The procedure outlined above was repeated for a range of sweep rates and in each case the error signal curves (such as that shown in figure 4.13) revealed that the linear region was approximately 70 kHz wide and that in this region the slope of the PDH error signal was approximately 50 $\mu V/Hz$. We can therefore write the nominal sensor gain as

$$H_o = 50\ \mu V/Hz. \tag{4.5}$$

Although the signal shown in figure 4.13 corresponded to the laser exciting the TEM_{00} mode of the cavity, similar curves also occur when exciting other cavity transverse modes. The presence of sustainable higher order transverse modes admits the danger of locking to one of them rather than the desired TEM_{00} mode. Should inappropriate locking occur, it can be detected by examination of the cavity transmitted power or the error signal. Locking to higher order modes is best

avoided, however, by correctly mode matching to the cavity. This ensures that it is very inefficient for laser light to couple into the higher order cavity modes.

4.2.4 Actuators

The frequency of the NPRO laser used in this project could be varied using either a piezoelectric transducer mounted on the gain medium or alternatively by changing its temperature. The PZT was useful for fast modulation and the heater for slow, high dynamic range variations of the laser frequency. In this section the two actuators are discussed and measurements of their transfer functions are presented.

4.2.4.1 Piezoelectric Actuator

The laser piezoelectric actuator was used for high bandwidth tuning of the laser frequency over a small dynamic range. The PZT is affixed to the gain medium so that voltages applied to the PZT cause strain in the gain medium, which in turn lead to a shift in the laser frequency through strain-induced birefringence [94]. According to the laser's manufacturer, the PZT actuator has a bandwidth of 100 kHz but a low dynamic range of order 200 MHz [95].

Accurate measurement of the PZT actuator's frequency response was difficult as the crystal was not accessible for an interferometric measurement such as the one discussed in section 4.3.3.1. A less direct method was therefore used to estimate the magnitude of the PZT response. In this experiment a sinusoidal drive voltage was applied to the PZT and the resulting frequency modulation was analysed using the reference cavity. A schematic of this experiment is shown in figure 4.14.

We expect that the PZT response should be linear, so sinusoidal modulation of the PZT voltage should lead to a sinusoidally varying laser frequency. Addition of a DC component to the drive allows this varying signal to be centered on the reference cavity resonance. We then used the reference cavity to analyse the laser frequency changes; detection of the laser carrier and the 17 MHz sidebands transmitted through the cavity allows us to determine the voltage drive required to produce a 17 MHz frequency shift in the laser output, yielding the PZT gain at the modulation frequency. The frequency of the sinusoidal drive was then changed and

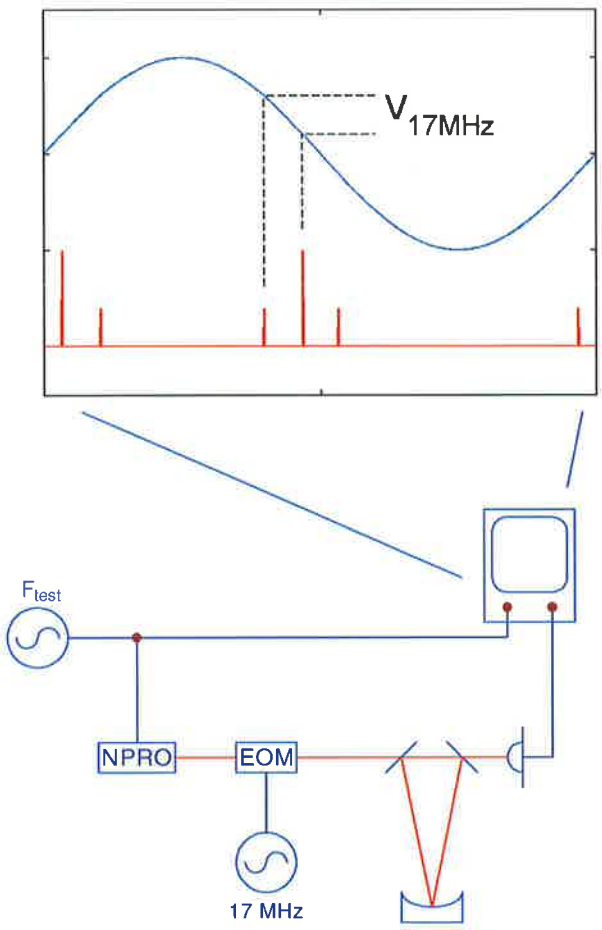

Figure 4.14: Experimental setup for determination of the laser PZT magnitude response.

the procedure repeated to construct the PZT magnitude response.

It should be noted that the output impedance of the function generator, F_{test} and the PZT capacitance form a low pass filter which will attenuate the drive signal at high frequencies. If this filter has too low a bandwidth it can affect the response measured *via* the above procedure. To find the capacitance of the PZT, it was driven through a 68 $k\Omega$ series resistor and the corner frequency of the resulting low pass filter was measured. The corner frequency (which was independent of dc offset) was 600 Hz, corresponding to a PZT capacitance of 3.9 nF. Combining this measured capacitance with the function generator's output impedance of 50 Ω results in a low pass filter with bandwidth in excess of 800 kHz. We therefore need not worry about contamination of the PZT magnitude response below 500 kHz.

Unfortunately, the procedure outlined above yields no direct information about the phase response of the actuator. While we could attempt to use Bode's gain-phase relationship to estimate the phase response, such a procedure would be fraught with danger as it assumes the absence of zeros in the right half of the s-plane (it assumes a phase-minimum plant). Lightly damped electro-mechanical systems often contain unstable zeros with corresponding high frequency resonances [96]. Consequently any estimated phase response would be of limited value at high frequencies.

Figure 4.15 shows the measured magnitude response. As can be seen the frequency response is flat to approximately 90 kHz, above which there is a complex series of resonances. There is also a very deep antiresonance at approximately 98 kHz, which was deeper than could be measured with the technique described above because a large enough PZT drive voltage could not be applied to drive the laser output through the required 17 MHz.

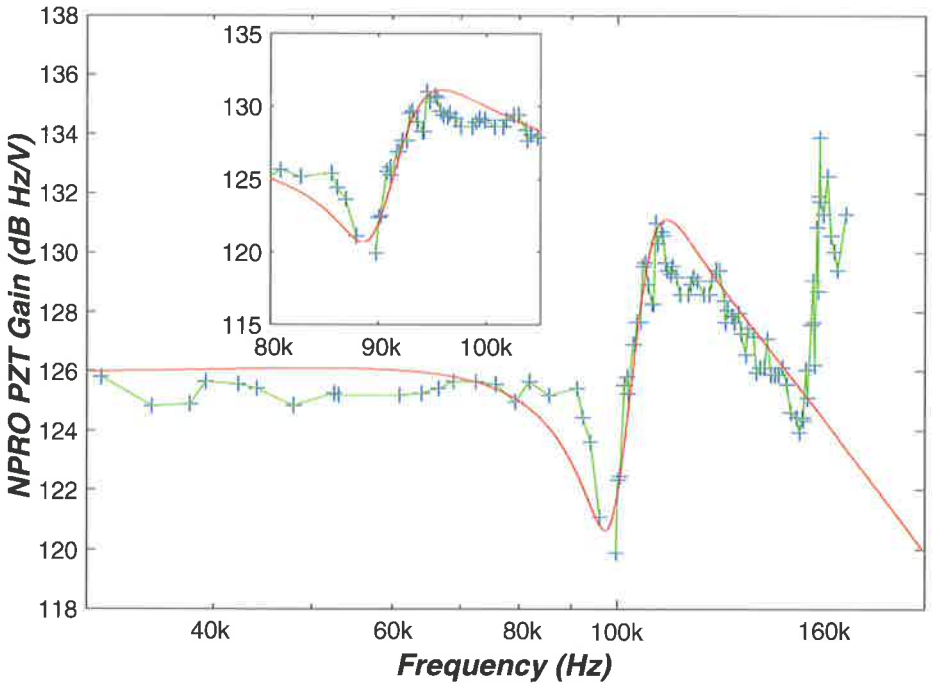

Figure 4.15: Magnitude response of the laser PZT actuator. Data points are shown as blue crosses. The red curve indicates the nominal response as described in the text. The inset shows the same data but in a small linear region around the first resonance.

The low frequency gain acquired by the above technique was confirmed with a

measurement of the slope of the PDH error signal at 1 kHz (a frequency well below the PZT resonance). The ratio of PDH output to the PZT drive was examined and the known PDH response of 50 $\mu V/\sqrt{Hz}$ used to calculate the PZT gain. At 1 kHz the PZT gain was found to be 1.96 MHz/V, a result which is in good agreement with the manufacturers quoted figure of 1-2 MHz/V [95]. The magnitude curve shown in figure 4.15 includes this calibration of the response.

A mathematical model of the actuator response was constructed for use in compensator design and its magnitude response is shown in red on the figure. The model consists of a 2^{nd} order anti-resonance and two 2^{nd} order resonances. The characteristics of the various resonances in the nominal model are shown in table 4.1. The second resonance was included to correctly roll off the plant phase to -180° at high frequency.

	Frequency (kHz)	Damping ratio (ζ)
Anti-resonance	99	0.05
Resonance 1	107	0.07
Resonance 2	120	0.5

Table 4.1: Resonance properties of the nominal NPRO PZT model

These resonance parameters yield a nominal PZT actuator transfer function of

$$G_{fo}(s) = \frac{1.3 \times 10^{18} \times (s^2 + 6.2 \times 10^4 s + 3.9 \times 10^{11})}{(s^2 + 7.6 \times 10^5 s + 5.7 \times 10^{11})\,(s^2 + 9.4 \times 10^4 s + 4.5 \times 10^{11})} Hz/V. \tag{4.6}$$

where the f subscript on G stands for "fast" and distinguishes it from the slow transfer function that will be discussed in the next section.

While the nominal curve is not of particularly high fidelity it does capture the salient features in the real response, namely the position and height of the first PZT resonance. It is this feature that will limit the bandwidth of the final control loop. The first anti-resonance was included in the model as it was otherwise impossible to model the sharpness of the first resonance. Even though the PZT resonance in

the vicinity of 160 kHz was very large it was not included in the model. The high frequency roll-off in the compensator required to ensure stability at the first resonance will ensure that this second resonance peak will remain well below unity gain. If a more accurate model of the system were required then experimental determination of the PZT phase response would be needed. It was felt sufficient to design the loop using this simplified model and then optimise loop performance with empirical tuning.

4.2.4.2 Thermal Actuator

A second method for laser frequency modulation was available *via* tuning of the laser crystal temperature. Thermal tuning provided a wide dynamic range but was very slow. The manufacturer of the laser specifies a dynamic range greater than 10 GHz with a tuning coefficient of 3.1 GHz/V, but with a bandwidth of less than one hertz [95]. Since the change in temperature is produced by changing the set-point of the laser's in-built temperature controller, the dynamics of the laser response are partially determined by tuning of the laser's internal PID controller.

The heater transfer function was measured while the laser was locked to the reference cavity. This locking was accomplished with the PZT, as will be described in the next section. For clarity we will discuss the experiment here and assume the existence of the appropriate compensator. Measurement of the heater transfer function was performed by applying a small modulation signal to the heater input and monitoring the corresponding change in the PZT modulation. Figure 4.16 shows the experimental layout. If the laser output is to stay at a constant frequency then the frequency variation produced by heating the laser gain medium must be cancelled by the frequency variation produced by the PZT: $X_s G_s = -X_f G_f$ where G_s is the laser transfer function from the heater input to output frequency and G_f is the transfer function between the PZT input and the laser frequency.

A measurement of the transfer function from the heater modulation input to the laser frequency is shown in 4.17. As the heater controller is very slow, the response was measured only over a small frequency range, 300 mHz to 4 Hz. Extrapolation of the gain curve to lower frequencies suggests that this is consistent with the man-

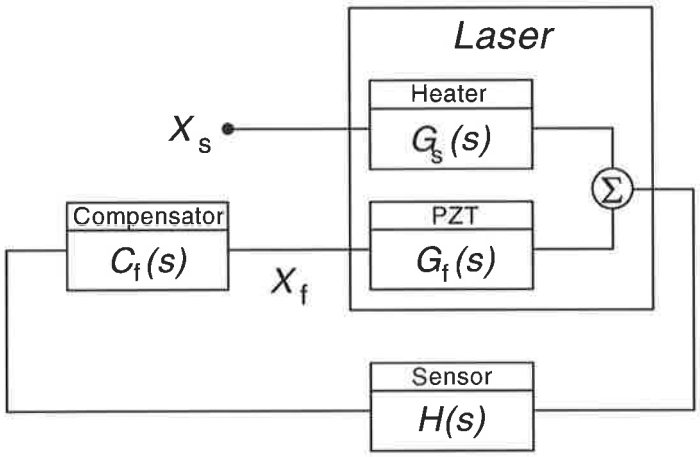

Figure 4.16: Block diagram of the experiment for the determination of the NPRO heater transfer function.

ufacturer's value of 3.1 GHz/V (190 $dBHz/V$) at DC. As can be seen in the figure, the bandwidth of this actuator is extremely low, with the phase dropping to -180° at 700 mHz.

The slow bandwidth of the heater was confirmed by measuring the transient temperature response when a step change was applied to the laser temperature set point. Figure 4.18 shows the response. The top curve shows the setpoint change, with the gain from command voltage to temperature being quoted as 1 K/V by the laser's manufacturer [95]. The bottom curve shows the error in the crystal temperature as a function of time.

4.2.5 Compensator Design

4.2.5.1 Introduction

We have seen that a free running NPRO laser is too noisy to meet our sensitivity goal (see section 2.3.3). In appendix B we show that the application of appropriate feedback can stabilise a system and reduce noise. The feedback reduces the laser frequency noise, $Y(s)$, by a factor $S(s)$: $Y_{closed\ loop}(s) = S(s) \times Y_{open\ loop}(s)$. The factor $S(s)$ is known as the sensitivity of the system [96, 97]. Since we know the typical open loop noise of a NPRO laser and the required frequency noise for the experiment, we can determine the necessary sensitivity as shown in figure 4.19.

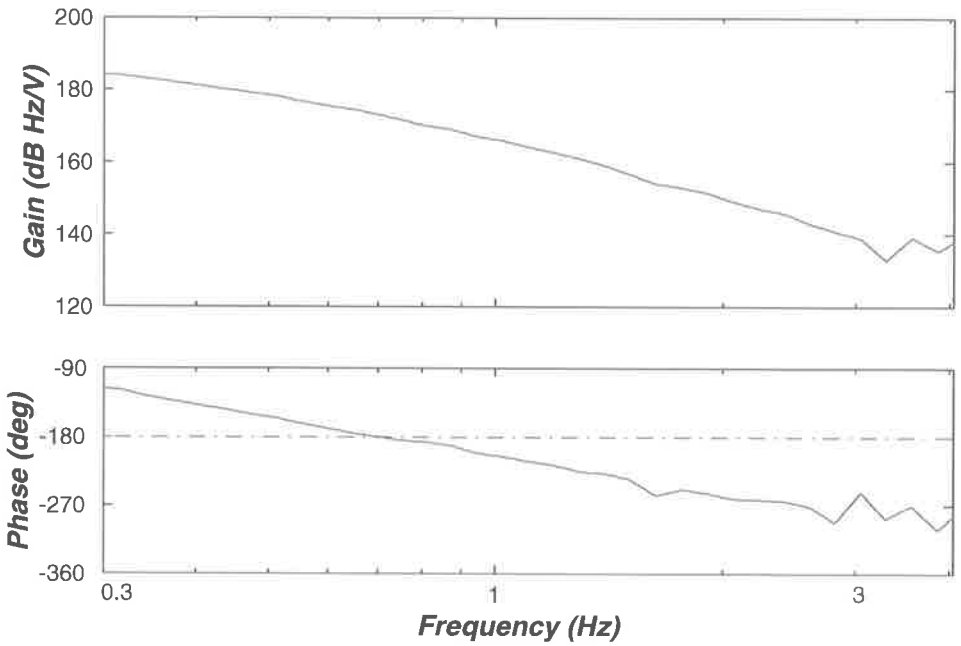

Figure 4.17: Transfer function from the laser heater modulation input to the laser frequency.

In a general servo loop, such as that shown in figure 4.20, the sensitivity function is determined by:

$$S(s) = \frac{1}{1 + C(s)G(s)H(s)}$$

$$\text{where,} \begin{cases} C(s) \text{ is the compensator transfer function.} \\ G(s) \text{ is the plant (laser and actuator) transfer function} \\ H(s) \text{ is the sensor transfer function} \end{cases}$$

The sensitivity requirement determined above thus leads to a determination of the required compensator gain $C(s)$:

$$C(s) = \frac{1 - S(s)}{G(s)H(s)S(s)}.$$

4.2.5.2 PZT Feedback Path

For simplicity in the design procedure we work with the nominal sensor and actuator transfer functions given in equations 4.5 and 4.6 respectively. Making the appropriate substitutions we obtain the target compensator gain shown in figure 4.21. While

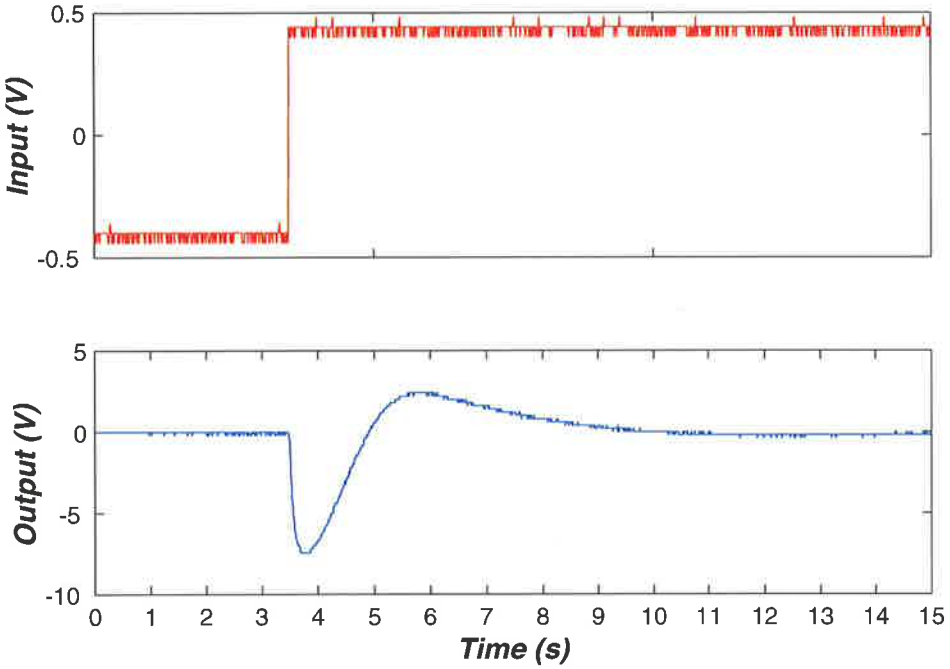

Figure 4.18: Transient response of the laser temperature in reponse to a step change in the heater setpoint.

the compensator gain remains in the green region of the curve the loop will provide frequency stabilisation sufficient to limit frequency noise to less than a tenth of the expected thermoelastic noise signal. Entry into the yellow region corresponds to gains that will keep frequency noise below the expected thermoelastic noise level.

Stability concerns prevent us from achieving the desired compensator gain at all frequencies however. We know that the first resonance of the PZT occurs at approximately 100 kHz so we must limit the unity gain frequency of the control loop to be somewhat below that. Since $G(s)H(s)$ is approximately 100 in the region of 100 kHz, we must ensure that the compensator gain has dropped to less than 0.01 by this frequency, as is shown on figure 4.21.

It is difficult to satisfy the stability requirement while meeting the design goal at the high end of the frequency band of interest. A design meeting such a goal would need a very rapid gain roll-off (faster than $1/f^3$) from 10 kHz to achieve a unity gain frequency less than 100 kHz. Such a control loop would prove difficult to tune and would likely be sensitive to environmental disturbances. We were therefore

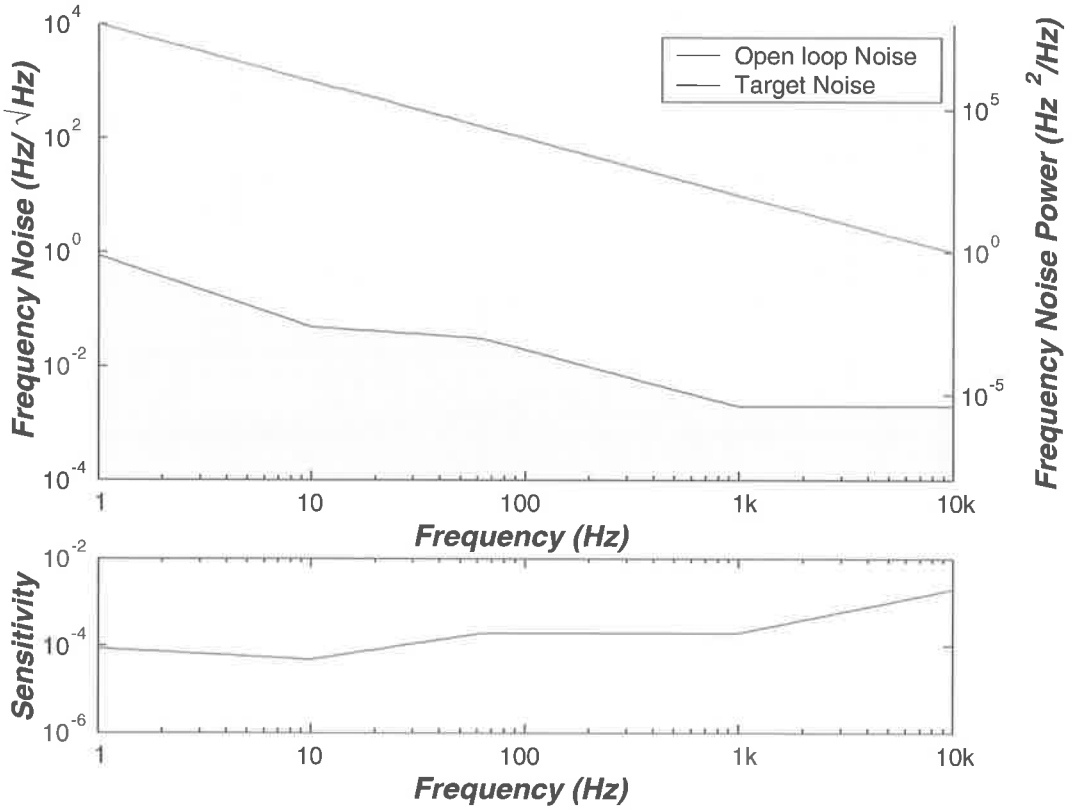

Figure 4.19: Determination of the sensitivity function required in the laser frequency stabilisation loop. The top axisgraph shows the open loop laser noise and the target noise level. The lower graph shows the required loop sensitivity.

forced to relax our design goal to one satisfying the noise suppression target over as much of the frequency range as was consistent with a reliable control loop.

The response of the final compensator design is shown in figure 4.22. As can be seen the compensator gain remains in the region of acceptable performance at all time. The gain curve levels out to a slope of -20 $dB/decade$ well below the PZT resonance so that the phase recovers before the unity gain frequency.

As described in section 4.2.3, the PDH error signal is linear only in the vicinity of resonance. The error signal is typically small (zero if there are no electronic offsets) away from resonance, but it also shows features due to higher order transverse cavity modes and due to the PDH modulation sidebands resonating with the cavity. When the system is operating outside the region having a linear error signal the large low frequency gain could cause the compensator to saturate at one of the supply rails.

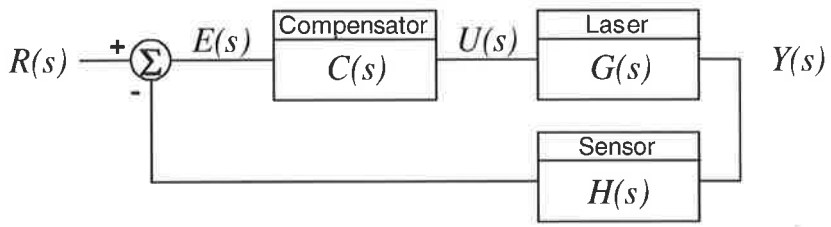

Figure 4.20: A general control loop. $R(s)$ is the reference signal, $E(s)$ the error signal, $U(s)$ the control signal and $Y(s)$ the system output. $C(s)$, $G(s)$ and $H(s)$ are the trasnfer functions of the compensator, plant and sensor respectively.

This is undesirable because such a saturated compensator would be ill-prepared to respond when resonance is approached. In practice such a system would be very poor at "locking" the laser frequency onto the cavity resonance.

To ameliorate this effect we lower the compensator's low frequency gain when the system is unlocked. We can therefore consider the required compensator as a simple gain-scheduled system, where the compensator exhibits two different gain curves for the two modes of operation. A low gain setting is used while the system is in "acquire" mode before switching over to a high gain "run" mode once lock has been achieved.

Figure 4.23 shows the nominal open loop compensator transfer functions for the two different modes. As can be seen the two curves differ only in the low frequency region. Because the two curves are identical in the region around the unity gain frequency the two modes have identical stability and high frequency performance. The change from acquire to run mode must be accomplished smoothly (so called bumpless transfer), so that there is no sudden jump in compensator output which would otherwise lead to a jump in the laser frequency and possibly could lead to loss of lock.

The compensator circuit is shown in figure 4.24. To change the compensator from acquire to run mode the switch, Sm, is opened to allow the low frequency boost stage to begin integrating. While the switch is closed, the low frequency boost stage provides the more modest gain allowable in the acquire mode. This compensator topology is similar to the limited integrator designs of [98] and [77], but is simpler

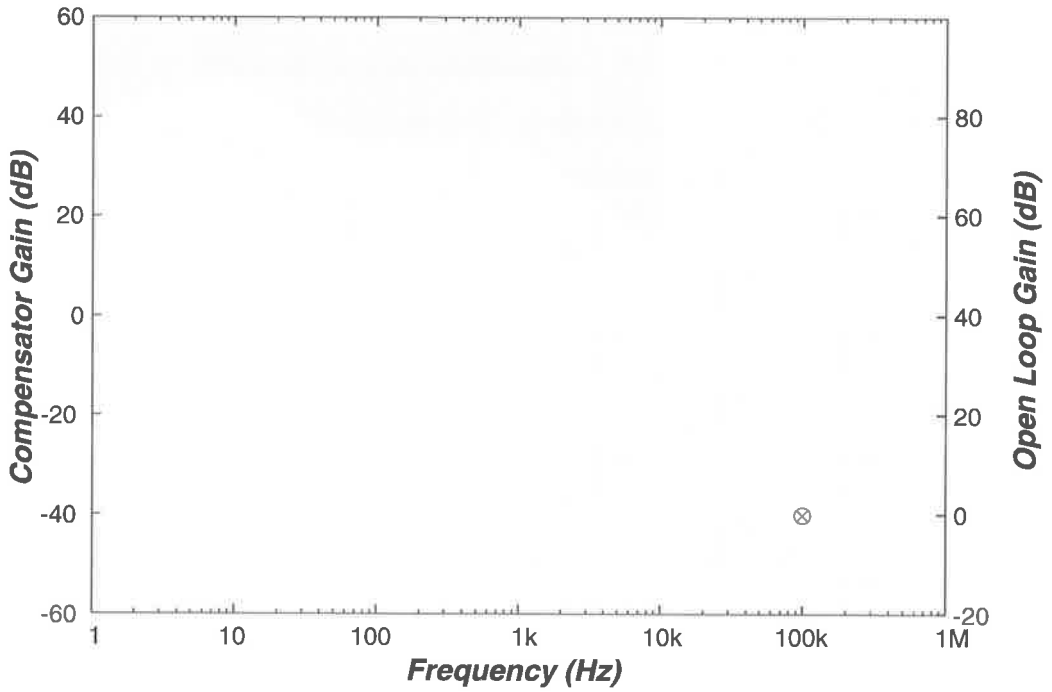

Figure 4.21: Requirements for the magnitude of the compensator gain. The green region on the figure indicates the desired area for the gain curve, the yellow area indicates permissible gain regions. The red cross indicates the maximum allowable gain to avoid instability in the control loop.

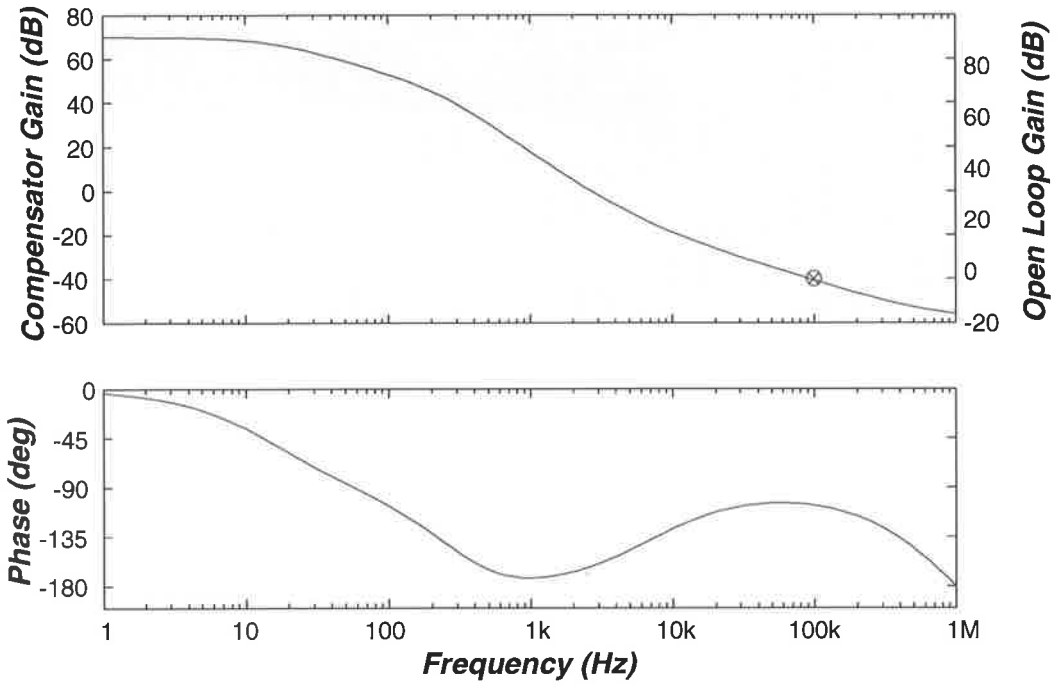

Figure 4.22: Bode plot of the fast frequency stabilisation compensator.

to optimise.

The compensator was connected to the system and the loop closed. The stability of the compensator was then optimised by adjusting the gain of the preamplifier stage of the compensator. The setting of this gain was a trade-off between noise suppression and reliability; a high gain setting improved the noise performance of the closed loop system but forced the loop into a less stable state by reducing its phase margin. As this loop was required to remain locked for long periods a conservative setting of the preamplifier gain was used.

Figure 4.25 shows the measured compensator transfer functions after the closed loop performance had been optimised. It is interesting to note that the final compensator gain is slightly higher than that of the nominal compensator, which suggests that the model for the laser PZT was pessimistic. The improvement in system gain is shown in figure 4.26. The unity gain bandwidth of the servo was found to be slightly less than 90 kHz.

The expected electronic noise at the output of the compensator was modelled using SPICE by following the procedure outlined in [99]. The noise was then multi-

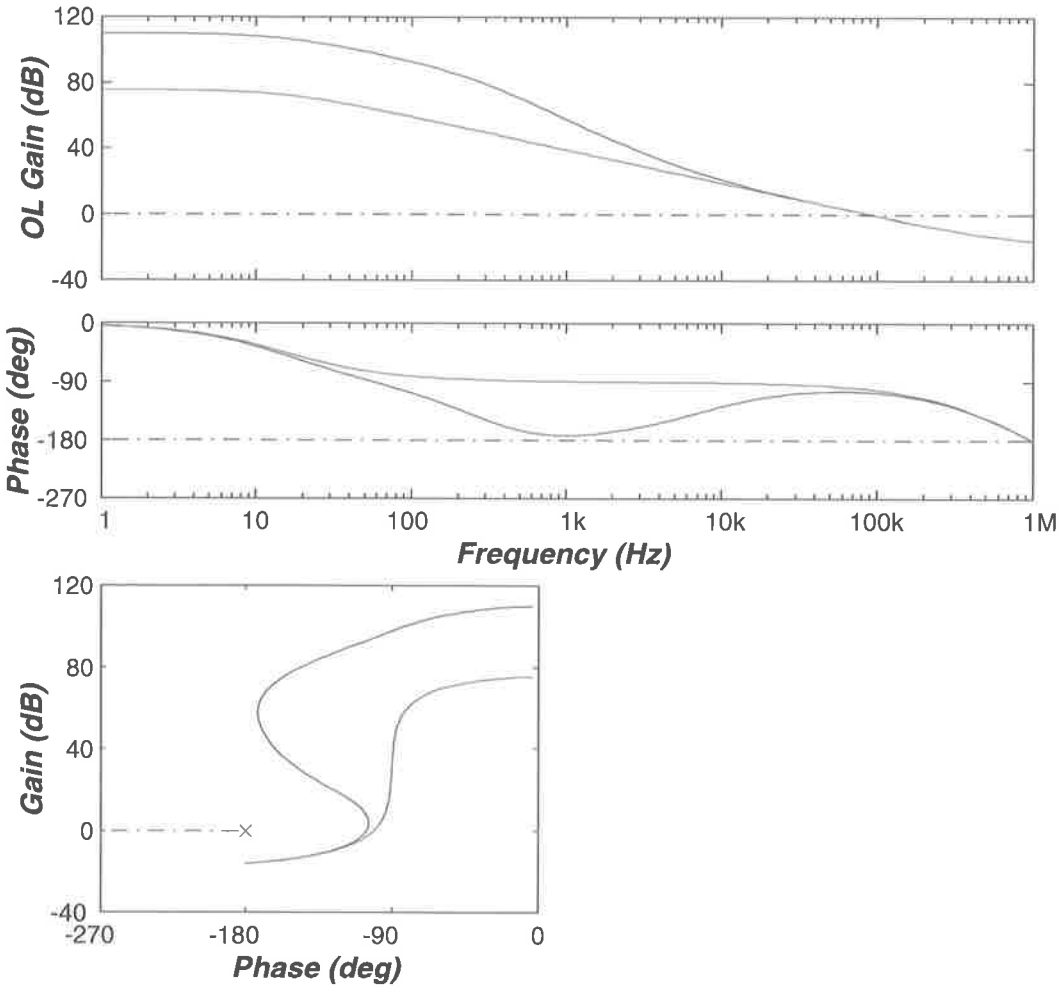

Figure 4.23: Bode plots of the two compensator modes. The blue curve shows the run mode, the green shows the acquire mode. The bottom plot is a Nichols plot of the open loop system.

plied by the loop's input sensitivity function (see appendix B) to find the effective frequency noise caused by the compensator electronic noise. The spectrum of this noise is shown in figure 4.27. As can be seen, the electronic noise does not compromise system performance.

4.2.5.3 Thermal Feedback Path

Locking of the laser using the PZT actuator alone was not sufficient due to the PZT's limited dynamic range, which resulted in a system that coped poorly with slow drifts. Consequently, a second feedback loop was added to extend the overall

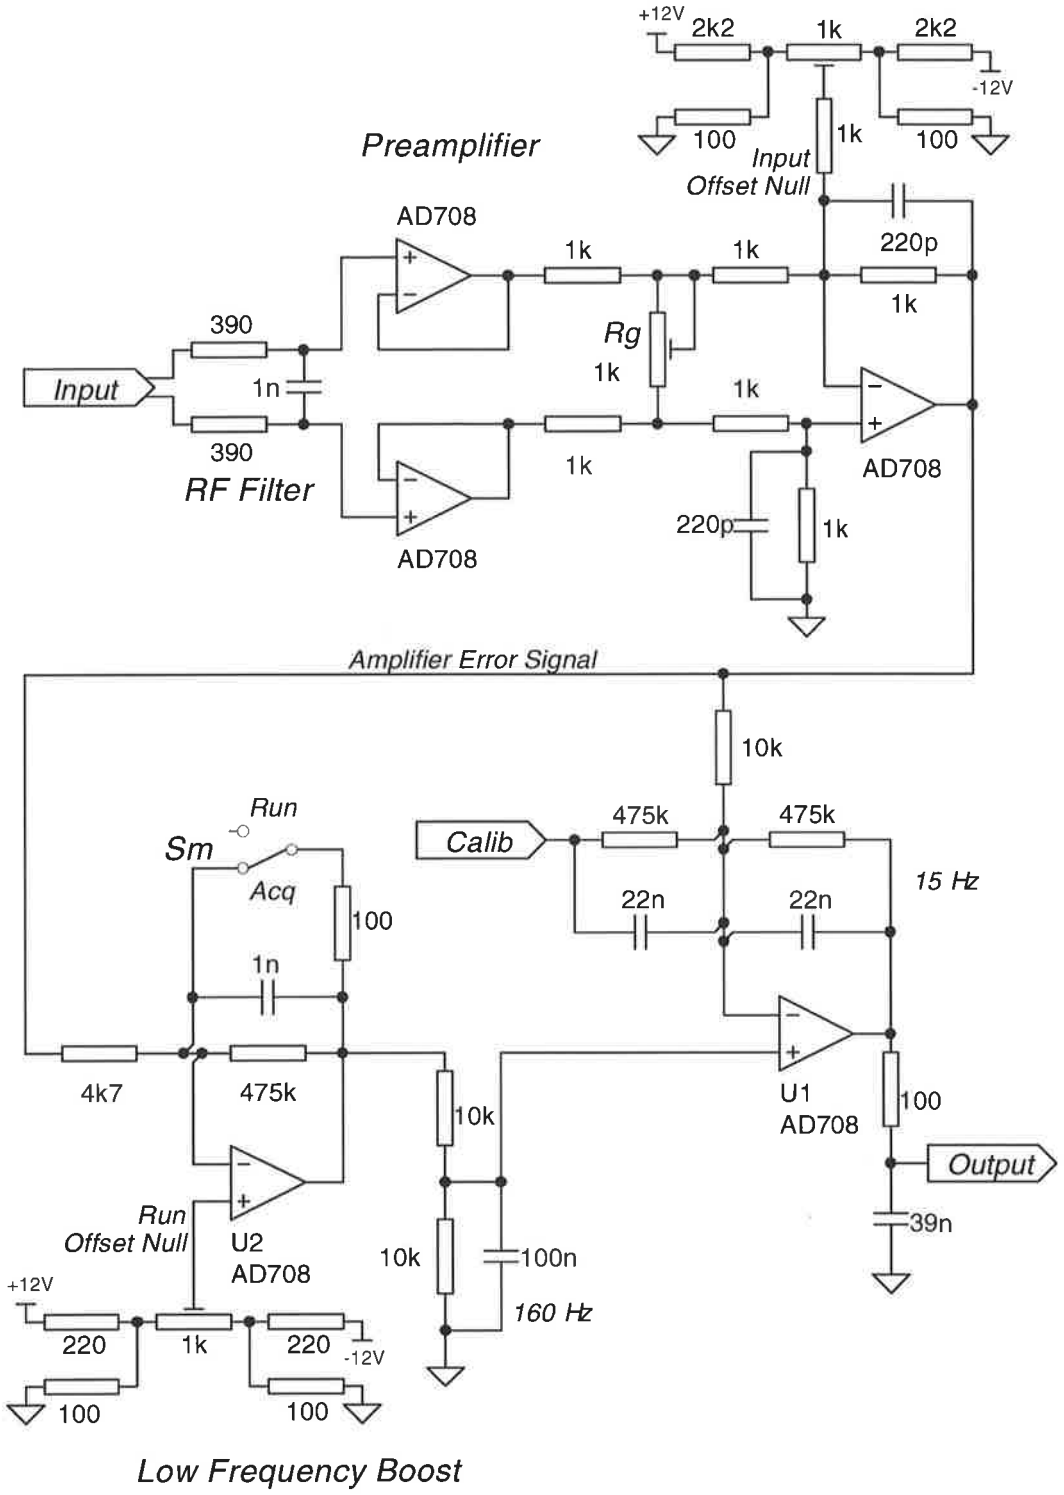

Figure 4.24: Schematic of the fast compensator used to drive the PZT acutator.

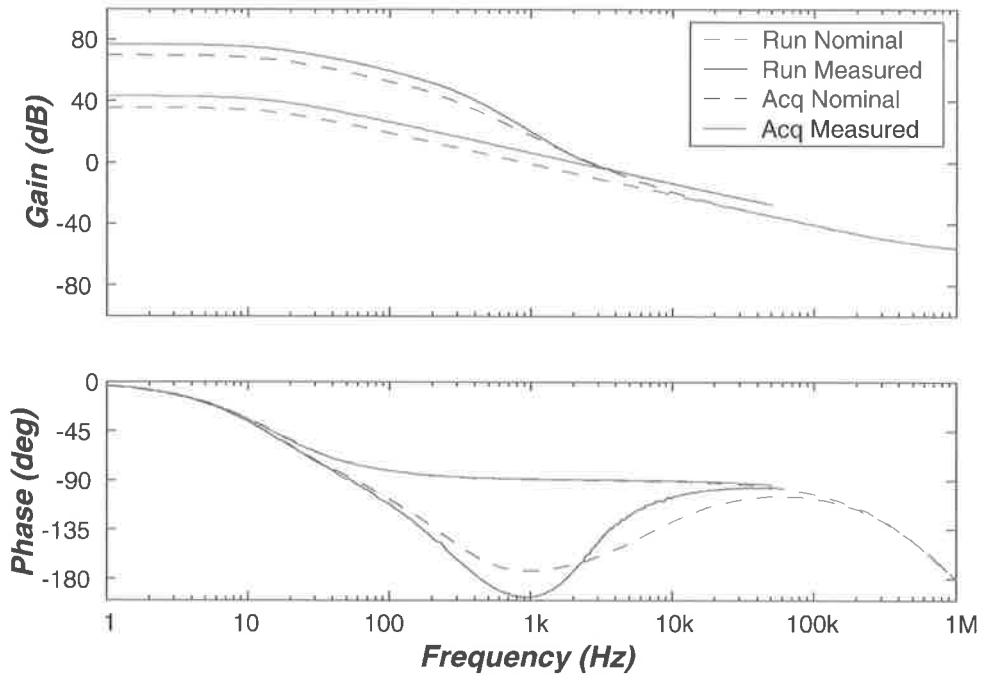

Figure 4.25: Measured response of the final compensator (solid lines) as compared to the nominal transfer function (dashed lines).

loop dynamic range. This second loop acted on the laser temperature to ensure that the PZT actuator stayed in the centre of its operational range. A block diagram of the entire system is shown in figure 4.28. The fast compensator C_f acts to keep the laser locked to the cavity, and the slow compensator C_s then acts to keep X_f at zero. This system is equivalent to that shown in figure 4.29 where the two feedback paths have been separated. This alternative representation makes clear that the system is a conventional split feedback system.

The slow compensator was required to cooperate smoothly with the fast compensator. The task of making the two compensators work well together was simplified during the design of the PZT compensator by ensuring that the fast compensator had constant gain at low frequencies (see figure 4.22). A slow compensator design which rolls off at $1/f$ in the region of crossover should therefore be stable.

Recall that in section 4.2.4.2 the transfer function from the heater input to laser frequency, X_f/X_s, was measured. This transfer function is shown in figure 4.30.

As can be seen from the figure the phase of the heater actuator transfer function

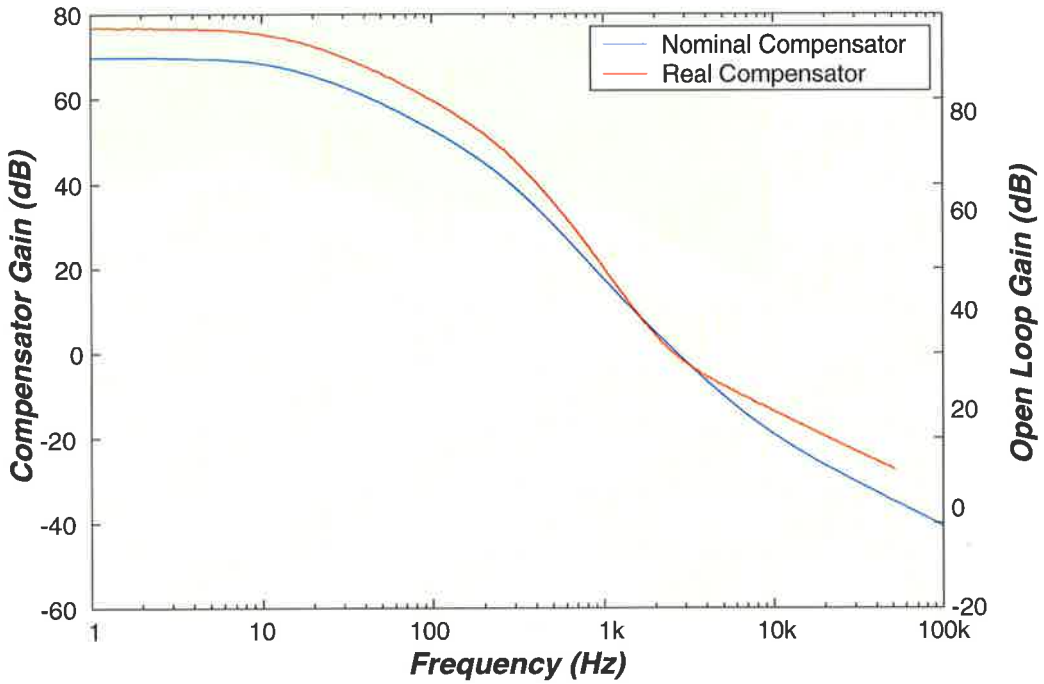

Figure 4.26: Comparison of the nominal compensator gain and that of the compensator as implemented. Both gain curves are for the compensator in the "run" mode.

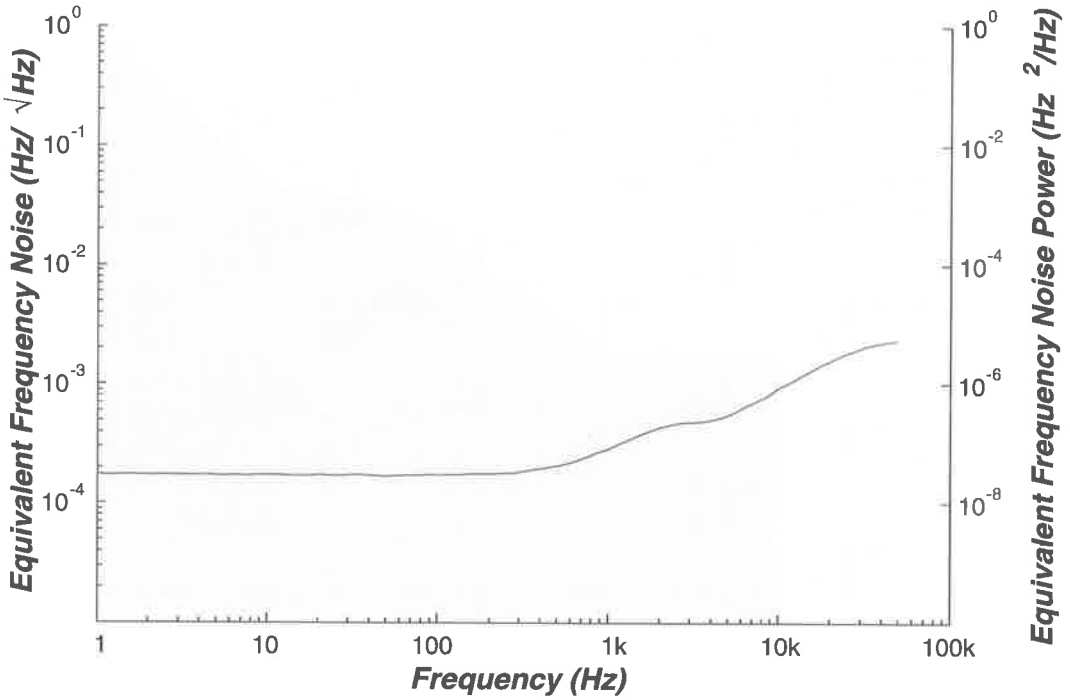

Figure 4.27: Effective frequency noise caused by electronic noise in the PZT feedback compensator. The background colours indicate the required sensitivity regions.

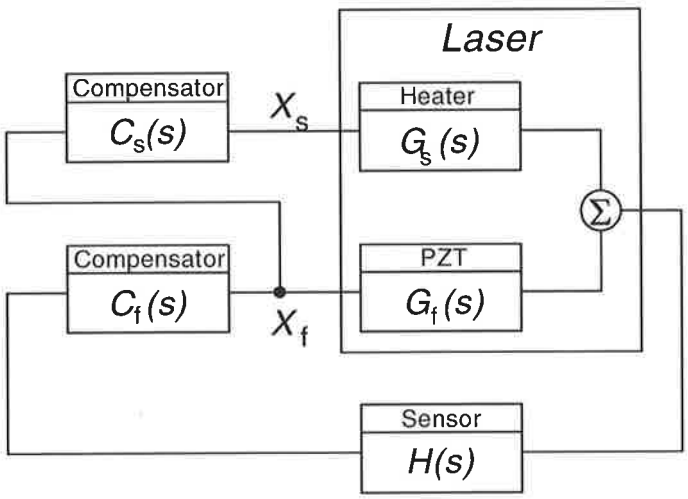

Figure 4.28: Block diagram of the complete frequency stabilisation system including the slow compensator, C_s, used to drive the laser heater.

crosses through -180 ° at 700 mHz. At this frequency the gain of the system is +47 dB. The slow compensator design must therefore have a gain of less than -47 dB at 700 mHz to produce a stable closed loop system. We would also like the PZT output to be driven to zero in the presence of no external disturbance so we use an integrator in the slow compensator. As was the case with the fast compensator, there is considerable danger of integrator saturation while the system is out of lock. This problem is particularly bad for the slow compensator because the slow time constants involved would lead to very long recovery times. We therefore again design the compensator with two operating modes: a "run" mode where the compensator looks like an integrator, C_{sr},and an "acquire" mode where the low frequency gain is reduced to prevent saturation, C_{sa}.

The blue curve in figure 4.31 shows the transfer function of a the compensator that meets these objectives. This simple transfer function is

$$C_{sa} = \frac{1.5 \times 10^{-4}}{s + 2\pi \times 72 \times 10^{-3}}$$
$$C_{sa} = \frac{1.5 \times 10^{-4}}{s + 0.45}$$

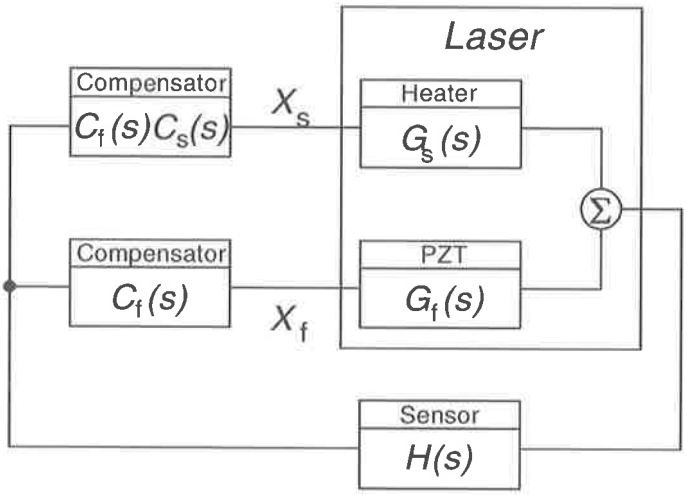

Figure 4.29: Alternative representation of the frequency stabiliser.

$$C_{sr} = 2.1 \times 10^{-3} \frac{1}{s}.$$

The transfer function of the acquire mode is also shown on figure 4.31. It has the same shape as the run mode in the region around unity gain but is flat below 72 mHz.

The slow compensator was implemented with the circuit shown in figure 4.32. JFET input amplifiers are used in the circuit because of the large resistor values required to realise the very low frequency corner at 72 mHz. The precision of this circuit is not crucial as any offsets simply led to a small offset in the PZT operating point. Nevertheless, an offset trim facility was included to ensure that the offset was not too large. The variable resistor in the preamplifier stage of the compensator allows the phase margin and hence the damping ratio of the loop to be optimised for good transient behaviour.

Figure 4.33 shows the response of the system when subjected to a step disturbance in the PZT output, which was produced by changing the dc offset on the PZT and hence the cavity resonance frequency. The fast compensator responds with a step of its own to ensure that the laser follows the cavity correctly. Once the fast

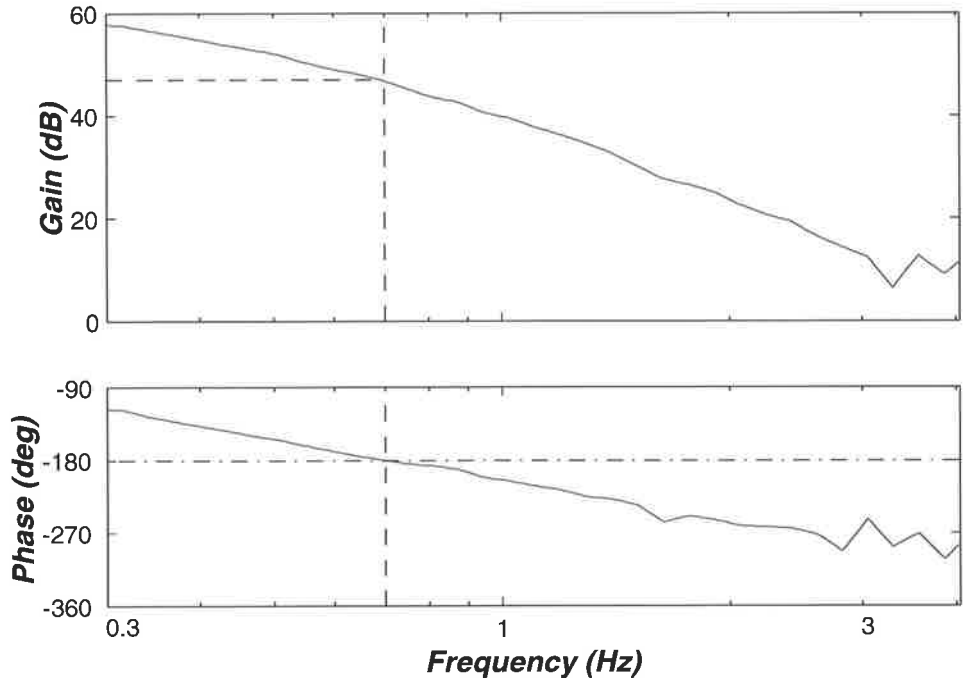

Figure 4.30: Bode plot of X_f/X_s. The dashed lines on the figure indicate the gain margin of 47 dB which occurs at a frequency at 0.7 Hz.

compensator output has moved from zero the slow compensator starts to integrate the error to drive the PZT back towards zero. Over the next fifteen seconds or so the heater is driven to the temperature that places the laser at the new cavity resonance frequency without significant DC at the PZT.

As can be seen the heater feedback exhibits a series of damped oscillations as it approaches its steady-state value. The figure shows that the frequency shift produced by these oscillations was compensated by the PZT. The size of the oscillations could be adjusted by changing the damping constant of the heater compensator. However, this adjustment requires a trade-off between speed of response and transient size. The setting shown in the figure was found to be adequate for producing stable long term lock of the system.

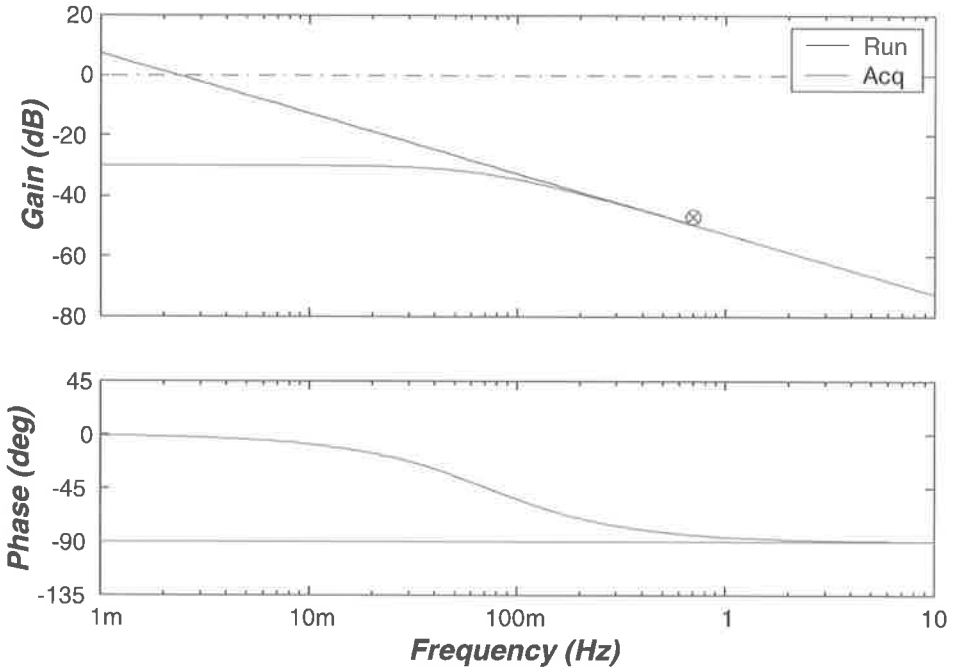

Figure 4.31: Nominal slow compensator transfer functions in run and acquire mode. The red cross within a circle indicates the -47 dB maximum allowable gain at 700 mHz.

4.2.6 Performance of the Laser Frequency Stabiliser

4.2.6.1 Reliability

In practice, the laser locking system was found to be very stable and was capable of keeping the laser locked to the reference cavity essentially indefinitely. The servo system was sufficiently robust that it did not unlock when the reference cavity was tapped with a screwdriver handle. Nonetheless, the system did occasionally unlock on large acoustic transients. The relocking procedure was straightforward and it typically took no more than twenty seconds to relock the system. The locking sequence was as follows:

1. Fast and slow compensators switched into "Acquire" mode;
2. PZT offset adjusted to bring the laser and cavity into resonance;
3. Fast compensator allowed to catch the resonance;
4. Fast compensator switched to "Run" mode to reduce noise;
5. Slow compensator switched to "Run" mode to drive the PZT to the centre of

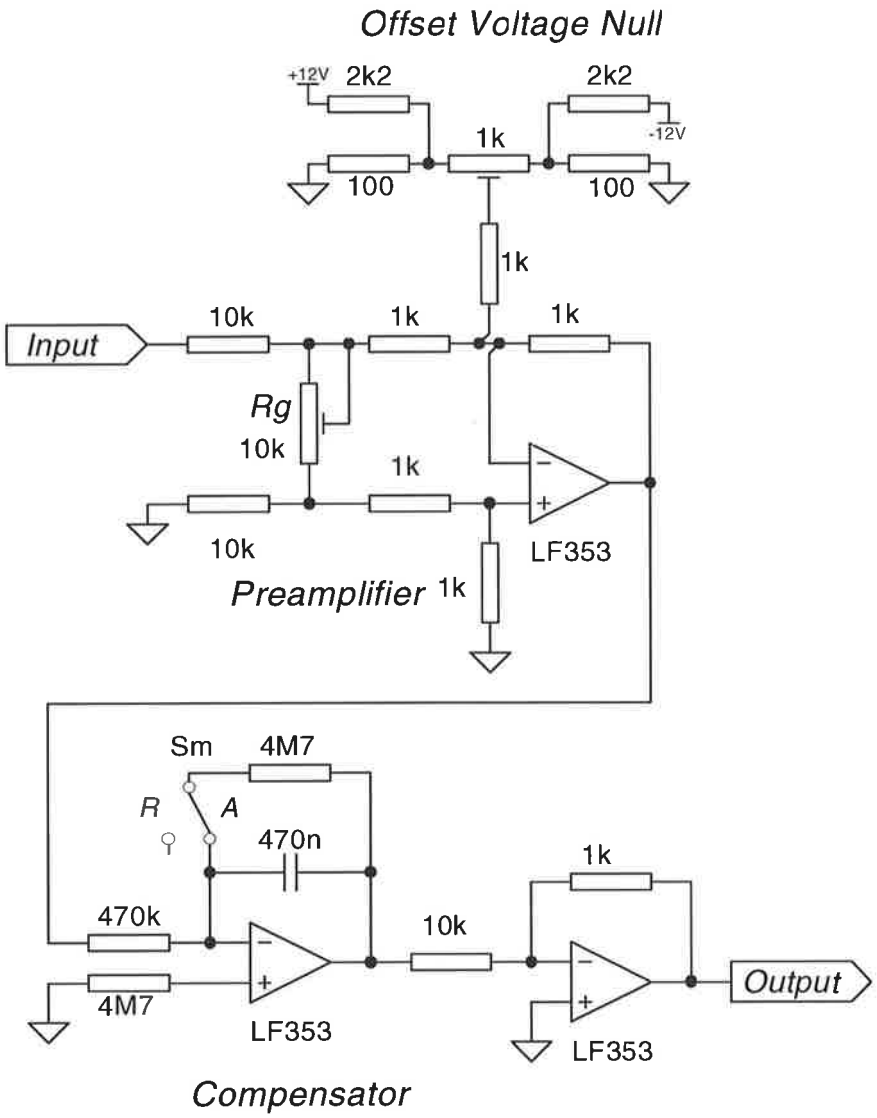

Figure 4.32: Schematic of the slow compensator used to drive the laser heater.

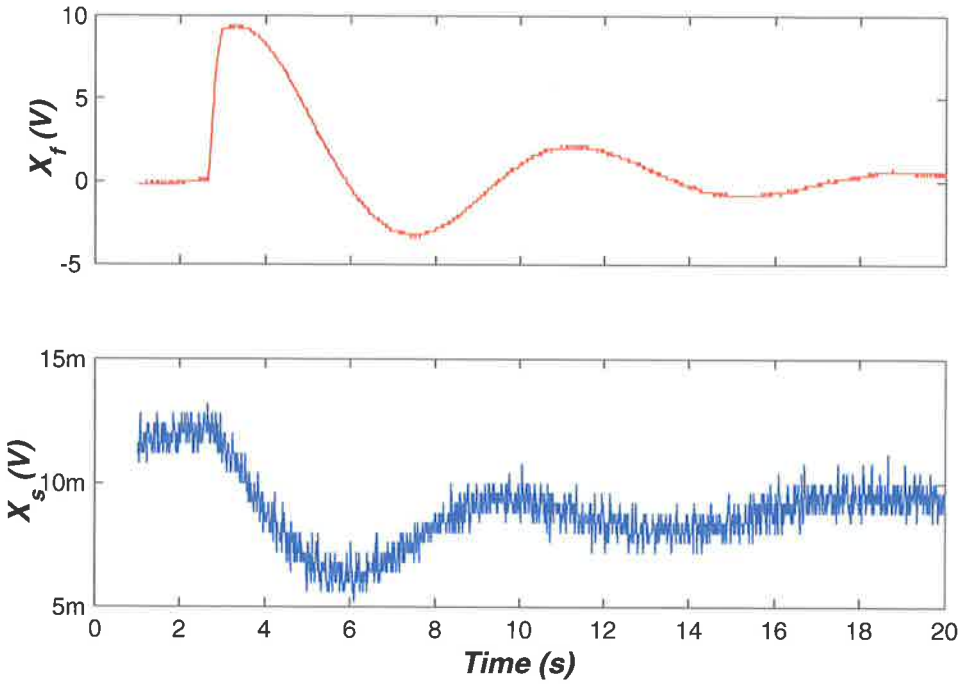

Figure 4.33: Response of the compensators when confronted with a step change in the reference cavity resonance frequency. The top curve shows the fast (PZT) compensator response. The bottom curve shows the slow (heater) response.

its range.

4.2.6.2 Measurement of Loop Sensitivity

The fast compensator shown in figure 4.24 includes a calibration port. A signal applied at the calibration point is injected into the compensator output with unity gain (though inverted). The disturbance injected into the control loop will be reduced by the sensitivity $S(s)$ of the loop. Thus, by measuring the reduction in the calibration signal we can determine the loop sensitivity, which will allow us to predict the closed loop frequency noise.

Consider the block diagram of the closed loop system as shown in figure 4.34. The dither signal $D(s)$ is added to the output of the compensator. The calibration transfer function $C_c(s)$ allows for any non-ideal transfer of the dither signal through to the output. In a closed loop system we know that the resulting signal at $X(s)$ will be given by

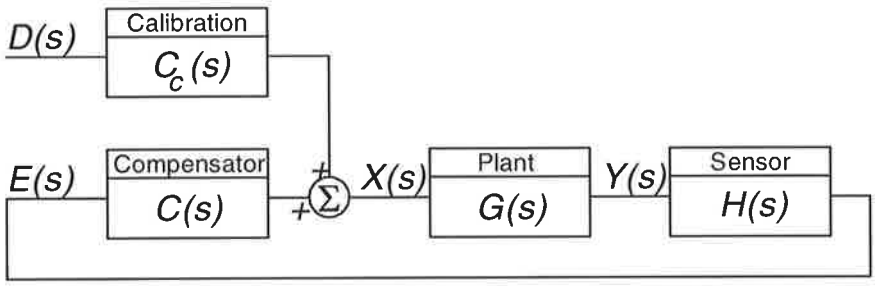

Figure 4.34: Block diagram of the closed loop system with its calibration input $D(s)$ shown.

$$\begin{aligned} X(s) &= D(s)C_c(s)S(s) \\ \frac{X(s)}{D(s)} &= C_c(s)S(s) \end{aligned}$$

Thus, $S(s)$ can be determined using measurements of $C_c(s)$ and the transfer function between D and X. However, the latter transfer function is difficult to measure due to the very small signal that is produced at X. It is considerably easier to measure the transfer function from $D(s)$ to $E(s)$ and use equation 4.7.

$$\begin{aligned} E(s) &= G(s)H(s)X(s), \\ \Rightarrow S(s) &= \frac{E(s)}{D(s)}\frac{1}{G(s)H(s)C_c(s)}. \end{aligned} \tag{4.7}$$

We have already measured the appropriate value for $G(s)$ in section 4.2.4.1 and $H(s)$ in section 4.2.3.4, where we found that $G(s) = 1.96\ MHz/V$ and $H(s) = 50\ \mu V/\sqrt{Hz}$ in the frequency region of interest. A measurement of $C_c(s)$ was performed using a signal analyser and the result is shown in figure 4.35. As expected the transfer function approximates an inverter with unity gain for most of the frequency range. At high frequencies the gain rolls off due to component mismatches in the output stage of the compensator.

The resulting estimate of $S(s)$ is shown in figure 4.36. The measured sensitivity closely follows the predicted curve except at high frequency where model uncertainty is likely to be high due to the presence of unmodelled high frequency PZT modes.

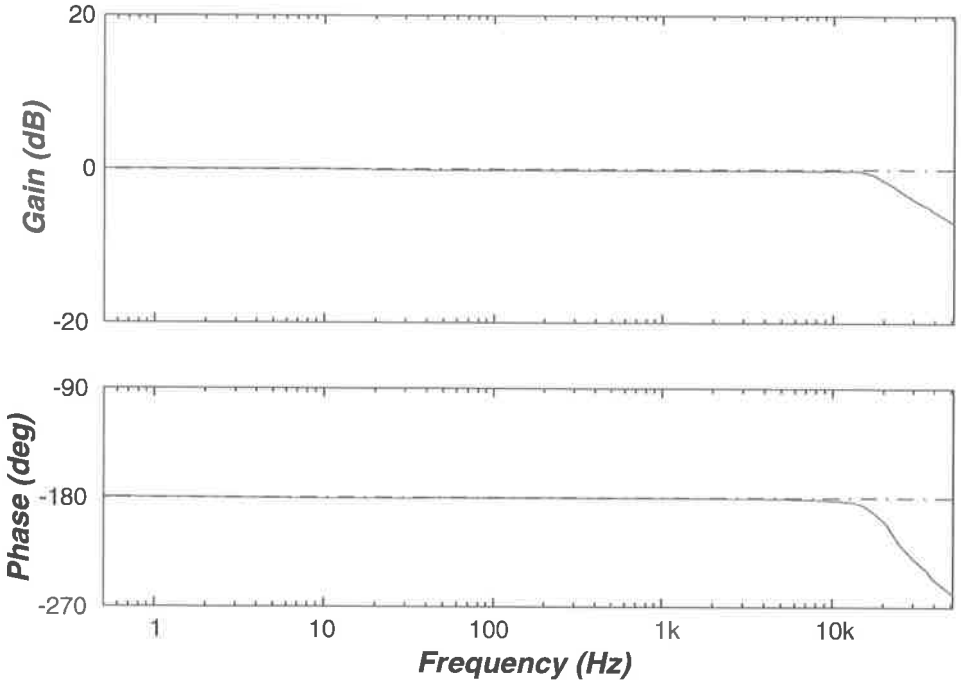

Figure 4.35: Transfer function from the calibration input of the compensator to the compensator output.

The direct measurement of this curve allows us to be confident in predicting the frequency noise resulting from the use of the servo system.

In figure 4.37 we see the prediction of the closed loop laser noise in the locked system. The noise is consistent with the experiment requirements except at the high end of the measurement range. Improvement of the noise in this region would require extension of the bandwidth of the system, possibly by adding a wideband frequency modulating element such as an electro-optic frequency modulator.

4.2.6.3 Closed Loop Noise Measurement

The signal from the PDH sensor provides an estimate of the relative frequency noise between the laser and the cavity. The spectrum of the PDH output voltage noise was recorded when the system was locked and converted to the equivalent frequency noise by dividing by the sensor gain. The measured noise spectra in acquire and run modes are shown in figure 4.38. The detector noise level measured with the photodetector covered and the shot noise level corresponding to the 60 mW of

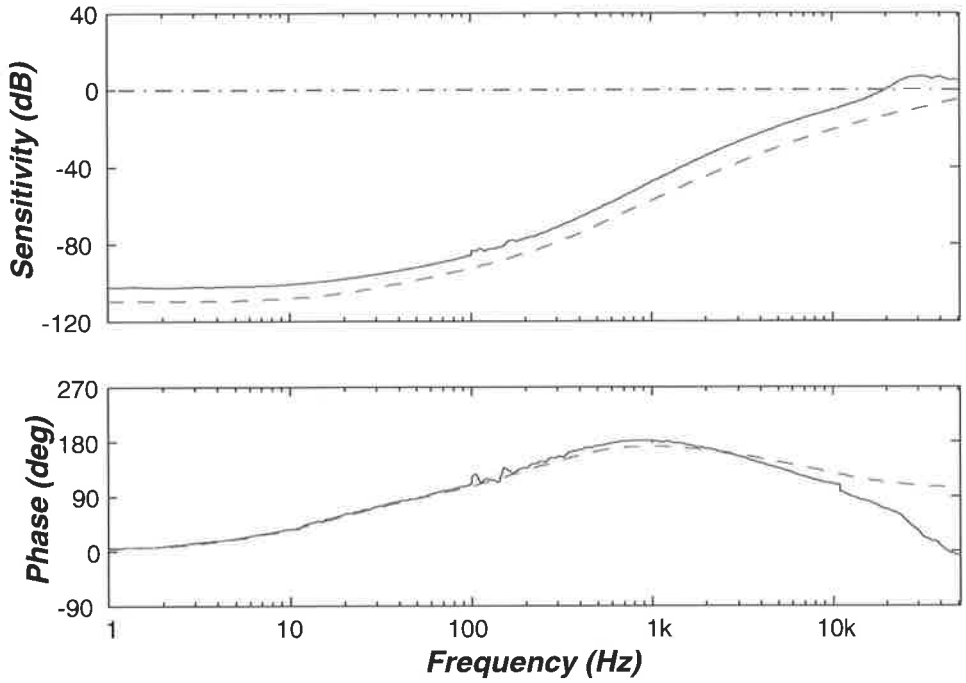

Figure 4.36: Measured sensitivity, $S(s)$, of the laser frequency stabiliser. The blue curve shows the measured sensitivity. The red dashed curve shows the predicted sensitivity.

detected power used in the experiment are also shown in the figure.

Figure 4.39 shows that the system does not perform quite as well as expected below 1 kHz. The noise curve is that of a white noise source at approximately 5×10^{-2} Hz/$\sqrt{Hz}$. As the compensator gain was increased from zero, the measured frequency noise improved until it reached this apparent floor. Further increases in compensator gain did not reduce the measured noise. We can therefore conclude that the noise is in the frequency sensor, where it is probably caused by the demodulation of RF pickup in the PDH signal line or of laser intensity noise produced in the EOM [79]. EOM's can produce amplitude modulation because of spatial or polarization modulation accompanying the desired phase modulation if the modulator is not perfectly aligned to the input beam [78,93]. Even with perfect alignment the EOM produces -60 dB residual intensity modulation when driven with a 1 radian phase modulation. While the alignment of the EOM was adjusted, the performance described above was the best achieved. It may be that further work on this part of the system could

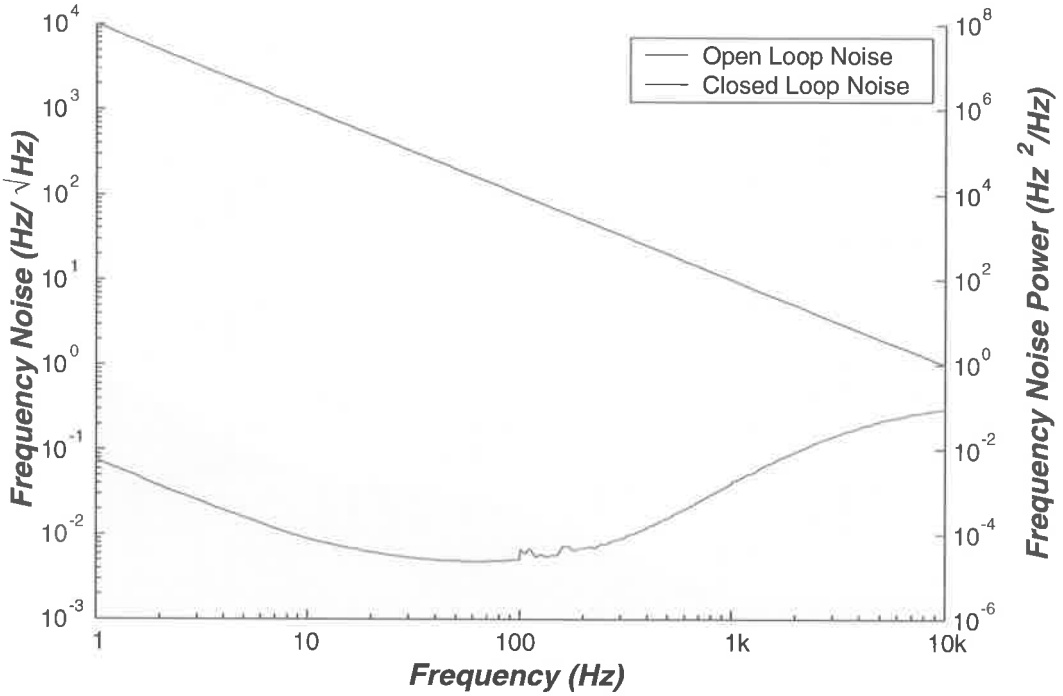

Figure 4.37: Predicted frequency noise of the laser locked to the reference cavity given our estimate of the open loop laser frequency noise and the measured sensitivity.

improve the amplitude modulation and thus the sensitivity limit. Modification to the layout of the RF electronics could also be made to reduce potential pick-up problems.

Nevertheless, the noise spectrum achieved satisfies the requirements for the experiment below 100 Hz and is in the acceptable range up to about 700 Hz. The figure also shows that the noise spectrum in the region 1-10 kHz is slightly better than that predicted by the theory. This suggests that the open loop laser frequency noise estimate of $1\times10^4/f$ was likely too pessimistic.

4.3 Test Cavity Frequency Readout

A simple Fringe Side Locking (FSL) technique is used to monitor the cavity resonance frequency as thermoelastic noise causes it to vary with respect to that of the probe laser. The resonance frequency of the reference cavity is locked to that of the

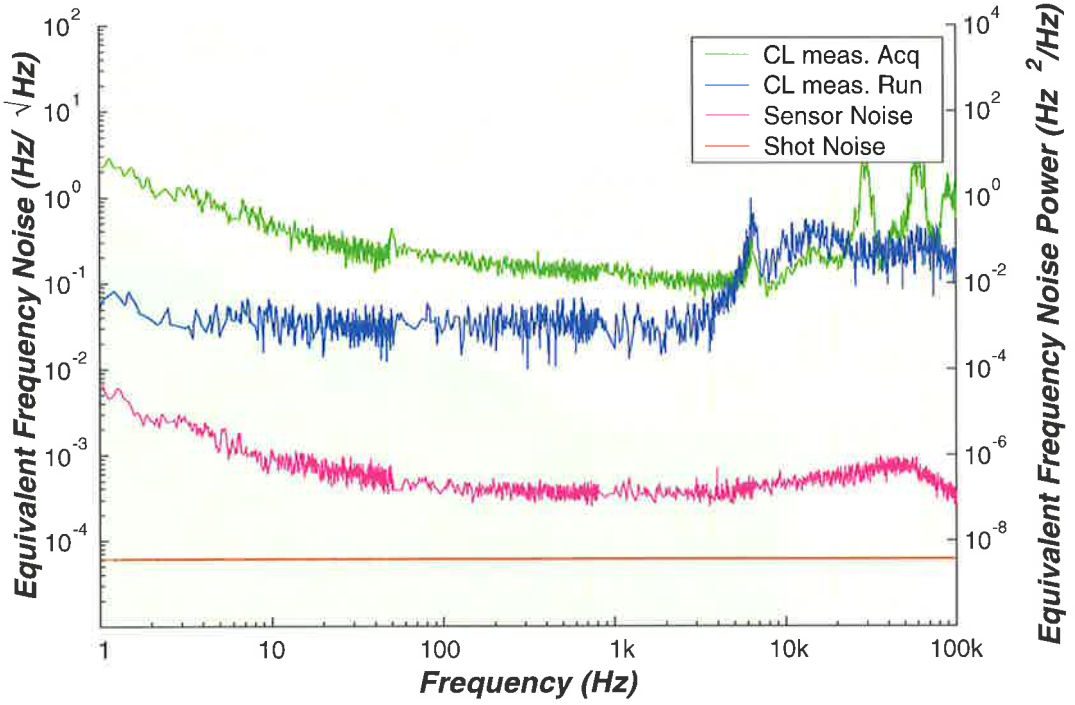

Figure 4.38: Frequency noise spectra in both acquire (green) and run (blue) modes. The magenta curve shows the detector noise floor and the red curve the shot noise limit for the experiemnt. Mains harmonics have been removed from the photodetector spectrum for clarity.

test cavity so that the probe laser power remains resonant in the test cavity. Details of the test cavity design are presented in section 4.3.1. The principle of fringe side locking is then discussed in section 4.3.2, as is its sensitivity to laser intensity noise. The actuator and compensators used to affect this read-out system are finally discussed in sections 4.3.3 and 4.3.4 respectively.

Within the locking bandwidth, the length noise in the test cavity is determined by examining the signal required to keep the two cavities in lock. At frequencies above the locking band, the thermoelastic noise level is measured at the FSL detector output.

4.3.1 Test Cavity Design

As discussed in section 2.2.1.1, the dimensions of the test cavity are determined by the need to ensure that thermoelastic noise in the test cavity remains the largest

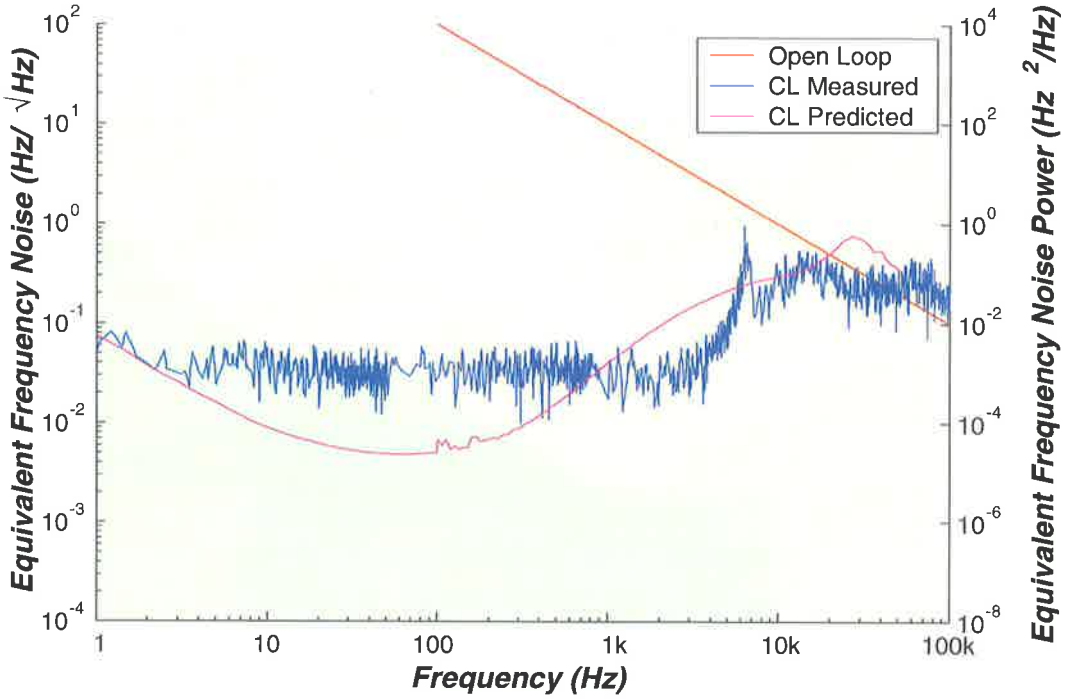

Figure 4.39: Measured closed loop frequency noise of the system. The blue curve shows the measured curve. Magenta shows the predicted noise curve and red the open loop noise estimate used in compensator design.

source of length noise in the experiment. In that section, we saw that the cylindrical sapphire substrate used to make the test cavity has a diameter of 50 mm and a length of 20 mm. A curved surface having radius of curvature of 250 mm was polished into one end of the substrate and mirrors with reflectivities of 99.5% at 1064 nm were coated onto each end.

While thermoelastic noise in sapphire is the target signal for this measurement system, the project philosophy is to allow the testing of many different test cavity materials and configurations. It was decided to keep the test cavity finesse requirement as low as possible to allow the examination of materials that exhibit high optical absorption. A nominal finesse of 1000 was selected as being achievable with a wide range of materials and cavity configurations.

Figure 4.40 shows the optical properties of the test cavity. The cavity has a free spectral range of 4.27 GHz and a transverse mode spacing of 520 MHz. As the test cavity was suspended, it was important to minimise errors due to motion of the

cavity. Such motion leads to variation in the efficiency with which the probe laser beam couples into the cavity and hence produce variations in the error signal [100]. To reduce these errors the spot sizes were made as large as possible. As shown in the figure, the beam size at the cavity input coupler is 172 μm, rising to 185 μm at the curved mirror.

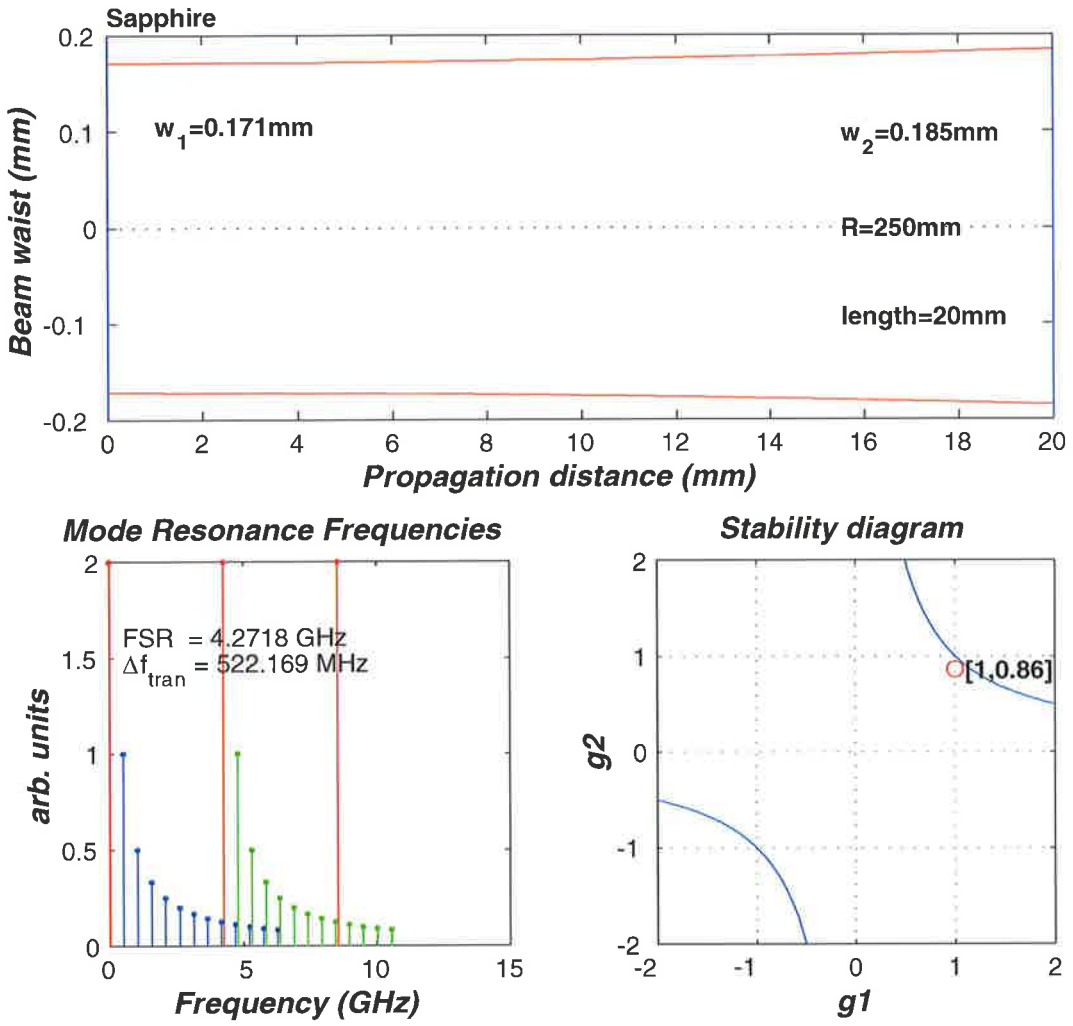

Figure 4.40: Optical properties of the test cavity. The top figure shows the beam radius of the mode inside the cavity. The lower left figure shows the frequencies of the cavity eigenmodes. The red lines indicate the various TEM_{00} modes, each separated by one cavity FSR. The blue and green modes are transverse modes corresponding to different longitudinal modes. The final figure shows the stability diagram for the cavity.

4.3.1.1 Dummy Cavity

The sapphire test cavity was not available for much of this project so a dummy cavity with similar optical properties was used for initial testing. The cavity was made using an aluminium spacer with standard dielectric mirrors clamped to the ends. These mirrors have the same curvatures as the sapphire mirrors but the spacer was lengthened to 35 mm to compensate for the fact that the refractive index of sapphire is 1.75 [101]. A photograph of the cavity is shown in figure 4.41.

Figure 4.41: Photograph of the dummy cavity.

The finesse of the dummy cavity was measured using the procedure described for the reference cavity (see section 4.2.2.1). Figure 4.42 shows a plot of the cavity's transmitted power as the NPRO (and its 17 MHz sidebands) were swept through the cavity resonance. The peaks are considerably wider than those in figure 4.6 because of the larger free spectral range in this case. Analysis of the data yields a finesse of about 2000 for the dummy cavity. This corresponds to mirror reflectivities of 99.92%, which is consistent with the manufacturers specifications [43].

4.3.1.2 Test Cavity Mode Matching

The mode exiting the reference cavity was modified to overlap the TEM_{00} test cavity mode. As was the case for the reference cavity mode-matching, the position of the mode matching optics was determined using Matlab code that modelled the optical system using ABCD matrices. The mode matching system, shown in figure 4.43, consists of a 250 mm focal length lens, L1, placed 530 mm from the reference cavity output coupler. The optical path distance between the two cavities is 845 mm. A

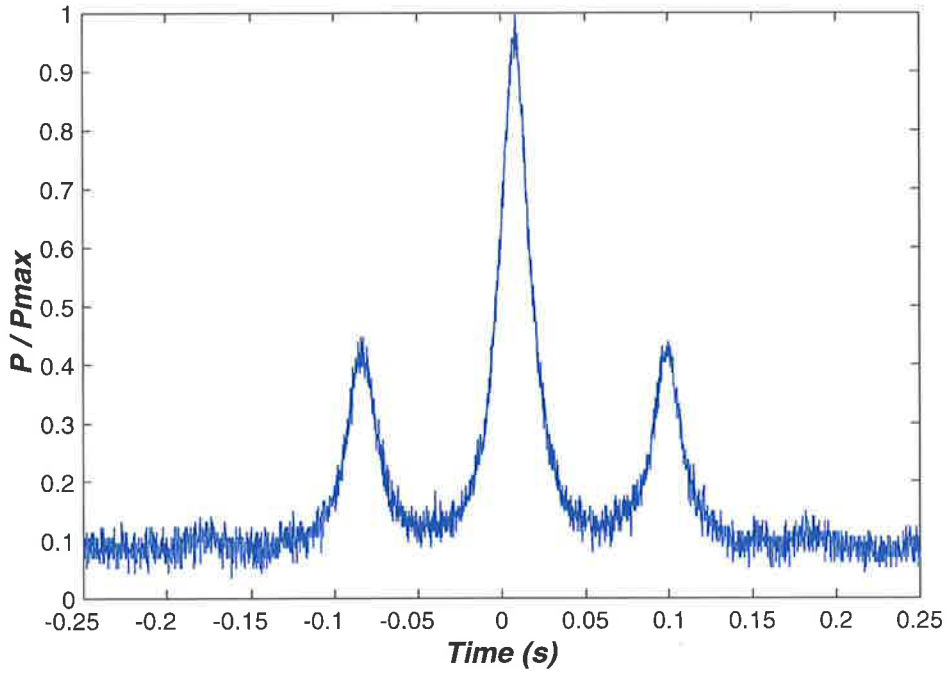

Figure 4.42: Transmitted power through the dummy cavity when a laser was tuned through the resonance. Peaks show the resonances of the main laser beam and its two sidebands.

translation stage and mirrors M3 and M4 are used to optimise the location of the input beam waist and the two beam steering mirrors M1 and M2 are used to align the probe beam with the test cavity. A plot of the beam radius between the two cavities is shown in figure 4.44.

4.3.2 Frequency Sensor

Fringe side locking (FSL) is used to detect variations in the test cavity resonance frequency. Fringe side locking has a significant advantage of being a modulation free technique. The use of PDH for our test cavity readout would have required the use of an extra electro-optic modulator (EOM) . However, if this modulator had been placed after the reference cavity, then it would have been in vacuum - a situation not recommended by the manufacturers [private communication, T. Kubo, New Focus, 2000]. A modulator after the reference cavity would not be optimal because it would have reduced the efficacy of the reference cavity's mode

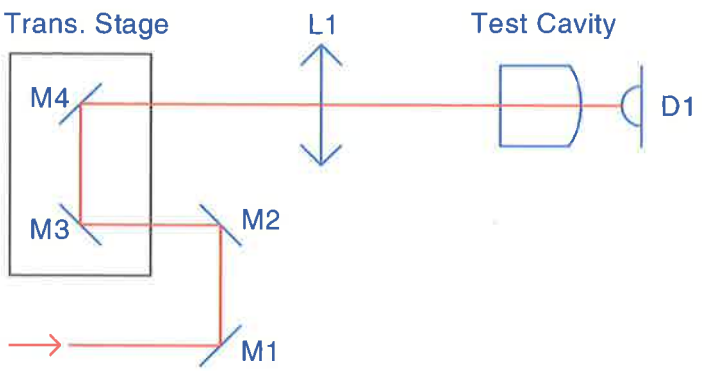

Figure 4.43: Optics used in modematching the prestabilised beam to the test cavity.

cleaning action by introducing geometry fluctuations into the probe laser beam at the modulation frequency. EOM's also invariably produce some polarisation rotation at their modulation frequency due to imperfect alignment. This polarization variation then leads to increased intensity noise when polarization optics are encountered [78, 79, 102].

An EOM could alternatively have been placed before the reference cavity and tuned to the FSR of the cavity so that the modulation sidebands would pass through the cavity. The cavity would have then have no effect on the phase modulation sidebands (nor on any undesirable amplitude modulation at this frequency), but it would provide filtering of geometric fluctuations of the beam. As discussed in section 4.2.2, the reference cavity FSR was chosen to allow this possibility and preliminary work was done on this topology. However, the system was not used in this project due to the complications inherent in dealing with high frequency RF signals (561 MHz). This method remains a viable technique for a future upgrade of the experiment if a greater sensitivity is required in future studies.

In contrast, FSL is simple, cheap and requires no specialised hardware or photodetectors. However, it has the disadvantages of a worse shot noise limited sensitivity than PDH and it is more sensitive to intensity noise. As will be shown below, the shot noise limit is not significant because we can use sufficient laser power to suppress the shot noise limit to below the required level. Sensitivity to intensity noise is more problematic and will be discussed in section 4.4.

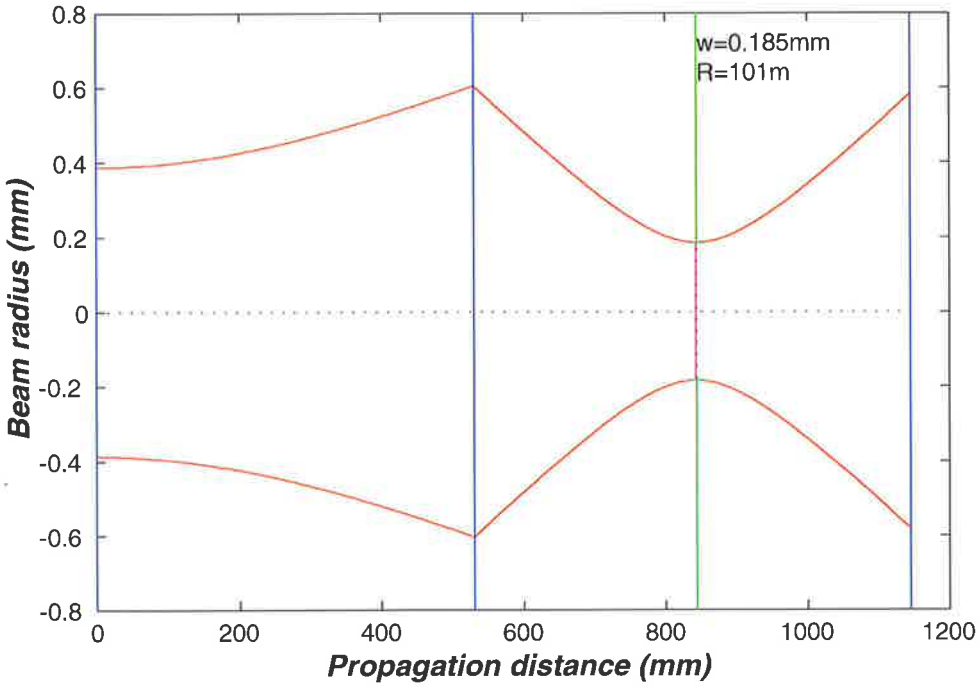

Figure 4.44: Beam radius between the reference cavity (z=0 mm) and the test cavity (z=845 mm).

Tilt locking [103] provides another alternative that could be used to measure test cavity frequency variation. In a tilt locking system, a slightly misaligned input beam is used to probe the test cavity mode. The power reflected from the cavity input coupler contains a mixture of cavity transverse modes. The resonant field leaking from the cavity interferes with the reflected field to produce an intensity pattern that depends on the leakage field's phase. When the cavity is on resonance, the phase of the leakage field is such that the resultant intensity pattern is spatially symmetric. However, when the cavity moves from resonance, the intensity pattern that is no longer symmetric. The degree to which the asymmetry is present is dependent on the frequency deviation from resonance and can be detected with a split photodiode to produce an error signal.

Like FSL, tilt locking is modulation free so would not introduce geometry fluctuations into the test cavity probe beam. However, the gain of a tilt locking sensor depends on the misalignment of the input laser beam relative to the cavity. In this experiment, the test cavity is suspended and therefore moves at the suspension sys-

tem resonant frequencies. Extensive modelling of the effect that this would have on a tilt locking error signal would be needed before it could be used with confidence.

4.3.2.1 Fringe Side Locking

In a typical fringe side locking (FSL) system, the power transmitted through the cavity is examined by a photodetector, as is shown in figure 4.45. As the cavity's resonant frequency changes with respect to the probe laser beam, the power transmitted though the cavity also changes. Thus, if we assume that the laser has no intrinsic power variability, we can infer changes in cavity frequency from the measured intensity variations. Similar systems can be constructed to examine the cavity in reflection or in a combination of transmission and reflection [104].

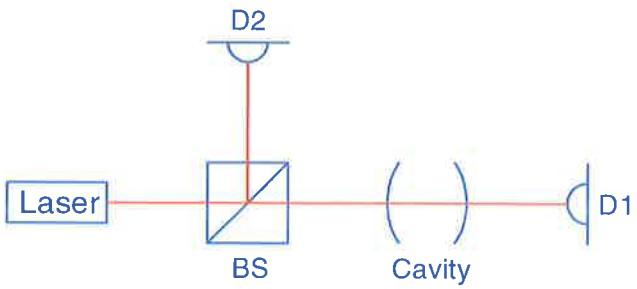

Figure 4.45: Schematic of a system using a fringe side frequency discriminator.

The operation of the FSL detector is illustrated in figure 4.46, which shows how variations in the frequency of the cavity relative to the laser, δf, couple to changes in the transmitted power, δP. If the frequency deviation is small we can regard the transfer to be linear with a gain that depends on the offset between the laser frequency and the cavity resonance. FSL systems are usually configured so that the transmitted intensity at the lock point is half that of the peak, an operating point that corresponds to a frequency offset from the centre of the resonance equal to half the FWHM linewidth.

The photodetector D1 detects the power transmitted through the cavity, which will vary according to the frequency difference between the probe laser and the cavity resonance. Unfortunately, any variation in probe laser intensity will also be seen at D1. This variation will be indistinguishable from the intensity variation that

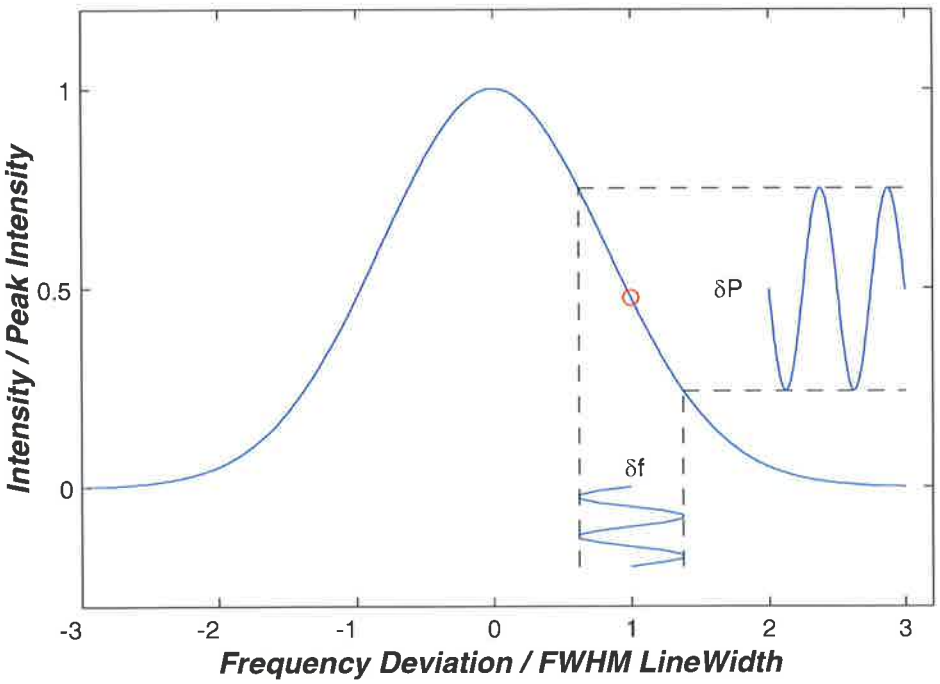

Figure 4.46: Transfer of frequency variation into intensity modulation using a fringe side locking frequency discriminator. The red circle indicates the quiescent lock point on the side of the fringe.

is caused by frequency difference fluctuation. However, the light intensity can be monitored using detector D2, and thus the laser intensity and frequency noise can be decoupled. This decoupling is most effectively conducted using a "noise-canceller" circuit as described in [105] and is limited only by shot noise in the two detectors.

An alternative approach is to use feedback from D2 to stabilise the laser intensity to the limit imposed by the sensor noise. This approach has the advantage of also reducing radiation pressure variations inside the test cavity and was therefore selected for this project. The implementation of the intensity stabilisation scheme is discussed in section 4.4.

In [106], [107] it is shown that at frequencies below that corresponding to the cavity round trip time, a cavity having mirrors of equal reflectivities will exhibit a

FSL response of

$$\begin{aligned} \frac{\delta P}{P} &= \frac{2T^2R\sin(4\pi L_{test}\Omega/c)}{1+R^2-2R\cos(4\pi L_{test}\Omega/c)}.\delta f \\ H_0 &= \frac{\delta V}{\delta f} = G\frac{2T^2R\sin(4\pi L_{test}\Omega/c)}{1+R^2-2R\cos(4\pi L_{test}\Omega/c)} \end{aligned} \tag{4.8}$$

where
- T is the (power) transmission coefficient of the cavity mirrors,
- R is the (power) reflection coefficient of the cavity mirrors,
- L_{test} is the length of the test cavity,
- Ω is the frequency offset of the lock point from the cavity line centre,
- G is the gain of the photodetector in V/W,
- H_0 is the nominal sensor gain.

4.3.2.2 Shot Noise

The shot noise limit of a fringe side locking detector is given in [79] as

$$S_{sn} = \frac{1}{FL}\sqrt{\frac{2hc^3}{\lambda P}} \tag{4.9}$$

Making the appropriate substitution for the test cavity length and assuming a worst-case test cavity finesse of 1000, we obtain the noise spectra shown in figure 4.47. The graph shows that the detected power must be of order 100 mW if the sensitivity goal of this project is to be reached. If this level proves to be too high then a longer test cavity will need to be used, or the power increased.

4.3.2.3 Design of the FSL Detector

Implementation of a FSL frequency discriminator requires the use of a photodetector that produces less noise than is produced by shot noise in the detected power. Because of the high transmitted power to be detected, a special high power photodetector was constructed that could be illuminated with 100 mW for extended periods. This detector and a low power version used for intensity stabilisation and the PDH sensor were modifications of the high power photodetector described in [108] and are further described in appendix E. Both types of detector were designed to have

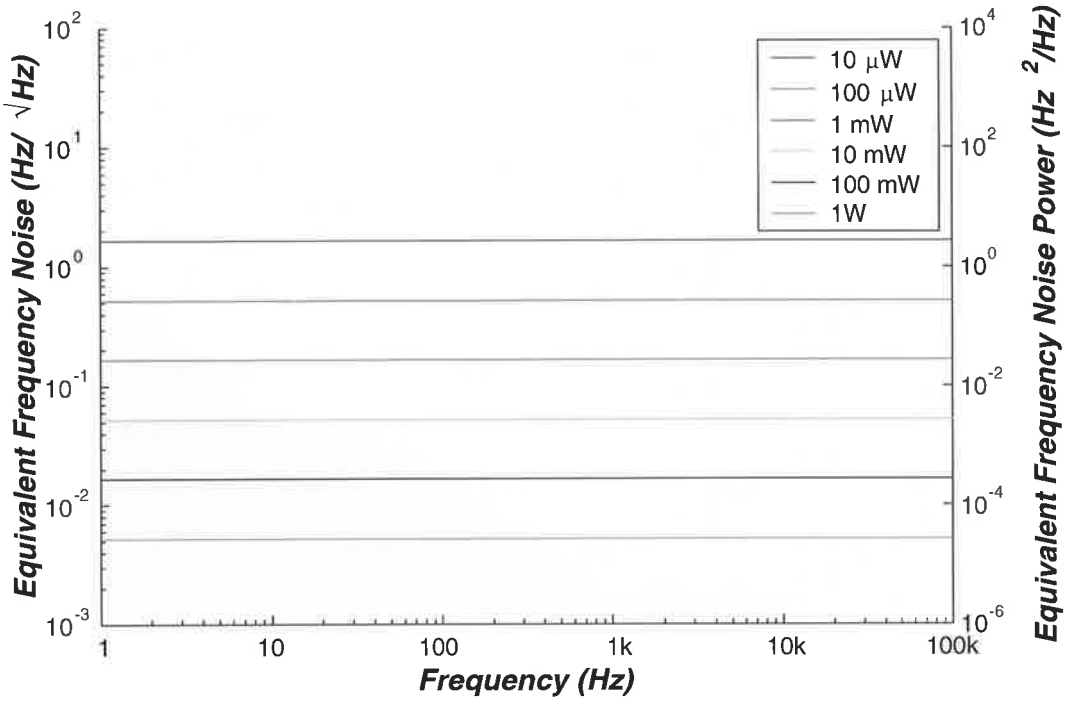

Figure 4.47: Shot noise limit for various probe powers on the fringe side detector.

bandwidths of 80 kHz.

The high power detector was designed so that an incident power of 100 mW would produce an output voltage of 10 V, corresponding to an offset from line-centre of half of the cavity linewidth. We assume a cavity finesse of 1000 and make the appropriate substitutions into equation 4.8 above, to obtain the nominal sensor gain, H_0.

$$H_0 \approx 4.7 \times 10^{-6}\ V/Hz$$

The frequency noise equivalent to the detector noise floor can be calculated by dividing the measured noise floor by H_0 and is shown in figure 4.48. Also shown is the frequency noise equivalent to shot noise level for 100 mW of detected power. As can be seen the noise floor is shot noise limited at frequencies beyond 100 Hz and is adequate to measure thermoelastic noise for all frequencies below 1 kHz. The figure shows that there is significant contamination of the detector signal by harmonics of the 50 Hz mains. These lines are narrow although the relatively poor resolution of

the graph makes the lines appear dense at high frequencies. However, the lines are well resolved for frequencies less than 10 kHz, so test cavity frequency noise will be visible between the lines. For clarity the graph includes the noise floor of the detector with the mains harmonics removed. The procedure for this line removal can be found in appendix G.

Temperature variations in the photodetector element can lead increased detector noise at low frequencies [109, 110]. When the detector is placed in a vacuum, the detector noise floor may well be reduced because of the removal of convective cooling variations. In a vacuum environment an actively cooled detector will be required to prevent excessive heating of the diode. The use of a temperature stabilised detector would further reduce thermal variations in diode temperature and thus improve the sensitivity of the FSL readout..

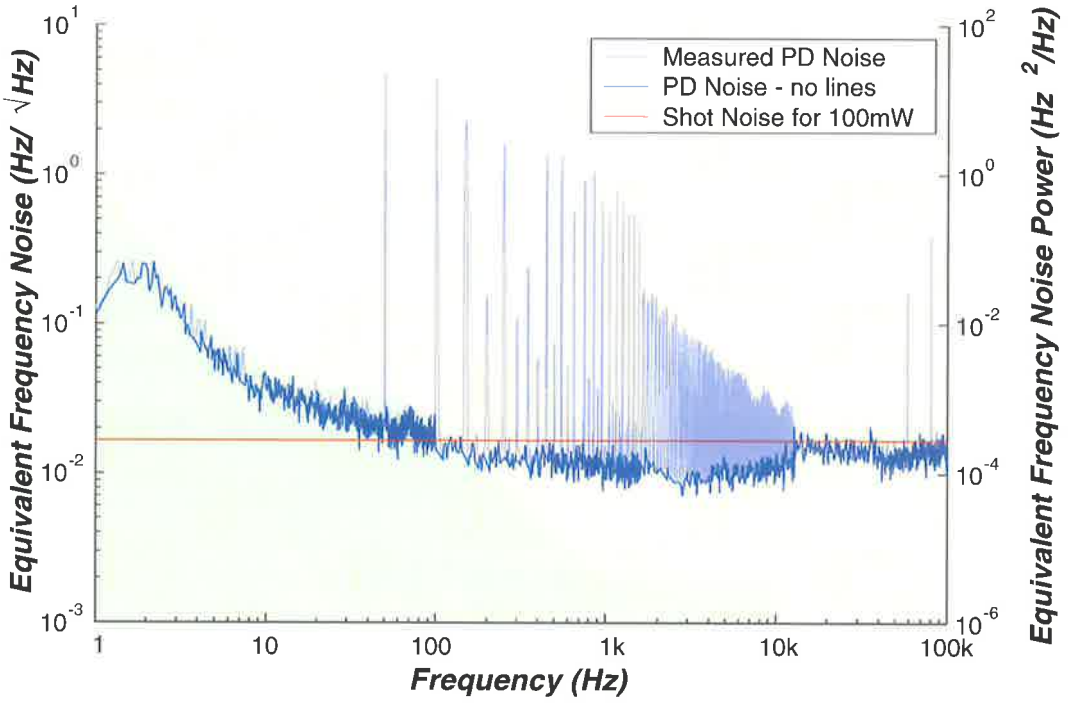

Figure 4.48: Detector noise floor and shot noise floor for the fringe side locking detector. The noise floor is shown with and without 50 Hz mains harmonics removed.

4.3.3 Actuator

As was discussed in section 4.2.2, the reference cavity was fitted with a piezoelectric transducer (PZT) so that its resonant frequency could be varied, thus allowing it to track the frequency of the test cavity. Design of an appropriate compensator to keep the two cavities locked required determination of the PZT transfer function and that of the high voltage amplifier used to drive the PZT.

4.3.3.1 Piezoelectric actuator

The reference cavity PZT response was measured using a simple Michelson interferometer as shown in figure 4.49. This measurement was conducted after the PZT had been mounted on the reference cavity spacer so that the rigidity of its mounting was the same as in the final system. This meant that the front side of the mirror was not accessible to a probe laser and thus the measurement was made through the back side of the PZT mounted mirror.

The interferometer included the PZT-mounted reference cavity mirror (M2) in one arm and a second arm was formed with standard HeNe mirrors M3 and M4. Because of the divergent effects of M2, the path from the beam splitter (BS1) to M4 was considerably longer than that from BS1 to M2. This ensured that the spot sizes produced by the two arms were the same, however the radii of curvatures of the wavefronts from the two arms were not matched. This produced a pattern of concentric rings on the screen at A1. However, the central spot of the pattern was sufficiently large to cover the photodetector element so this did not compromise the experiment.

A signal analyser was used to provide a small sinusoidal signal that was then amplified using a high voltage amplifier (HVA) capable of voltage swings of 400 V. Details of the amplifier design can be found in appendix E. The HVA had a DC offset adjustment facility that enabled the PZT position to be tuned until the interferometer sat halfway up a fringe. The sinusoidal variation of the PZT position produced a variation in the intensity on D1 that was compared to a sample of the HVA output, allowing determination of the PZT gain and phase at the test frequency. Variation of the test frequency then allowed the full frequency response

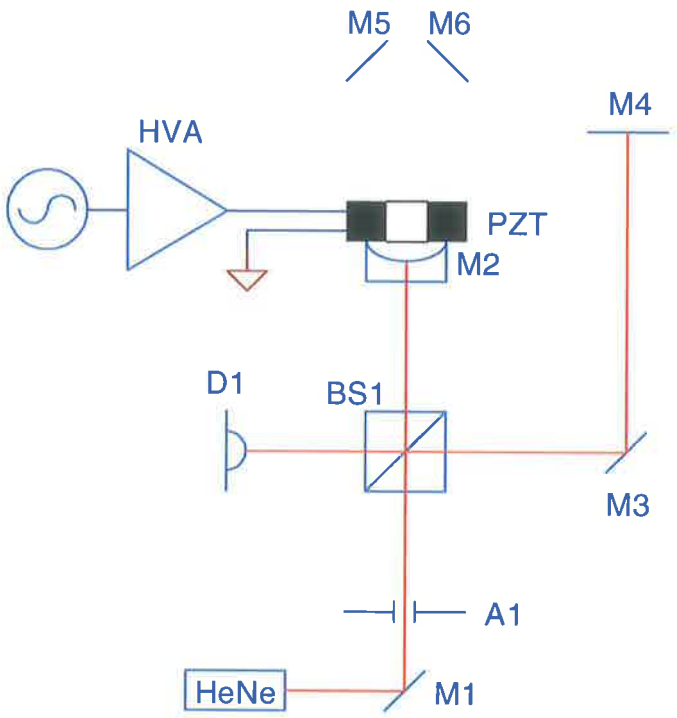

Figure 4.49: Schematic of the system used to measure the transfer function of the reference cavity PZT. HVA represents a high voltage amplifer used to rive the reference cavity PZT.

of the PZT to be determined. The measurement was repeated a number of times and averaged to form an estimate of the PZT response.

The frequency response obtained by the method above was augmented by a measurement of the low frequency gain of the PZT, in which the reference cavity was probed with the NPRO laser and then swept through a succession of resonances. The voltage at which each resonance occurred was recorded and the differences between each resonance yielded an estimate of the voltage required to drive the reference cavity through one FSR (561 MHz). Examination of these results led to an estimate of 9.8±0.2 $\times 10^4$ Hz/V for the DC gain of the PZT. The transfer function derived from this procedure is shown in figure 4.50.

A nominal model of the PZT response was calculated for use in compensator design and is shown in 4.10.

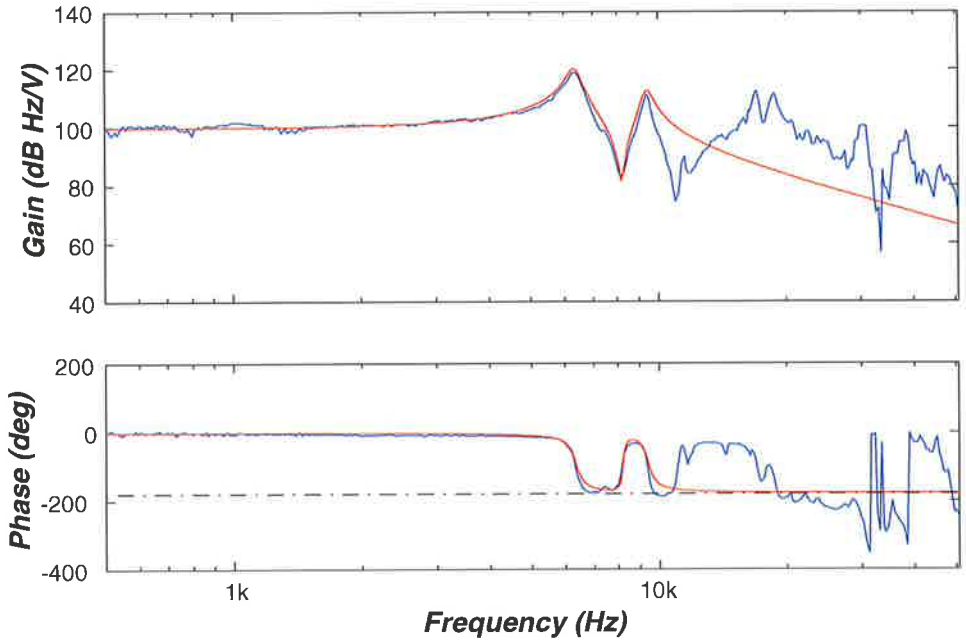

Figure 4.50: Transfer function of the reference cavity PZT (not including its driver amplifier). The experimental curve is shown in blue. The red curve shows the nominal model.

$$G_{PZT}(s) \approx 200 \times 10^{12} \frac{s^2 + 1030s + 2.7 \times 10^9}{(s^2 + 2800s + 1.6 \times 10^9)(s^2 + 3500s + 3.5 \times 10^9)} Hz/V. \tag{4.10}$$

The model included two resonances and one anti-resonance with properties as shown in table 4.2. It was not necessary to model higher frequency behaviour of the PZT as these features were well outside the intended bandwidth of the control loop and were therefore unimportant.

	Frequency (Hz)	Damping ratio (ζ)
Resonance 1	6340	0.035
Anti-resonance	8200	0.01
Resonance 2	9400	0.03

Table 4.2: Resonance properties in the nominal reference cavity PZT model

4.3.3.2 High Voltage Amplifier

The high voltage amplifier used to drive the PZT is described in appendix E. It has a gain of 100 and an output range of 400 V. The frequency response of the amplifier was measured while it was connected to the PZT and with its output set to 200 V. The result of this measurement is shown in figure 4.51. Note that the response is independent of frequency within the expected control bandwidth of the servo.

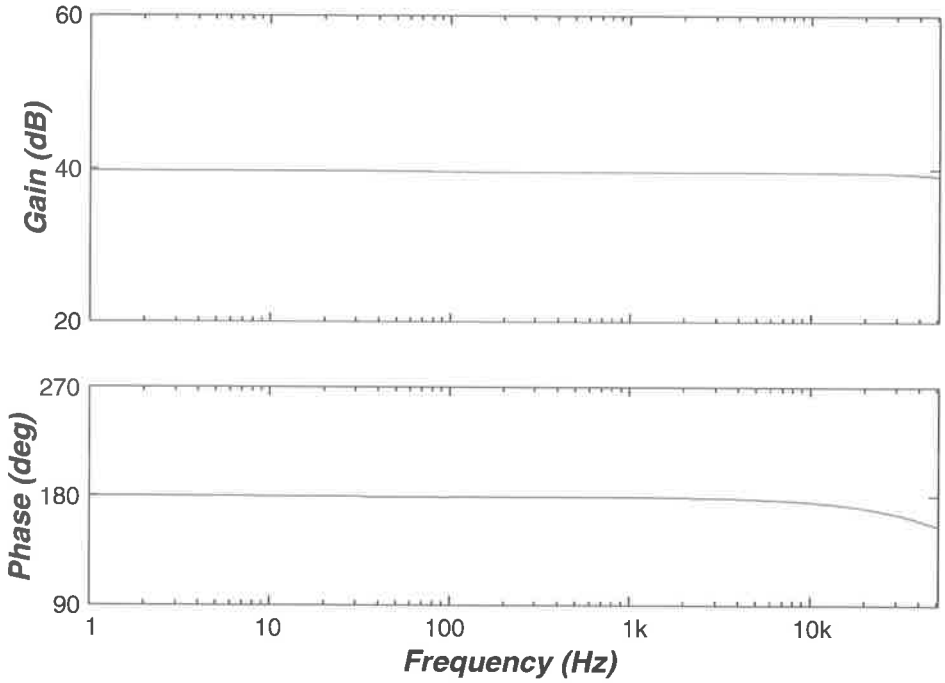

Figure 4.51: Frequency response of the reference cavity High Voltage Amplifier.

Voltage noise acting on the PZT acts as a source of reference cavity length noise. The spectrum of this length noise may be estimated by the product of the HVA output voltage noise, S_V, and the PZT's transfer function, G_{PZT}:

$$S_{PZT} = G_{PZT}^2 S_V$$

The HVA output noise was preamplified by the circuit shown in figure 4.52 to bring it above the signal analyser noise floor. The gain-bandwidth product of the OP27 operational amplifier limited the bandwidth of this circuit to about 80 kHz,

but this bandwidth was sufficiently large that the results were valid over the entire frequency range measurable by the signal analyser.

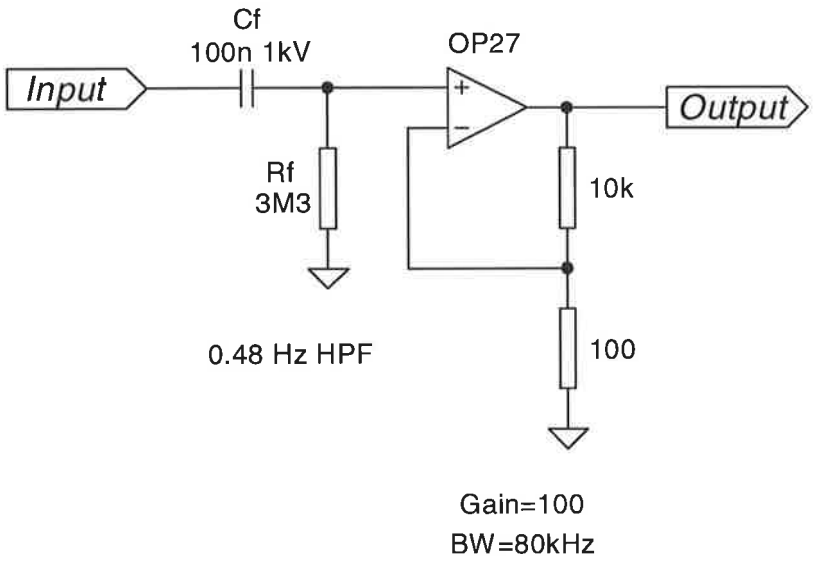

Figure 4.52: Schematic of the preamplifier used in the measurement of the high voltage amplifier noise

Figure 4.53 shows that the equivalent frequency noise of the high voltage amplifier was far too high. The noise floor of the preamplifier and the signal analyser are also shown in the figure to confirm that the measured HVA noise level is real.

The HVA was therefore fitted with a low-pass output filter to reduce the effects of its voltage noise. A high voltage filter with a corner at 1 Hz ensured that the HVA noise was reduced into the green region of the target sensitivity graph. This filter was made with a high voltage (1000 V) polyester 100 μF capacitor and a high voltage 1.5 $k\Omega$ resistor and is shown in figure 4.54. The noise of the amplifier with the filter fitted could not be directly measured, as it was lower than the noise floor of the preamplifier. The expected output noise can be obtained by multiplying the unfiltered noise by the measured low pass filter transfer function and this result is shown in figure 4.55.

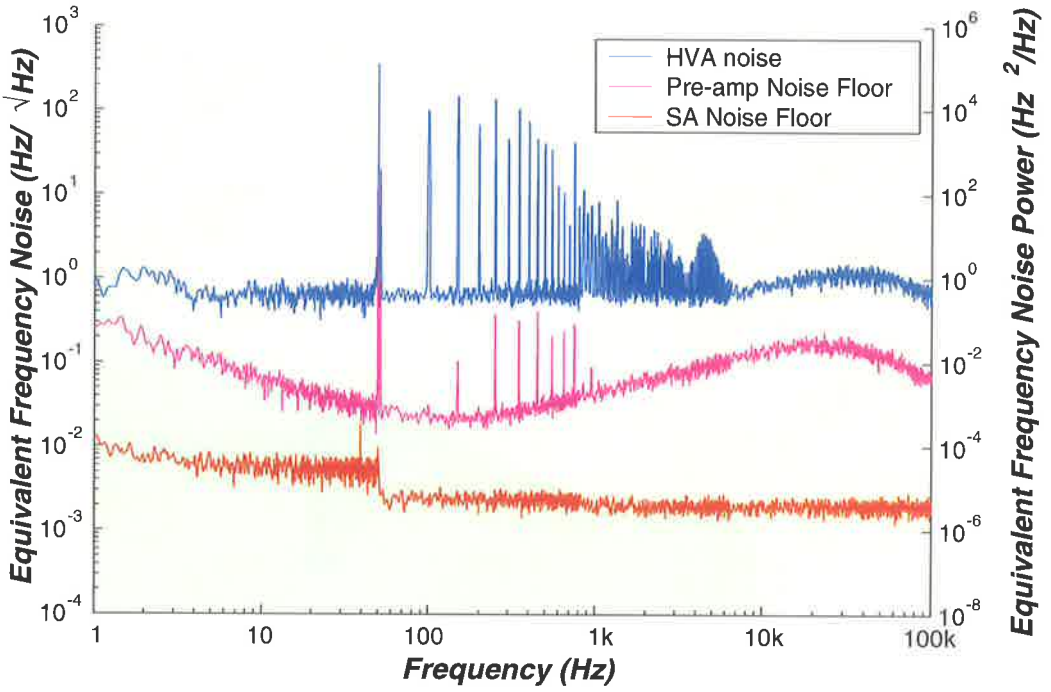

Figure 4.53: Equivalent frequency noise of the HVA voltage noise shown in blue. The magenta curve indicates the preamplifier noise floor, the red curve the instrumental noise floor.

4.3.4 Compensator Design

The compensator is required to keep the difference between the reference and test cavity resonant frequencies small. The precision with which this is accomplished is not particularly important so long as the frequency difference stays less than the cavity linewidths and thus preserves a usable error signal. We can determine the expected magnitude of the fluctuations in the frequency difference using

$$F_{rms} = \sqrt{\int_0^B S_{length}(f).df} \tag{4.11}$$

where $\begin{cases} S_{length} \text{ is the power spectral density of the frequency noise produced by the test cavity leng} \\ B \text{ is the bandwidth in which the noise is measured} \end{cases}$

Since the test cavity thermoelastic noise will be the dominant source of noise, the rms frequency deviation is expected to be 38 Hz measured in a bandwidth of 1

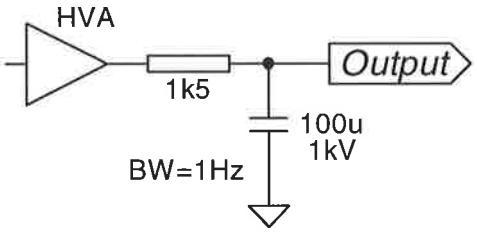

Figure 4.54: The schematic of the post-HVA filter.

kHz. Given that the width of the reference cavity resonance is of the order of 100 kHz and that of the test cavity is of order 4 MHz, it is clear that thermoelastic noise will not drive the two cavities out of resonance.

However, there may be long term drift in the cavity resonances, primarily those caused by temperature changes. The magnitude of such drifts are difficult to predict but their long time constants should make them simple to control with the use of the reference cavity PZT. Rather than define a minimum gain specification for the servo, it was decided to build a simple servo with as much gain as was consistent with stability and noise performance. Such a servo should vastly exceed the steady state requirements for the system while maximising performance during the locking transient.

The nominal transfer function of the HVA (including the 1 Hz LPF), the actuator and the FSL sensor is shown in figure 4.56. To ensure stability, the loop gain must be less than unity at the PZT resonances in the region 6-10 kHz. Thus the compensator can have a gain of up to 22 dB at the first resonance frequency. The transfer function of a simple compensator that will adjust the loop gain is shown in figure 4.57. It has a gain of 20 dB at moderate frequencies and an additional pole at 22 kHz attenuates the system response beyond the unity gain frequency. Below 1 Hz the compensator gain increases so that the system looks like an integrator at low frequencies. This ensures that the two cavities have no dc offset at steady state (assuming no dc offsets elsewhere in the servo loop). The open loop gain for the whole system is also shown in the figure.

The circuit for the compensator is shown in figure 4.58. *Rset* sets the reference voltage to which the transmitted cavity power is compared, and it is used to vary

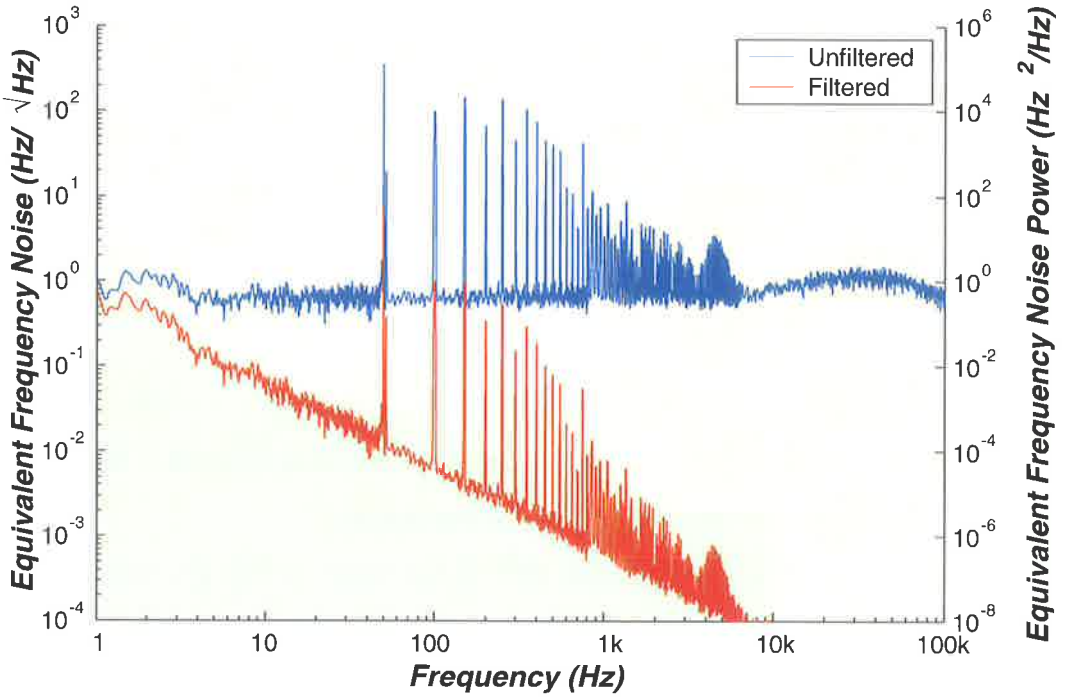

Figure 4.55: Inferred equivalent frequency noise of the HVA voltage noise with and without the post filter present.

the locking position on the side of the fringe. *Rg* allows the preamplifier gain to be adjusted and *Radj* allows trimming of the frequency at which the low frequency boost cuts in. Adjustment of *Radj* allows the corner frequency of the compensator to be adjusted to match the corner of the post-HVA filter, so that the overall response is that of a simple integrator. The switch $Sref$ allows the setpoint to be derived from an external source rather than the internal voltage. This external setpoint could be derived from the photodetector measuring the cavity input power to produce an additional measure of isolation from laser intensity noise.

As was discussed during the design of the laser frequency stabiliser, integrator saturation can easily occur when the system is unlocked. Thus, the test cavity tracking compensator must be a two stage system that allows the integrator action to be bypassed when the system is out of lock. In run mode, the compensator adds the extra low frequency gain to make the combination of the compensator and the HVA post-filter look like an integrator. In acquire mode, the low frequency gain boost is removed, which results in a simple lag compensator. In each case the

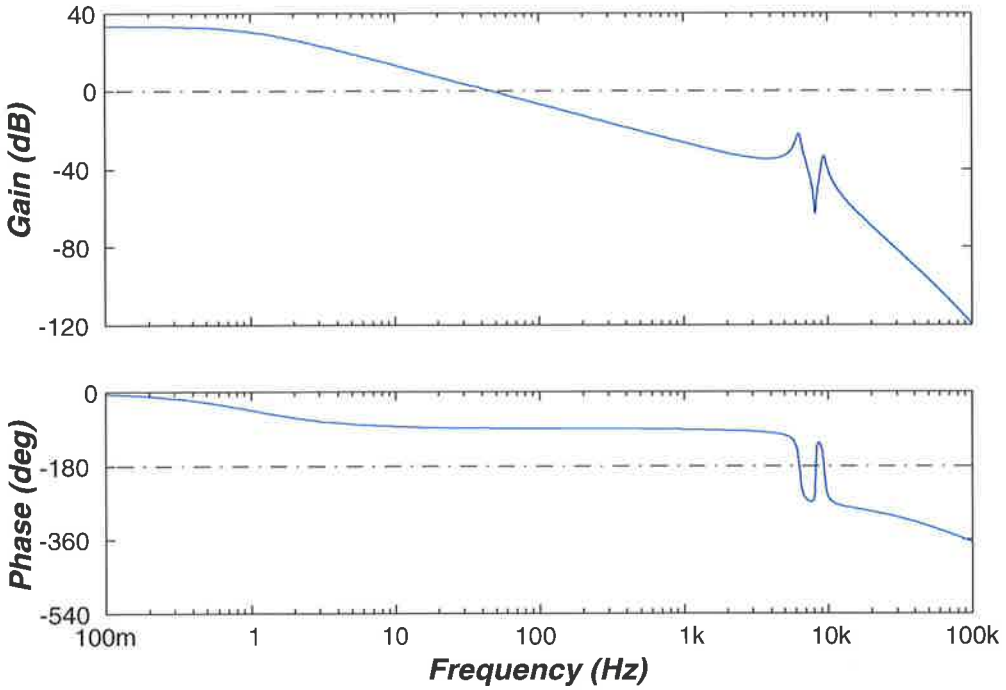

Figure 4.56: Transfer function of the high voltage amplifier, PZT actuator and FSL sensor.

compensator provides a gain of 10 at medium frequencies so that the unity gain frequency (and hence the system stability) is unchanged. Figure 4.59 shows the transfer functions for the two operating modes. The nominal open loop responses in each of the two compensator modes are shown in figure 4.60.

Electronic noise in the compensator was modelled using the SPICE simulator. Noise arising in passive components is automatically handled by SPICE, and noise in the operational amplifiers was introduced using the procedure outlined in [99]. This procedure yielded a voltage noise spectrum referred to the input of the compensator and this spectrum was multiplied by the loop's input sensitivity to find the expected noise contribution in the closed loop case. The resultant equivalent frequency noise spectrum is shown in figure 4.61. Note that this spectrum is of the noise contribution of the compensator alone, the HVA noise having been discussed above. As can be seen, the compensator electronic noise does not limit the experiment's ability to sense test cavity thermoelastic noise.

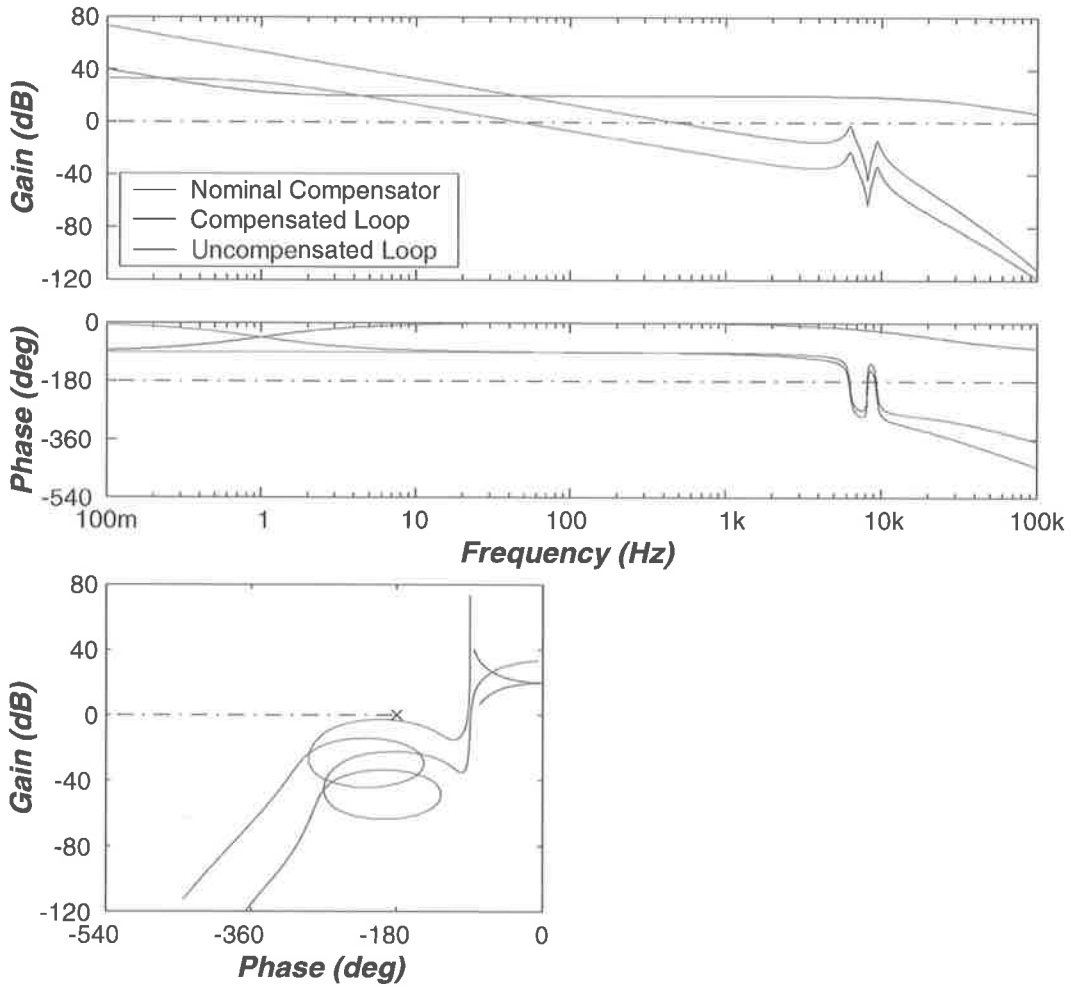

Figure 4.57: Bode and Nichols plots of the nominal compensator response and the resultant nominal open loop FSL system. The Nichols plot shows that the system will be closed loop stable. The uncompensated system is also shown for comparison.

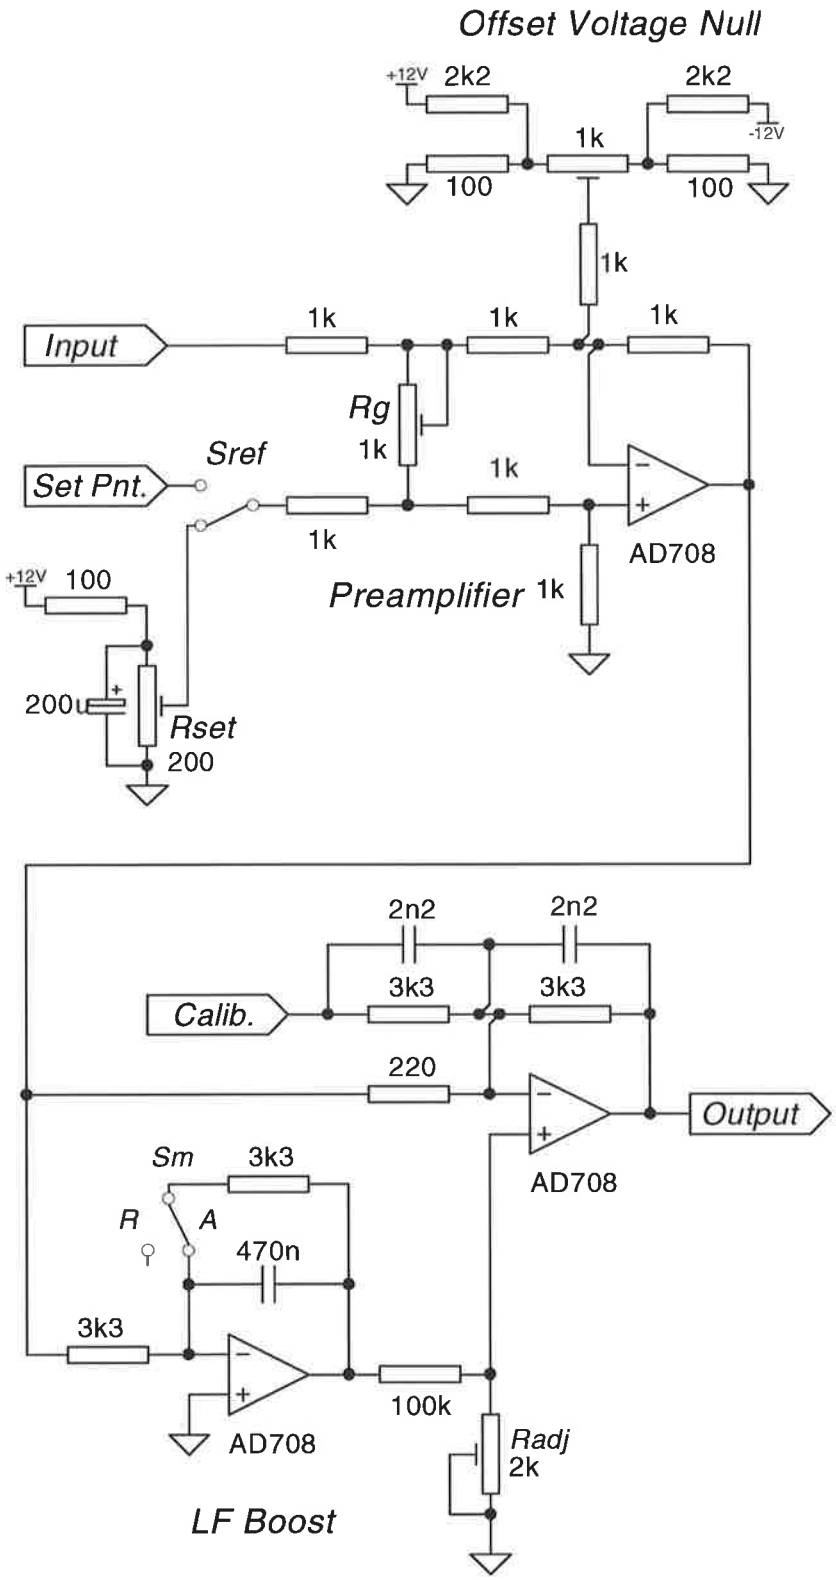

Figure 4.58: Schematic of the FSL compensator used to force the resonant frequency of the reference cavity to follow that of the test cavity.

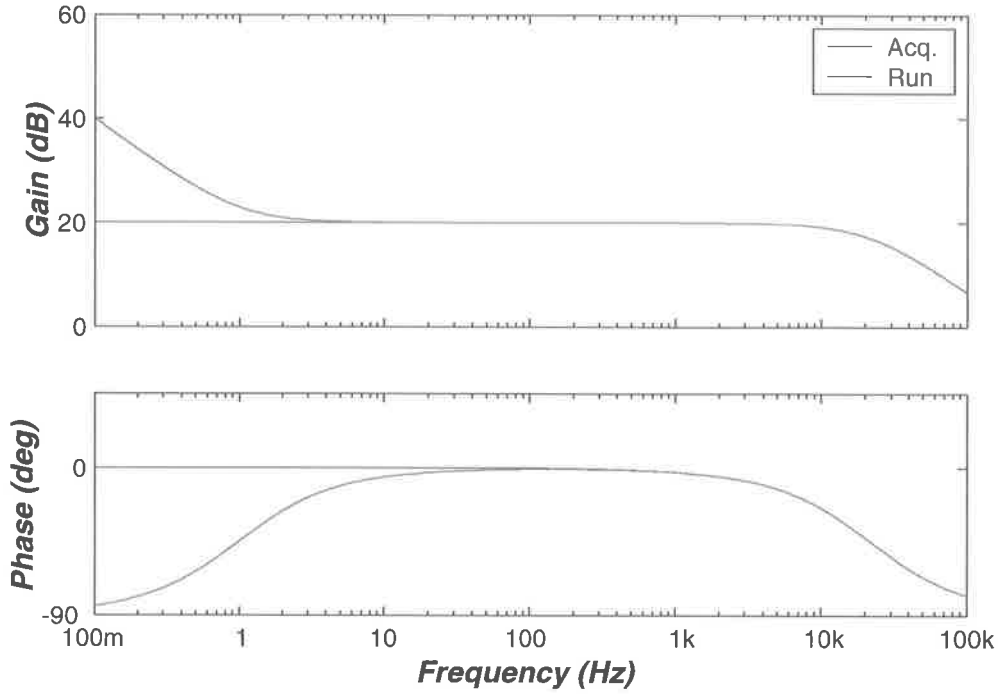

Figure 4.59: FSL compensator transfer functions in acquire and run modes.

4.4 Laser Intensity Stabilisation

As discussed earlier, laser intensity noise can contaminate the readout by producing variation in the power on the FSL detector. In section 4.4.2 we will see that the laser has a relative intensity noise of S_{RIN} =$5\times10^{-5}/f$. If we divide this noise by the FSL sensor gain, H_0, we obtain the equivalent frequency noise that the intensity variation causes in the test cavity readout, S_{coup}:

$$\begin{aligned} S_{coup}^{1/2} &= S_{v,intensity}^{1/2}\frac{1}{H_0} \\ &= V_0 S_{RIN}^{1/2}\frac{1}{H_0} \\ &\approx \frac{100}{f}\ Hz/\sqrt{Hz}, \end{aligned}$$

where $\begin{cases} S_{v,intensity}^{1/2} \text{is the voltage noise seen at the output of the FSL detector} \\ V_0 \text{ is the dc voltage produced by the FSL detector.} \end{cases}$

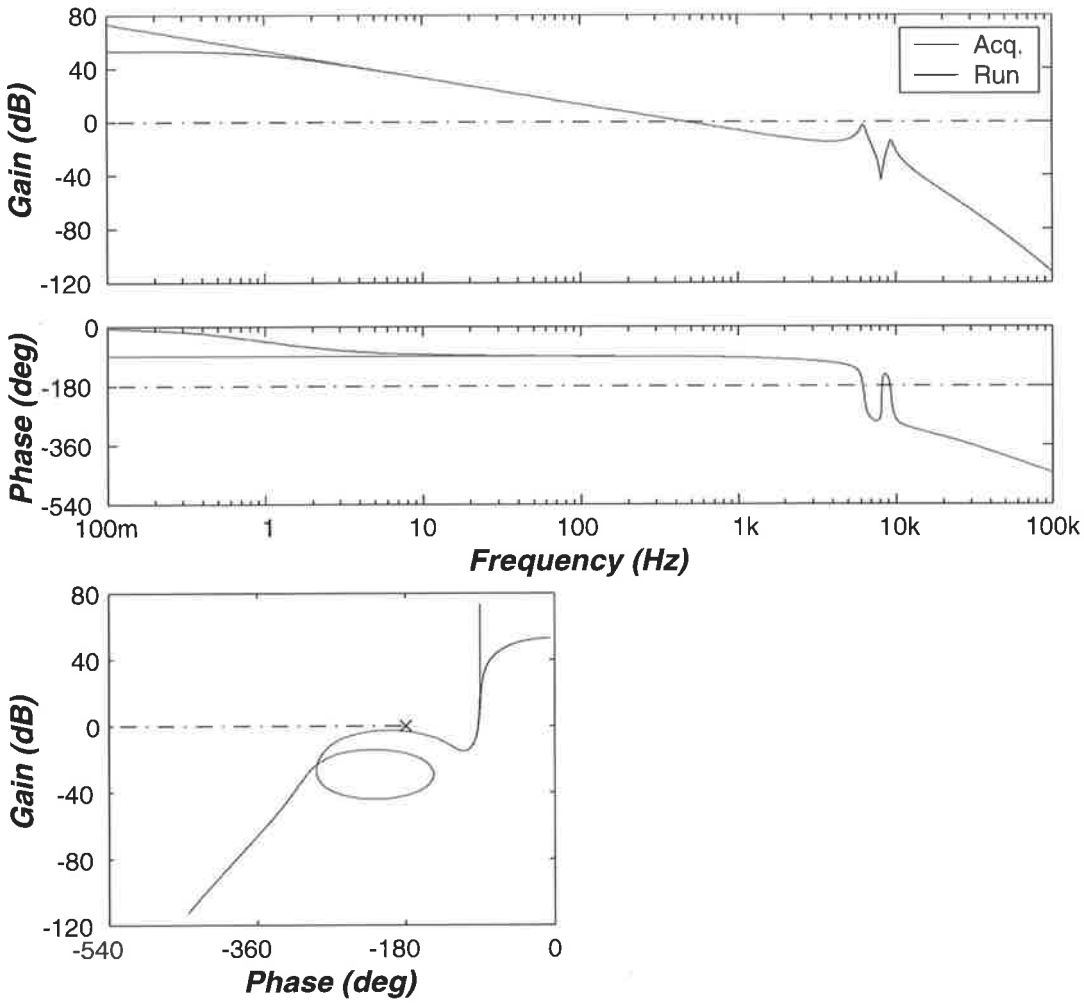

Figure 4.60: Nominal open loop FSL system responses with the compensator in acquire and run modes. The plot at the bottom shows the Nichols plot of the system and shows that the closed loop system will be stable.

This noise spectrum in shown in figure 4.62. As can be seen the intensity noise causes an unacceptable level of noise in the readout system. The laser intensity noise must therefore be reduced using a feedback loop. The reduction in intensity noise will also reduce radiation pressure effects in the reference and test cavities, although these are not expected to prove troublesome in this experiment (see sections 2.3.2.4 and 2.3.4.5).

The intensity stabilisation is achieved by sampling the light before the reference cavity and applying feedback to the laser pump diodes. This modulates the laser power and thus removes intensity fluctuations. The sensor is briefly discussed in

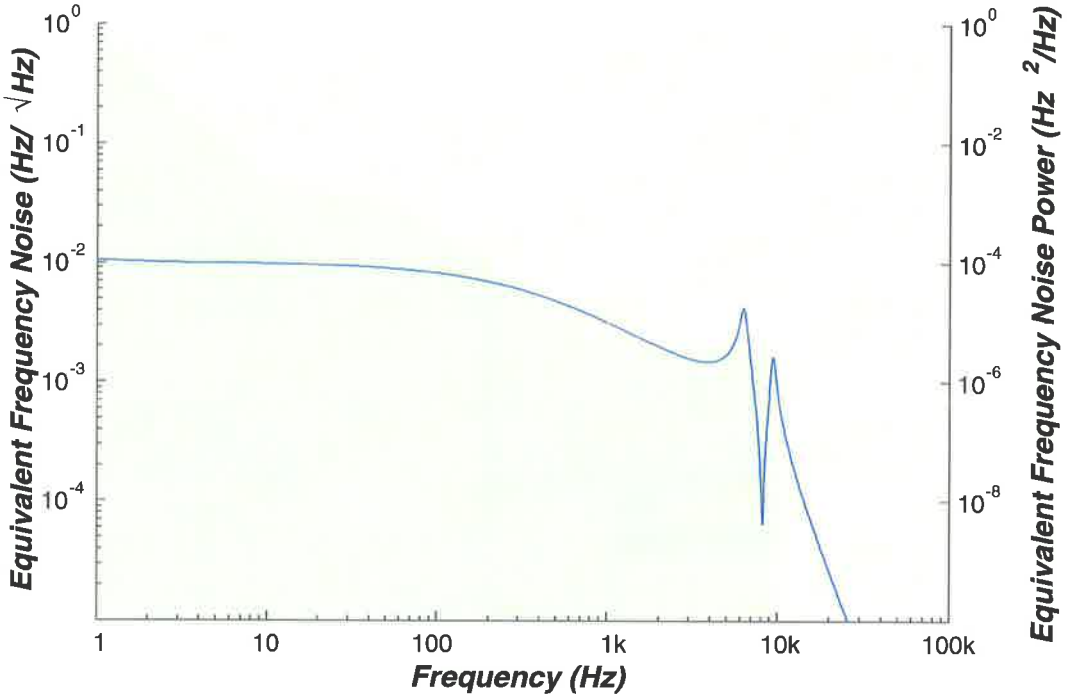

Figure 4.61: Equivalent frequency noise caused by electronic noise in the FSL compensator. The background colours indicate the required sensitivity regions.

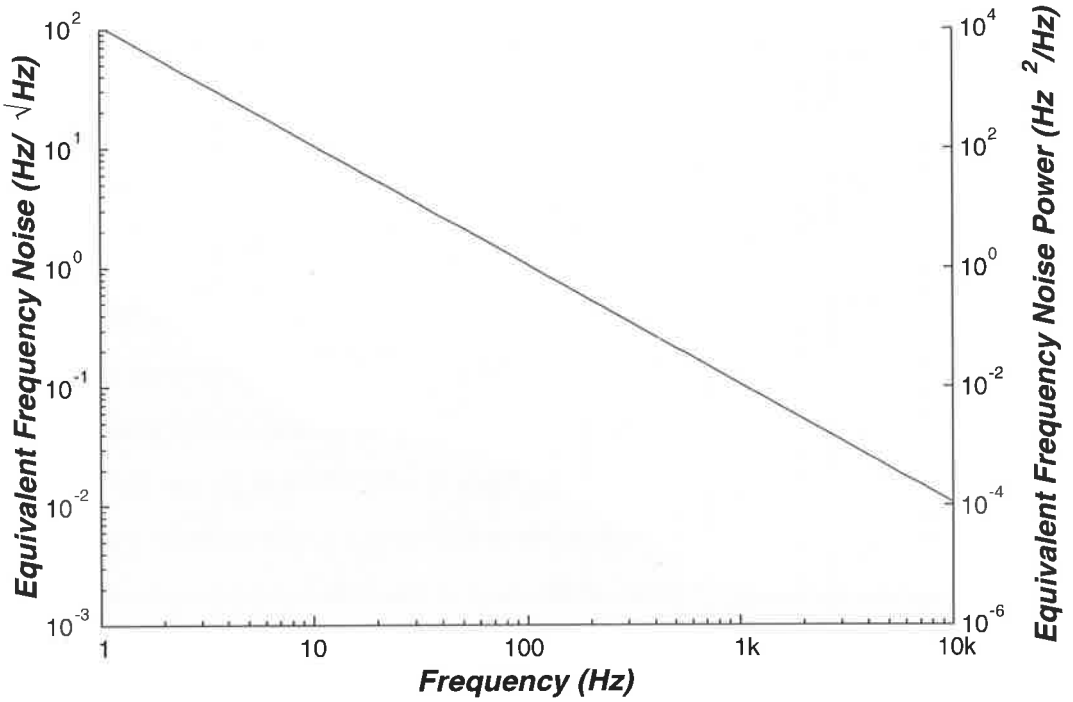

Figure 4.62: Equivalent frequency noise due to the intensity noise of the free running laser.

section 4.4.1, before the actuator and compensator are discussed in sections 4.4.2 and 4.4.3. The performance of the intensity stabiliser is presented in section 4.4.4.

4.4.1 Sensor

The sensor used to measure the laser power employs the photodetector design used for laser frequency stabilisation, as described in appendix E. The photodetector gain is 100 V/W and it has a bandwidth of 80 kHz.

4.4.2 Actuator

The laser used in this study is equipped with two intensity modulation inputs. The first input controls the setpoint of the laser's internal intensity stabilisation system. It is located on the laser control unit, where it interfaces directly with the laser's current setting circuitry. This input operates from DC up to approximately 10 kHz but was found to exhibit some non-linearity.

A second intensity modulation input is provided on the laser head. A voltage applied at this point produces a laser modulation current that is summed with that produced by the laser control unit. This input has a higher bandwidth than the first but is AC coupled. It can optionally be used by an inbuilt "noise-eater" that controls laser intensity noise and in particular drastically reduces the size of the high frequency relaxation oscillation.

Figure 4.63 shows the intensity noise of the laser with and without the noise-eater engaged. As can be seen from the figure, the operation of the noise-eater makes the intensity stabilisation task considerably easier by performing a large amount of noise reduction in the region above 100 Hz. Thus it was decided to use the noise-eater and implement the intensity stabilisation using the set-point modulation input only.

The open-loop intensity noise of the laser is modelled as having a relative intensity noise spectrum of $S_{RIN}^{1/2} = 5 \times 10^{-5}/f$ with the noise-eater operational. As shown in figure 4.63, this model overestimates the noise at high frequencies but should be adequate for design purposes.

The transfer function measured between the set-point modulation port and the output of the sensor is shown in figure 4.64. This transfer function is the product of the actuator and sensor transfer functions. A striking feature of this transfer function is the non-linearity: high input amplitudes increase the gain at the resonance and move it to lower frequencies. When the noise-eater was switched off the non-linearity disappeared, so it appears that the non-linearity is produced by interaction between the set-point modulation and the noise-eater electronics. The benefits of the noise-eater were considered sufficient to persist with its use, despite the complications caused by the non-linearity.

A nominal transfer function, G_o,was constructed to model the actuator and is also shown in figure 4.64. As can be seen, G_o approximates the low frequency behavior of the plant reasonably well but is less accurate at high frequencies. Note that the nominal plant transfer function does not include the additional phase roll-off that occurs at high frequencies in the real transfer functions. Consequently our plant model is not valid at frequencies beyond about 20 kHz. The nominal transfer function between the modulation input and the sensor output can be expressed as

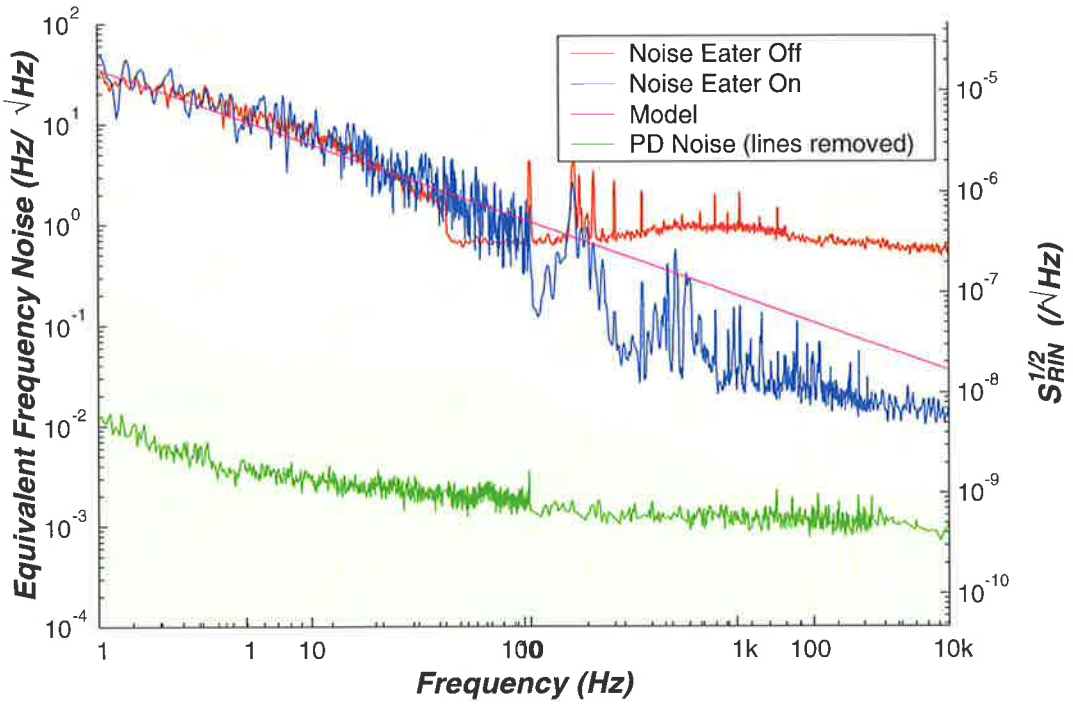

Figure 4.63: Open loop intensity noise of the laser with and without the noise eater operational. The right axis shows the intensity noise and its equivalent frequency noise is shown on the left axis. The magenta curve shows the $S_{RIN}^{1/2} = 5 \times 10^{-4}/f$ model for the noise. The noise produced by the photodetector is shown in green. Mains harmonics have been removed from the photodetector spectrum for clarity. The coloured regions show the required sensitivity regions.

$$G_o(s) = \frac{53 \times 10^{12}(s + 1.6 \times 10^3)}{(s + 1.3 \times 10^5)(s + 160)(s^2 + 1.0 \times 10^4 s + 2.9 \times 10^9)}$$

4.4.3 Compensator Design

As the plant transfer function is non-linear we could not expect a simple classical controller design to be optimal (or even stable) once implemented. However, a design derived from the assumption of plant linearity provides a reasonable starting point from where empirical tuning can be used to optimise the loop to cope with the plant nonlinearities.

Figure 4.65 shows the open loop intensity noise of the laser along with the target noise level. As discussed in appendix B, the ratio of these two curves yields the

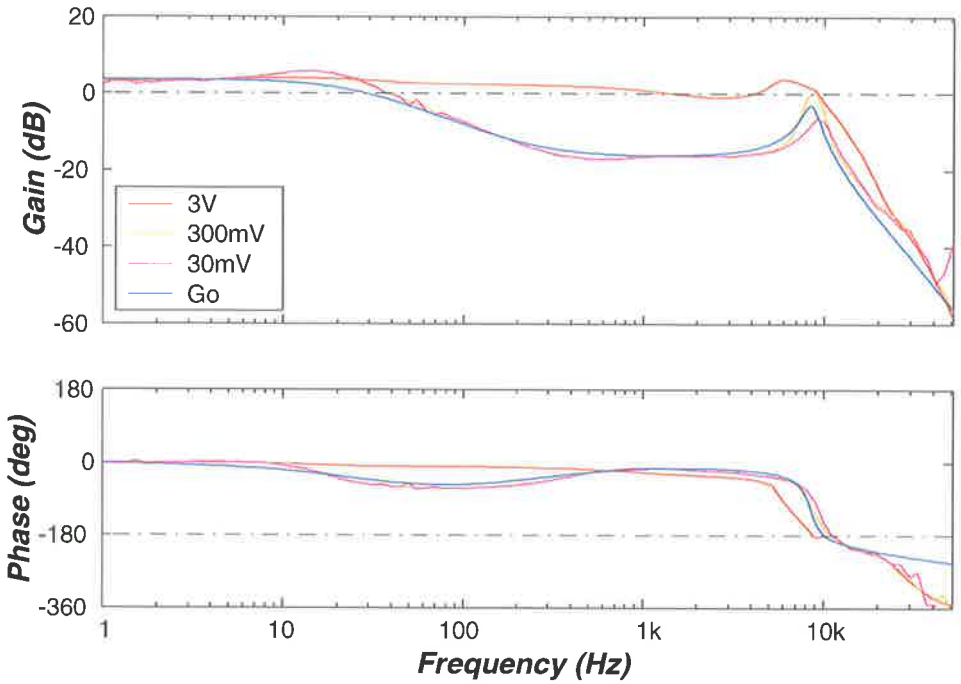

Figure 4.64: Transfer functions of the intensity modulation system for different sized input signals. The nominal transfer function is shown in blue.

required loop gain, which is shown in the bottom part of the figure. The design task then becomes one of meeting (or exceeding) the required gain for as large a bandwidth as can be achieved while maintaining loop stability.

A compensator was designed to keep the gain as high as reasonable throughout the low frequency region while achieving a unity gain frequency just below the actuator resonance at about 10 kHz. The unmodelled additional high frequency phase lag in the real plant transfer function meant that attempting to design a wider bandwidth controller would be futile. Figure 4.66 shows the nominal plant transfer function and the nominal compensator transfer function with the corresponding nominal loop gain. As can be seen the nominal system will be closed loop stable.

The nominal loop gain achieved by the servo design is shown in figure 4.67, along with the target loop gain. As can be seen, the system exceeds the required gain at low frequencies. However, at frequencies above about 300 Hz, the gain is no longer sufficient to meet the requirements. Thus for frequencies above 300 Hz, we expect the intensity noise to exceed our desired level. The gain of the compensator

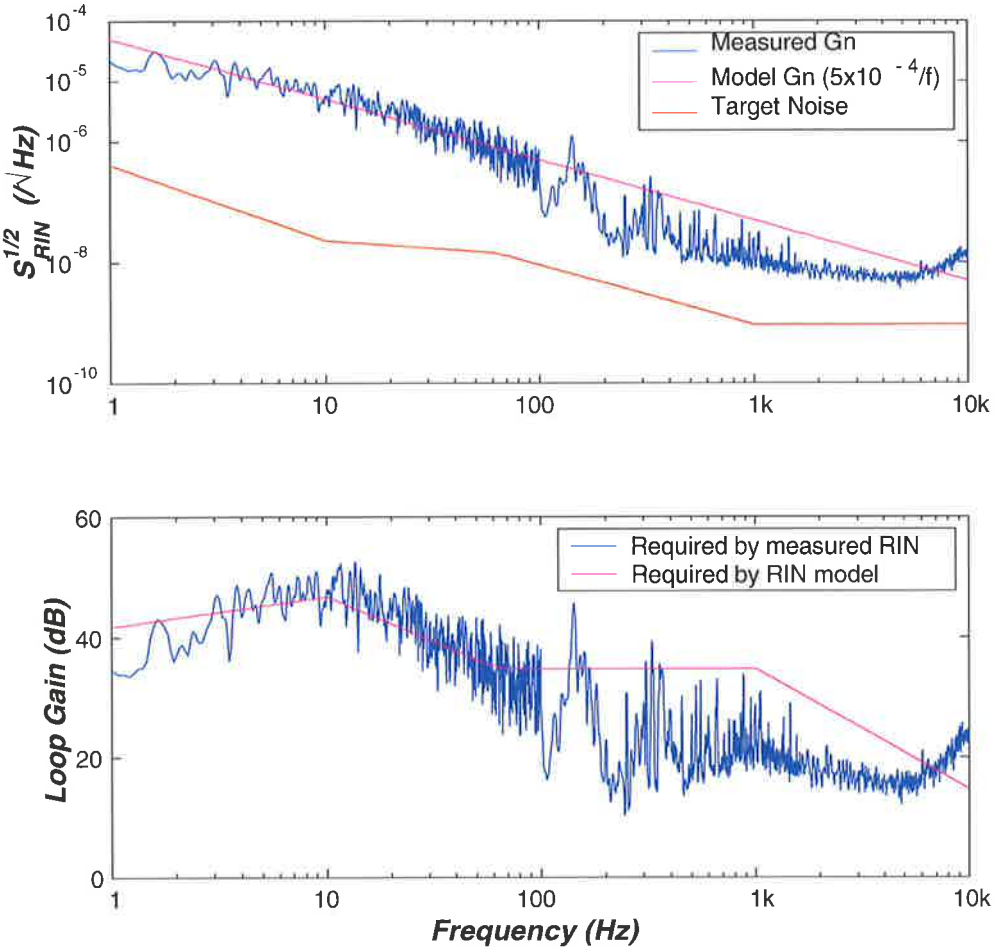

Figure 4.65: The top graph shows the open loop intensity noise, G_n (both the measured curve and its nominal model) and compares it to the target noise. The required loop gain is shown in the bottom graph with results for the measured data and the simplified model presented.

in the high frequency region is limited by the necessity to keep the loop stable, so only modest improvements in the performance could be made by further tailoring of this control loop. If the performance of this loop proves insufficient then it will be necessary to introduce an intensity actuator with a higher bandwidth. The two available options are to use the laser's noise-eater input or to add an external an intensity modulating element such as an EOM or AOM.

The compensator circuit is shown in figure 4.68. The variable resistor, R_r, allows the setpoint of the servo loop to be adjusted. Rr was adjusted so that the detector photocurrent was stabilised to 10 mA. Additional details of the construction of the

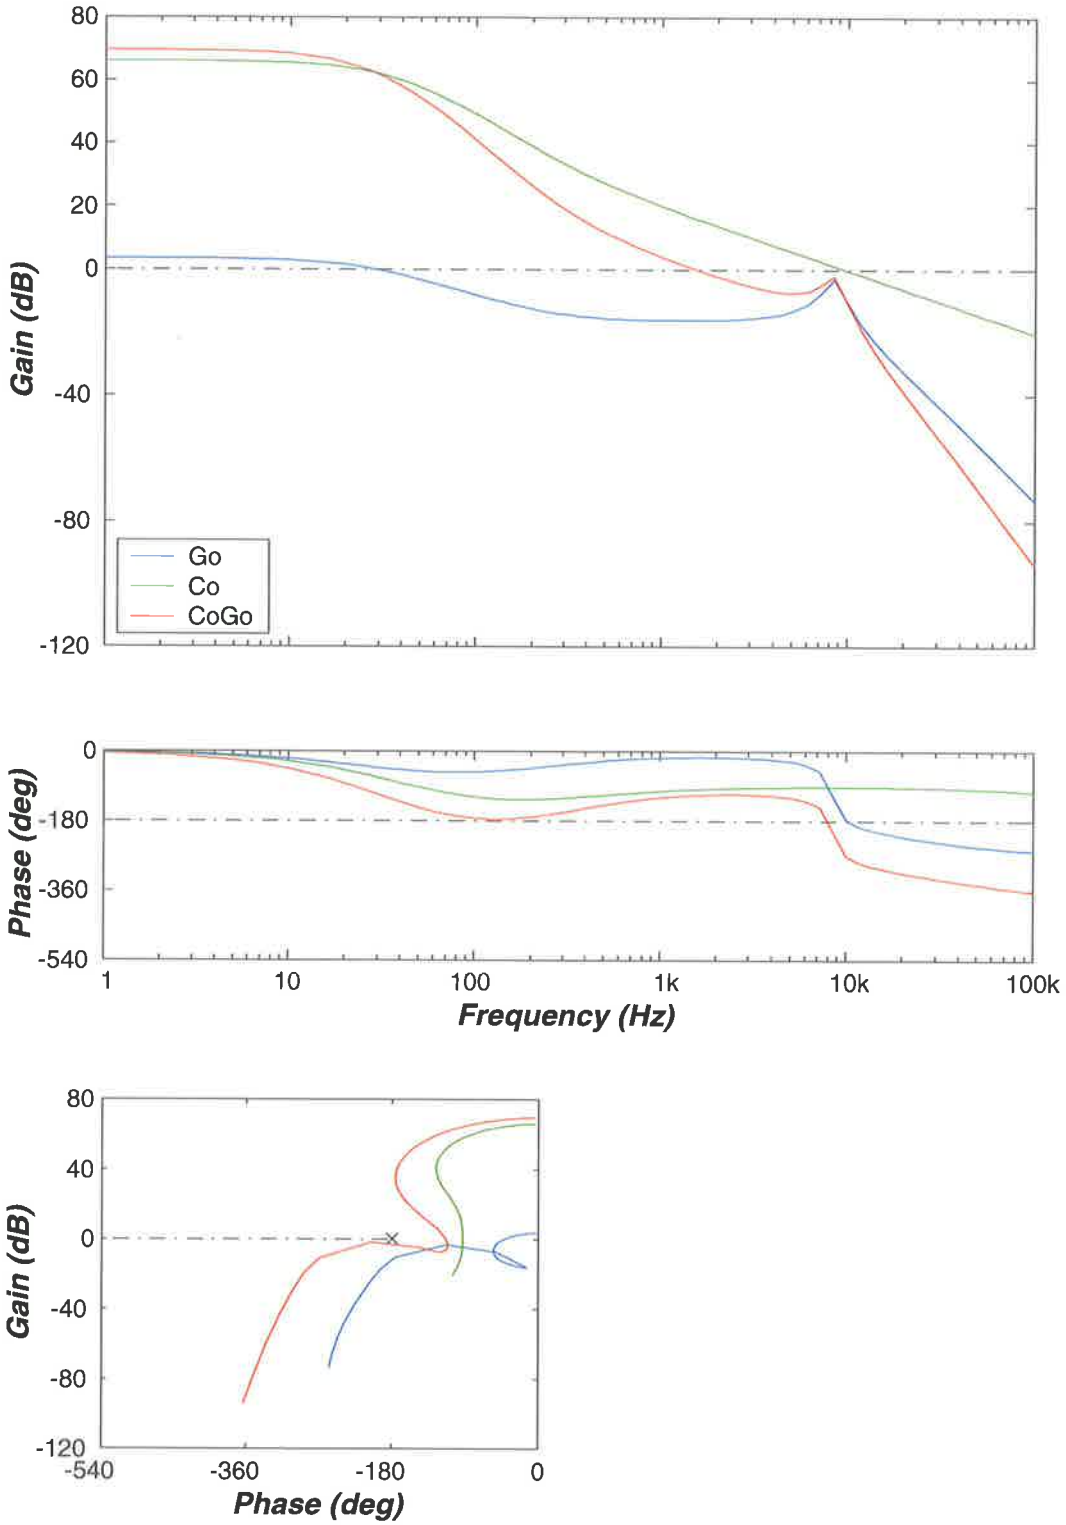

Figure 4.66: Nominal transfer functions for the plant (including the sensor), G_0, compensator, C_0, and the open loop gain (C_0G_0) for the intensity stabilisation system. The Nichols plot shows that the system should be closed loop stable.

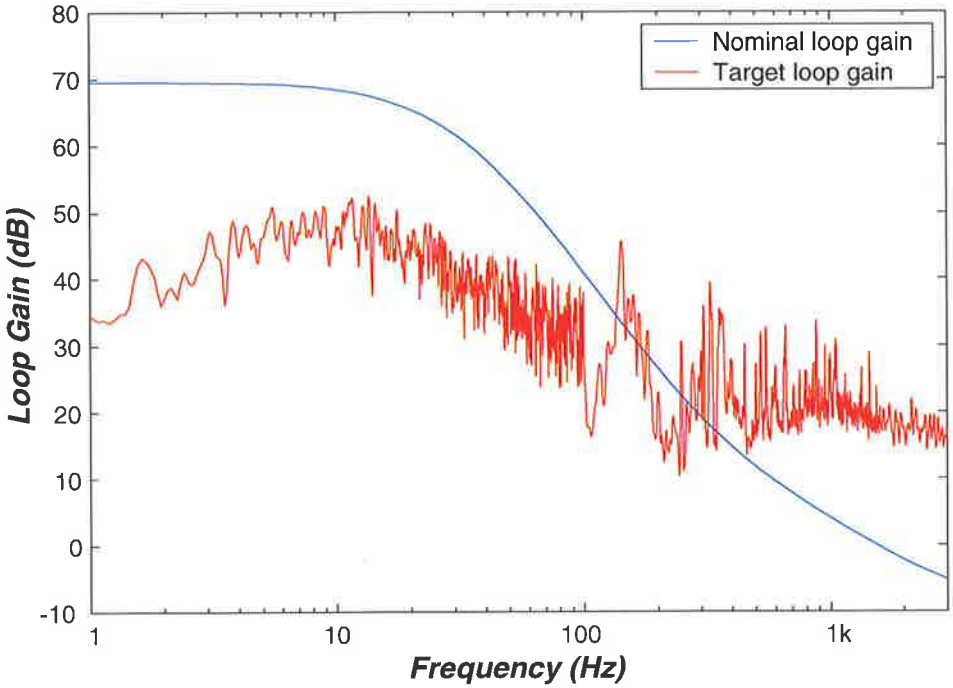

Figure 4.67: Nominal and target curves of the intensity stabilisation servo's loop gain.

compensator can be found in appendix E.

The presence of the non-linearity complicates the operation of the loop, as the resonance becomes more pronounced when larger control signals are required. These large control signals are required when the servo is switched on and the compensator establishes control over the laser intensity. Once the intensity has been stabilised, the amplitude at the compensator output is reduced markedly. The amplitude of the resonance is thus considerably lower in normal operation than it is during the start up phase.

The large resonance requires the use of a relatively low value for the compensator gain during start-up. Once the loop is running, the gain can be increased to further reduce the system noise. The sliding gain control necessary to carry out this operation is controlled by the variable resistor, R_m. This resistor enables the gain of the servo to be smoothly changed from zero up to its full value. Experiments showed that the use of this variable gain servo enabled us to achieve an ultimate gain five times higher than the conservative gain value that ensured unconditional

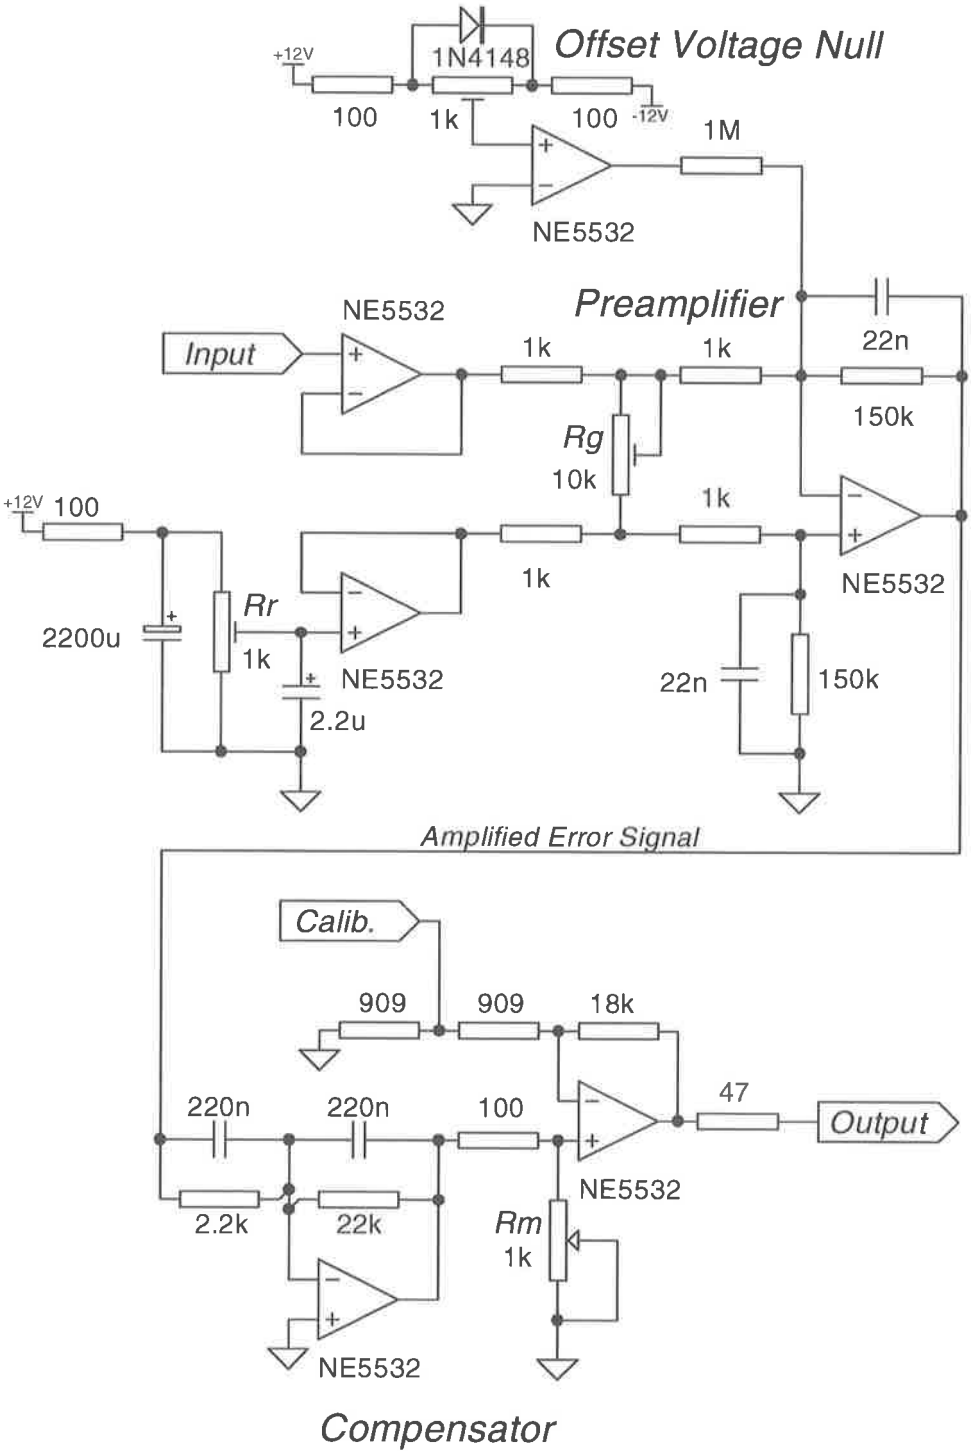

Figure 4.68: Schematic of the intensity stabilisation compensator.

system stability. The disadvantage of the high gain configuration is that it allows the intensity stabiliser to lose "lock" in the presence of unusually large noise transients.

The compensator was tuned so that it was stable for extended periods. The resulting system rarely lost lock under normal laboratory situations, but when it did the situation was easily remedied by briefly lowering the gain. Figure 4.69 shows the measured compensator transfer function C while at its full gain setting along with the nominal compensator transfer function C_0. As can be seen in the figure the tuned compensator had a higher gain than the nominal compensator (no doubt due to the use of a conservative model in the design), but otherwise the two curves are in excellent agreement.

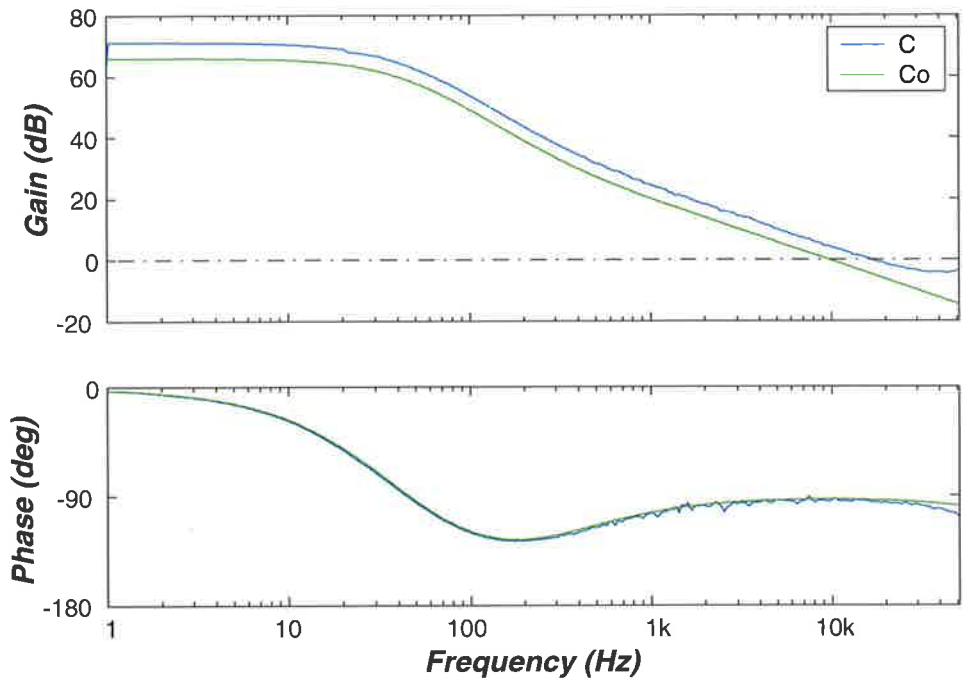

Figure 4.69: Measured transfer function of the intensity stabilisation compensator compared with the nominal design compensator. C is the measured compensator transfer function and C_0 is its nominal form.

4.4.4 Intensity Stabiliser Performance

The laser intensity noise was measured in the closed loop system and the result is shown in figure 4.70. The figure also shows the equivalent frequency noise that

would be expected in the test cavity readout due to the presence of the intensity noise. As can be seen the closed loop noise closely matches the level predicted by the nominal sensitivity. As expected, the noise level leaves the desired green region at frequencies above 300 Hz, a result consistent with that shown in figure 4.67. The noise remains less than the expected level of thermoelastic noise for all frequencies below 1 kHz.

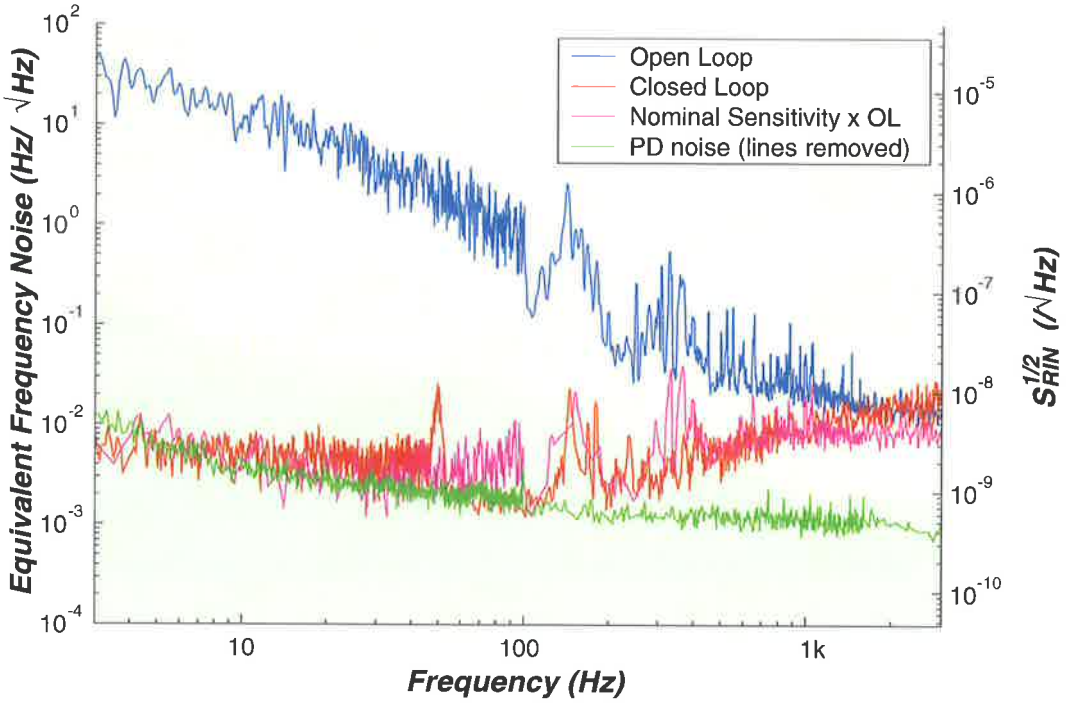

Figure 4.70: Comparison of the RIN of the open and closed loop systems. The predicted closed loop noise is also shown for comparison. The photodetector nosie floor is also shown, with 50 Hz mains harmonics removed for clarity. The right y-axis shows the RIN explicitly; that on the left shows its equivalent frequency noise.

The noise achieved at the in-loop sensor is not necessarily a good indicator of the real intensity noise suppression achieved however. Measurements of out-of-loop intensity noise usually yield intensity spectra considerably above the in-loop levels [109,111,112]. A number of different noise sources can contaminate the in-loop sensor signal and these will be impressed onto the laser beam via the servo loop. The most important of these noise sources are;

1. Beam jitter: Movement of the laser beam across the photodetector can lead to

variations in apparent power due to clipping or because of spatial variations in the responsivity of the detector [109].

2. Detector temperature variation: Photodetector responsivity is a function of temperature, so any change in photodetector temperature will appear as a power variation [109, 110].

3. Voltage reference noise: Noise in the reference with which the detected power is compared is indistinguishable from laser intensity noise [110]. Any noise arising from this reference will therefore be impressed onto the laser intensity.

4. Photodetector noise: Photodiodes can exhibit excess electronic noise at low frequencies, particularly if exposed to air after their hermetic seals have been broken (which is required if they are to be used in vacuum) [105].

The in-loop and out-of-loop intensity noise levels were compared by placing a detector after the reference cavity and comparing its output with that of the in-loop sensor. In this location the out-of-loop detector measured the intensity noise seen at the test cavity. The two detectors used were of identical design and were tuned to be as similar as possible. Both detectors were set to measure 10 mA of photocurrent. The diameter of the laser at the out-of-loop detector was 600 μm on the 1 mm diameter detector. The resulting spectra are shown in figure 4.71.

A number of issues could be responsible for the poor low frequency performance of the servo. Several groups have found unexpectedly poor low frequency intensity stabilisation due to a variety of reasons. Indeed, [112] reports that an unknown noise source prevents *any* out-of-loop noise reduction by their intensity stabilisation system below 5 Hz. Several possible noise sources are examined in [113] [114] and [110], but the achievable noise levels are still unexpectedly high.

In our case, voltage reference noise is likely to be a contributor to the problem, as the compensator circuit uses a simple 78L12 regulator that is filtered with a 0.8 Hz, single-pole low pass filter. Others have used precision voltage references and two-pole filters to reduce the voltage reference noise [109, 110, 112]. A future version of our system will use a temperature stabilised, battery-powered voltage reference to improve this noise signal. As has been shown in [109], the reduction of detector

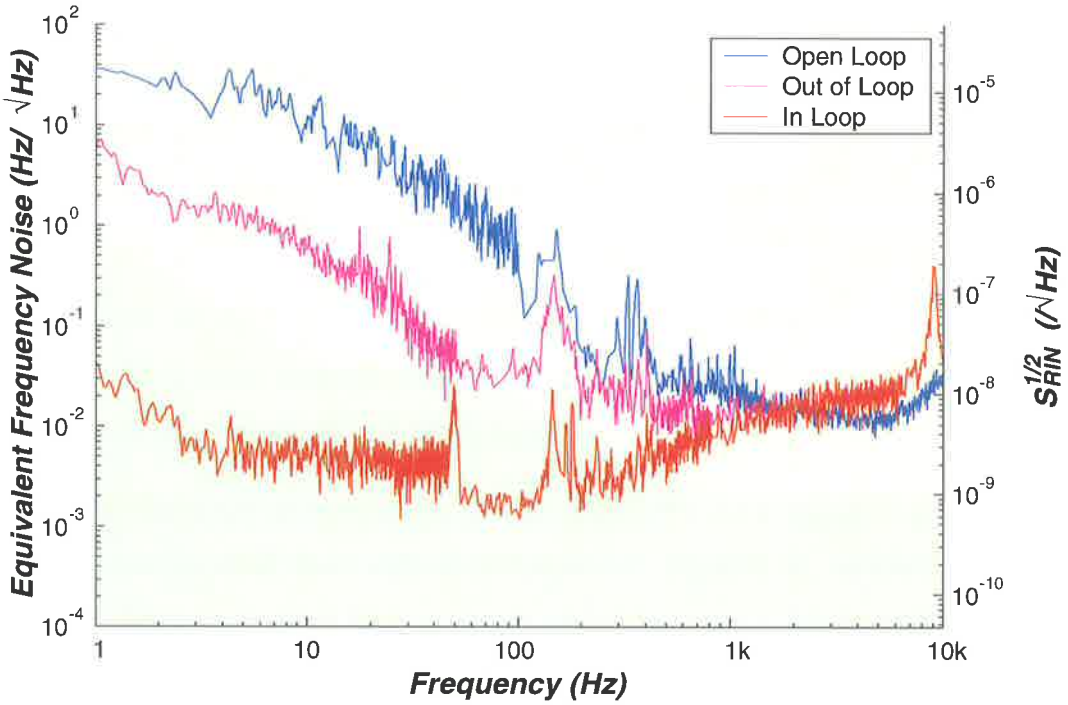

Figure 4.71: Comparison of the RIN spectra for the closed loop intensity noise measured by the in-loop and out of loop sensors. Also shown is the open loop noise. The background colours indicate the required sensitivity regions.

temperature fluctuations is also likely to improve the low frequency excess noise by a significant factor. Placing the system in vacuum is thus expected to improve the temperature stability of the detector by preventing convective cooling. If necessary, the use of a temperature stabilised photodetector could also be considered [110].

As has been discussed in chapter 3, we also expect the vacuum environment to have significantly lower levels of seismic and acoustic excitation of the optical components in the system. While this should lead to a small reduction of the jitter of the laser beam across both the in-loop and out-of-loop detectors, beam jitter is likely to remain a concern because a significant portion of the optics will remain outside the vacuum envelope. If the photodetector improvements discussed above do not adequately reduce the low frequency noise then it may be necessary to use the photodetector after the reference cavity to provide the intensity stabilisation error signal [115]. The mode cleaning action of the reference cavity would then reduce the beam jitter. Preliminary experiments indicate that this approach is feasible for

our system, but it has not yet been adopted because of the potential complications caused by the reference cavity coupling laser frequency noise into intensity noise.

The broad noise feature around 150 Hz is also of concern. This is probably due to a mechanical resonance in the photodiode mount, the amplitude of which should be reduced once the system is in a vacuum.

As was discussed above, the high frequency performance of the intensity stabilisation could only be improved by increasing the bandwidth of the actuator. While this could be done by using the noise eater input of the laser, such a system would prevent use of the built-in noise eater. A simpler upgrade would be to add an external intensity modulating element, such as an AOM [110]. This modulator would then be driven in a split feedback system, such as that described in the section on laser frequency stabilisation.

There are thus a number of potential improvements that can be made to the intensity stabilisation servo loop. Which of them, if any, need to be implemented will be determined by the results of measurements made after the system has been placed into vacuum. With these improvements it is clear that the intensity stabilisation should prove adequate.

4.4.5 Intensity/Frequency Cross-coupling

The cross-coupling between the laser's intensity and frequency noise was examined to ensure that stabilisation of one of the noise sources did not introduce additional noise into the other. If the cross-coupling from intensity to frequency were found to be significant any intensity modulation signal applied to the laser would appear on the laser frequency multiplied by the cross-coupling transfer function. Conversely, if frequency to intensity coupling were significant then frequency stabilisation would lead to an increase in measured intensity noise. In the presence of significant interaction between the two loops a full multiple-input multiple-output (MIMO) control system design would need to be considered.

4.4.5.1 Coupling of an Intensity Modulation Input to Laser Frequency

NPRO lasers are known to exhibit coupling from intensity to frequency [94] because changes in the pump power are coupled to changes in crystal length via thermal absorption . The resulting optical path length alters the resonant frequency of the crystal and hence of the laser radiation. This coupling is sufficiently strong that stabilisation of laser intensity can lead to approximately 20 dB of frequency noise suppression in an NPRO [114].

To determine the effect of the measured coupling we consider an intensity modulation signal X_i and calculate the resulting frequency noise Y_f. Consider the system shown in figure 4.72, where G_{if} represents the coupling from the intensity modulation input to the laser frequency.

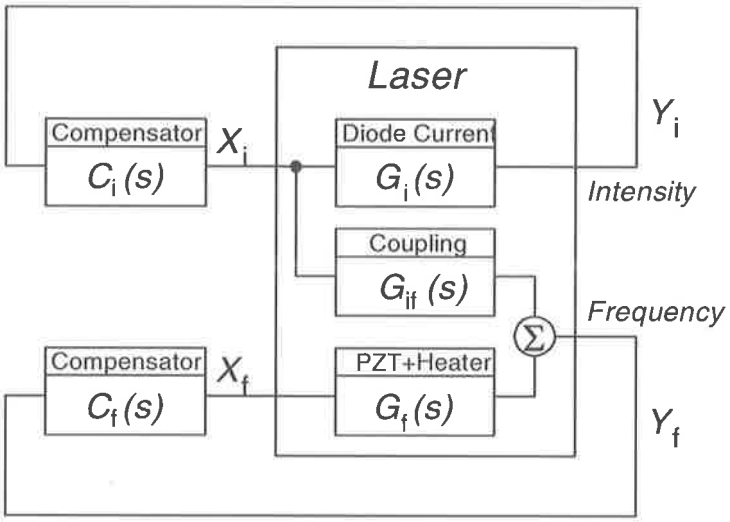

Figure 4.72: Block Diagram of the intensity modulation to laser frequency coupling mechanism.

We know the closed loop intensity noise, Y_i, is as shown in figure 4.70 above. The noise at the intensity modulation port, X_i, can thus be found by multiplying this noise by C_i, the transfer function of the intensity compensator. This noise can be multiplied by the cross coupling transfer function to find the resulting (open loop) frequency noise. Finally, multiplying by the frequency stabilisation loop sensitivity yields the closed loop frequency noise caused by cross coupling, as shown below.

$$\begin{aligned} X_i(s) &= C_i(s)Y_i(s) \\ Y_f(s)|_{OL} &= G_{if}(s)X_i(s) \\ &= G_{if}(s)C_i(s)Y_i(s) \\ Y_f(s)|_{CL} &= S_f \times Y_f(s)|_{OL} \\ &= S_f G_{if}(s)C_i(s)Y_i(s) \end{aligned}$$

were S_f is the sensitivity of the frequency stabilisation loop.

The coupling was measured by modulating the laser intensity and measuring the resultant signal with the PDH frequency discriminator. The results of this experiment are shown in figure 4.73. The graph shows that the coupling is quite strong at low frequencies but drops off beyond about 10 Hz. This low corner frequency is consistent with a coupling mechanism having a thermal origin. It should be noted that the transfer function beyond 500 Hz is of doubtful quality because of the low signal amplitudes observed.

The calculated closed loop frequency noise caused by cross coupling is shown in figure 4.74. As can be seen this noise source satisfies the sensitivity requirement of the experiment.

4.4.5.2 Coupling of Frequency Modulation Inputs to Laser Intensity

A measurement of the reverse coupling path was also made to see whether modulation of the laser frequency would introduce intensity noise into the laser output, an effect predicted in [94]. The laser frequency was modulated by changing both the laser PZT and laser temperature and in each case there was no discernible correlation between the modulation signal and laser intensity. We can conclude that the frequency modulation exhibits negligible coupling to laser intensity.

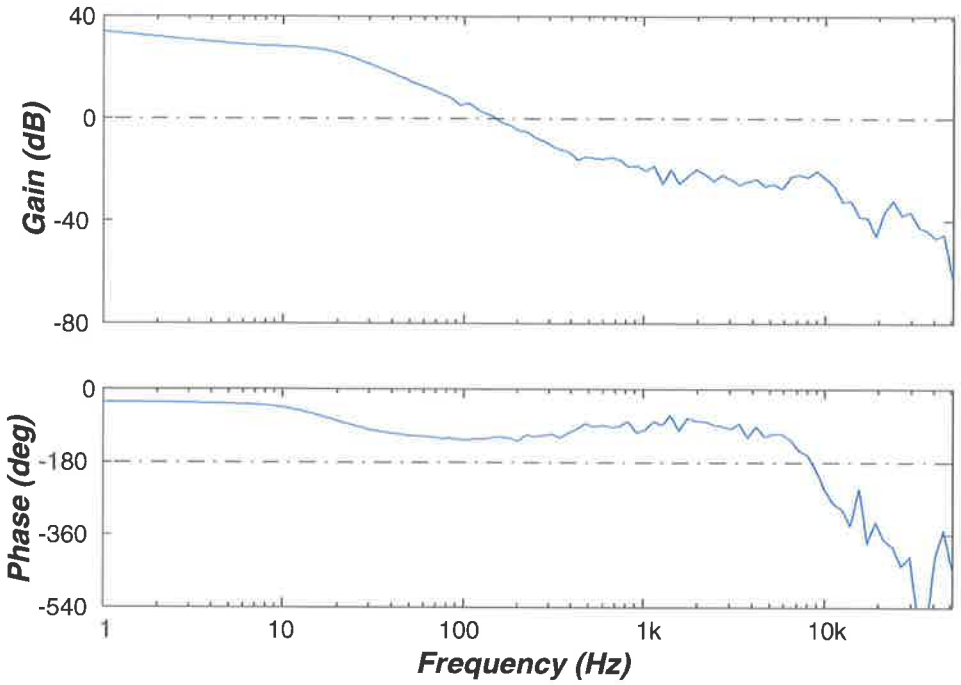

Figure 4.73: Transfer function of coupling from the intensity modulation input to the closed loop PDH frequency discriminator output ($S_f G_{if}$).

4.5 Conclusion

This chapter described the optical design of the experiment, the servo systems required to stabilise the frequency and intensity of the laser and the system used to read the length noise of the test cavity.

Frequency stabilisation of the NPRO laser was discussed in section 4.2. The laser was successfully locked to a ring cavity having a finesse of 4100 and a length of 0.53 m. This locking system uses a Pound-Drever-Hall frequency discriminator and applies feedback to the NPRO PZT and to the temperature of the laser crystal. We showed that the stabilisation could adequately reduce laser frequency noise for all frequencies below 500 Hz. This performance does not extend to higher frequencies, probably because of the presence of excess laser intensity noise at the PDH modulation frequency [79]. Furthermore, it may be necessary to improve the bandwidth of the feedback by incorporating a wideband electrooptic frequency modulator. Extension of the bandwidth would allow more servo gain in the region near 1 kHz and

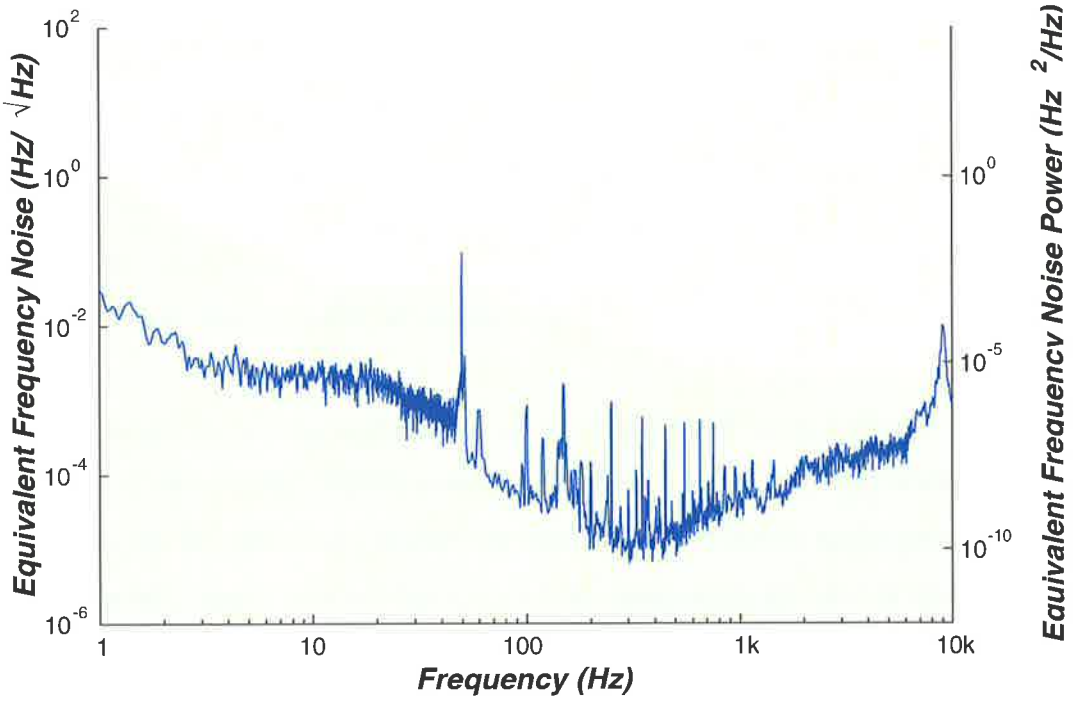

Figure 4.74: Frequency noise caused by crosscoupling from the intensity modulation port. The background colours indicate the required sensitivity regions.

would therefore improve the system noise at the high end of the frequency band.

In section 4.3 we presented a design for the test cavity and described a servo system used to lock the frequency of the reference cavity to that of the test cavity. This servo uses a fringe side locking sensor to infer the frequency difference between the two cavity resonances and feedback is applied to the reference cavity PZT with a bandwidth of 500 Hz. The FSL readout scheme was shown to be sufficiently quiet to measure test cavity thermoelastic noise provided that the intensity of the laser was stabilised appropriately. Some practical difficulties arose in the implementation of the test cavity monitoring scheme because of the cavity's high free spectral range and its lack of tunability. These difficulties and systems to overcome them will be the focus of chapter 6.

The intensity stabilisation of the laser was described in section 4.4. The laser light is sampled before the reference cavity and fed back to the laser pump source, resulting in a substantial reduction in laser intensity. The noise reduction at the in-loop sensor was shown to meet the sensitivity requirement of the experiment.

However, the intensity noise measured after the reference cavity was shown to be noisier than that at the in-loop sensor, probably due to laser jitter, thermal effects in the two detectors and electronic noise in the detector's voltage references. Nevertheless, the intensity performance at the out-of-loop sensor was shown to meet the target sensitivity for frequencies between 20 Hz and about 700 Hz. We expect that improving the voltage reference electronic noise and the move into vacuum in the next phase of the project will result in a substantial improvement in the out-of-loop intensity measurement.

The expected noise levels produced by the various servo systems in the system are shown in figure 4.75 . The total noise is dominated by laser intensity noise at low frequencies, which should be significantly reduced by improving the voltage reference source and moving the experiment into vacuum in the next phase. Indeed, it should be possible to measure thermoelastic noise at frequencies even lower than our target frequency of 10 Hz. Measurement of thermoelastic noise at these frequencies if probably only possible using our monolithic test cavity technique, as it is much less sensitive to environmental noise than alternatives. At frequencies above 400 Hz, the system noise is dominated by the limited bandwidth of the reference cavity locking servo. As discussed above this bandwidth can be extended by adding a broadband electrooptic modulator to the reference cavity servo. Thus we expect that our system should be able to measure sapphire thermoelastic up to, and possibly slightly beyond, 1 kHz.

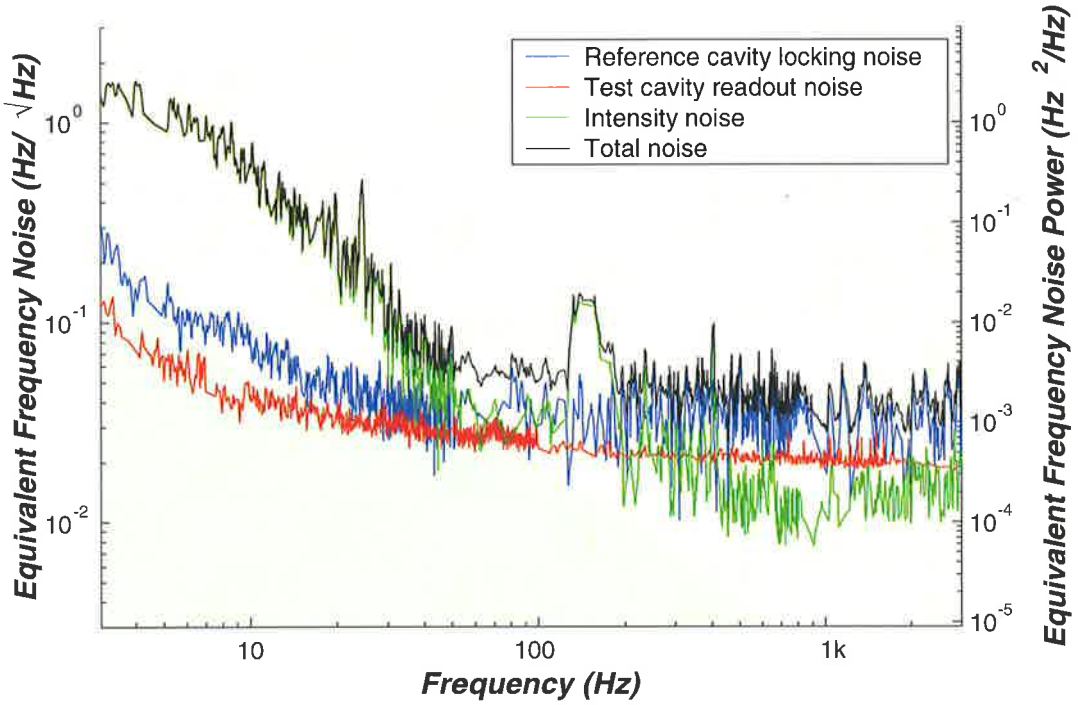

Figure 4.75: A summary of noise arising in the three main subsystems. For clarity, 50 Hz mains harmonics have been removed from these curves. The coloured regions in the background indicate the desired sensitivity regions.

Chapter 5

Sapphire Test Cavity Results

This chapter describes several preliminary experiments carried out using a sapphire test cavity. It opens with a measurement of the sapphire cavity finesse and shows that the result is close to the desired value of 1000 (see section 4.3.1). Difficulties in tuning the probe laser over the entire FSR of the test cavity made mode matching the test cavity difficult. These difficulties are discussed in detail and the mode matching that was achieved is presented. The lack of easy tunability also lead to significant problems with locking the reference cavity to the probe laser frequency. These problems are discussed in detail. The chapter concludes with preliminary noise measurements made by measuring the FSL detector output as the probe laser drifted through test cavity resonance.

5.1 Sapphire Cavity Finesse

The finesse of the monolithic sapphire cavity was measured using the technique described in section 4.2.2.1. The reference cavity was removed from the system during this measurement so that the laser could be swept rapidly through the test cavity resonance without concern for maintaining reference cavity lock. Figure 5.1 shows the power transmitted through the cavity during a typical sweep through the sapphire test cavity resonance.

The width of the resonance was calculated using the spacing between the two 17 MHz modulation sidebands. The calculated test cavity FSR of 4.27 GHz was

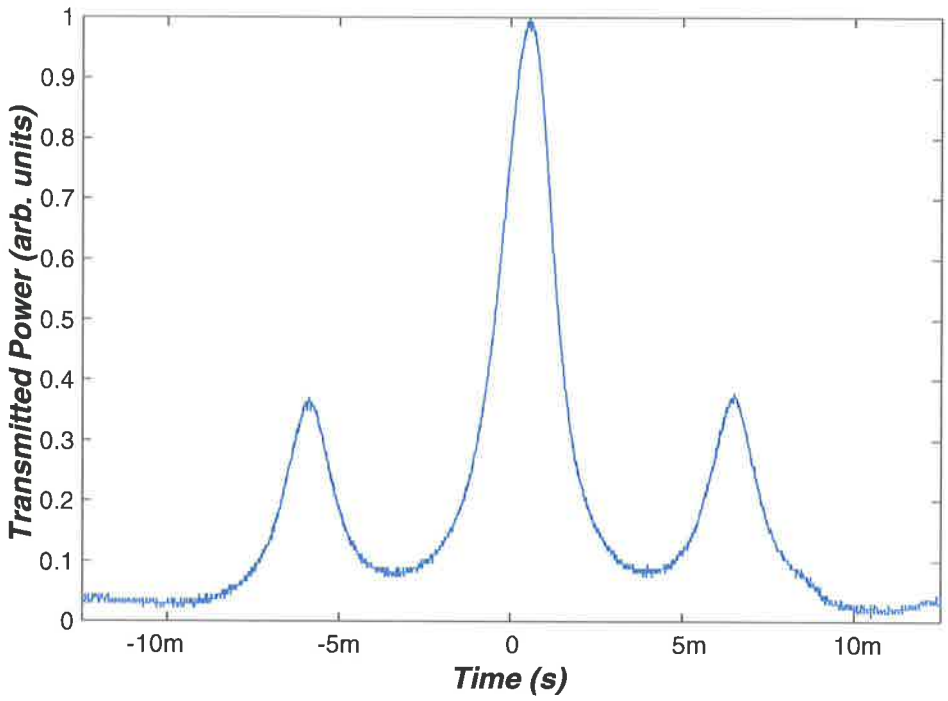

Figure 5.1: Transmitted power as the laser was swept through a resonance of the sapphire test cavity. The figure shows the carrier and the two modulation sidebands sweeping through the cavity resonance.

then divided by this linewidth to determine the cavity finesse. Nine separate measurements were averaged, resulting in a cavity finesse of 920 with a standard error of 50. This result is sufficiently close to the design finesse of 1000 that the cavity should perform acceptably.

5.2 Mode Matching to the Test Cavity

Mode matching to the test cavity was done using the optical system described in section 4.3.1.2. Two steering mirrors were used to position the beam on the test cavity input coupler and a lens was used to produce the required beam waist. A movable stage was used to vary the optical path length so that the waist could be located accurately at the front surface of the test cavity.

While the reference cavity was removed during the mode matching procedure so that it did not need to be locked to the laser during this time, its removal did

not damage the quality of the mode matching because the laser beam was well mode matched to the reference cavity (see section 4.2.2.2). The reference cavity was replaced after the test cavity modematching had been completed and the matching was found to be unchanged by its presence.

The conventional mode matching technique described in section 4.2.2.2, requires that the frequency difference between the laser and the cavity resonance be rapidly swept by at least one cavity free spectral range. This allows examination of the higher order transverse modes during the sweep and consequently enables their minimisation *via* adjustment of the mode-matching optics. Unfortunately the frequency difference could not be rapidly scanned over an entire cavity FSR in this case because the dynamic range of the laser PZT actuator was only 200 MHz [95], much less than the 4.27 GHz FSR of the cavity.

Thus, it was necessary to use thermal tuning of the laser gain medium, which was implemented by applying a triangular waveform to the laser crystal's temperature controller. The sweep rate achievable by this mechanism was very slow, particularly given the very large range of the sweep. As a result the mode matching process proved painstaking as a complete temperature sweep was required for each adjustment of the mode matching optics. Consequently, the mode matching process took many hours to complete, even to the modest level shown below.

Figure 5.2 shows the power transmitted through the test cavity as the laser frequency was varied. The regular spacing of the higher order modes indicates that both odd and even order modes are present due to misalignment of the beam and errors in waist size and position respectively. However, perfect mode matching is not required as long as there is sufficient coupling to the TEM_{00} mode that an error signal can be obtained with the FSL detector. The mode matching shown in figure 5.2 was adequate for preliminary experiments.

5.3 Lock Acquisition

Severe difficulties were encountered during tuning of the compensator used to lock the reference cavity to the test cavity. To set up this compensator, the frequency of

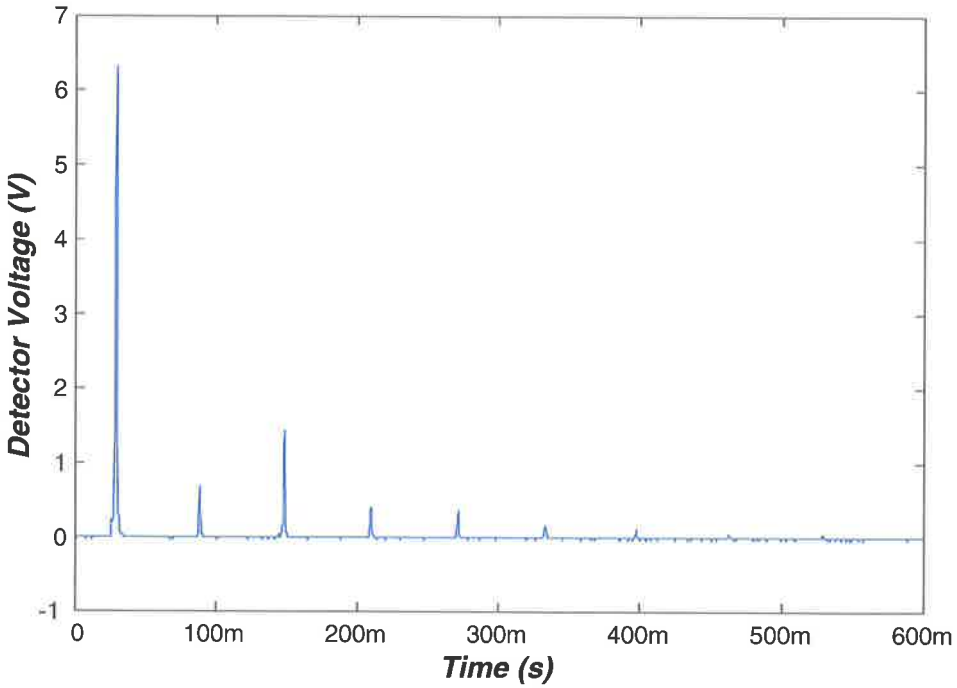

Figure 5.2: Power transmitted through the test cavity after mode matching. The large peak on the left of the figure corresponds to the cavity TEM_{oo} mode, while the smaller peaks are produced by coupling of the probe beam to higher order test cavity TEM modes.

the probe laser was repeatedly driven through the test cavity resonance. Since the test cavity has a fixed length, this was accomplished by changing of the reference cavity length. This tuning of the reference cavity length was done thermally, because the dynamic range of the reference cavity PZT was only 40 MHz, much less than the 4.27 GHz required to tune through a test cavity FSR. Heating of the cavity was provided by four small lamp bulbs that were placed in contact with the reference cavity spacer and driven with a high current amplifier. Details of the amplifier design can be found in appendix E

The heater could vary the frequency of the reference cavity resonance until the probe laser came into resonance with the test cavity. At this point the servo would lock the reference cavity to the test cavity resonance by driving the reference cavity PZT actuator. In practice however, it was found that thermal inertia in the heater and reference cavity quickly exhausted the actuator dynamic range, causing lock to be lost. As a result of these difficulties, only transitory locking of the two cavities

was obtained during the work described in this thesis.

Thus, thermal tuning would need to be very slow so that thermal overshoot could be minimised. This corresponds to a low heater power, so low in fact that the process was found to be impractical for use in manual lock acquisition. An automatic lock-acquisition system that it is expected to overcome these difficulties by using a microcontroller will be discussed in chapter 6

5.4 Preliminary Noise Spectrum

A series of spectra were recorded as the probe laser drifted through resonance with the test cavity TEM_{00} mode. Such resonances lasted of the order of a few seconds until lock was lost. Because the system was not accurately locked, the spectra were contaminated by the drift processes. In addition the non-constant power on the detector meant that the FSL sensor response changed as the data was acquired; the root mean square of the FSL gain was therefore used to estimate the conversion from discriminator voltage to length noise. Despite these difficulties, the spectra provided a useful diagnostic and enabled the determination of an upper limit for the system noise. The spectrum shown in figure 5.3 was the best obtained and serves as our estimate of the system sensitivity. During this measurement the system was still in air and was thus subject to acoustic excitation and significant degradation of Q's due to gas damping. Despite the expected degradation by these effects, the measured spectrum was within a factor of ten of the expected thermoelastic noise spectrum over the measured frequency range.

As only a brief time was spent drifting through each resonance, the spectra were recorded very quickly. This short sampling time had a number of unavoidable consequences:

- The small available collection time meant that measurement of noise at low frequency was not possible.
- The spectrum analyser needed to be set with poor frequency resolution (a factor of four worse than all other spectra presented in this thesis) to speed data

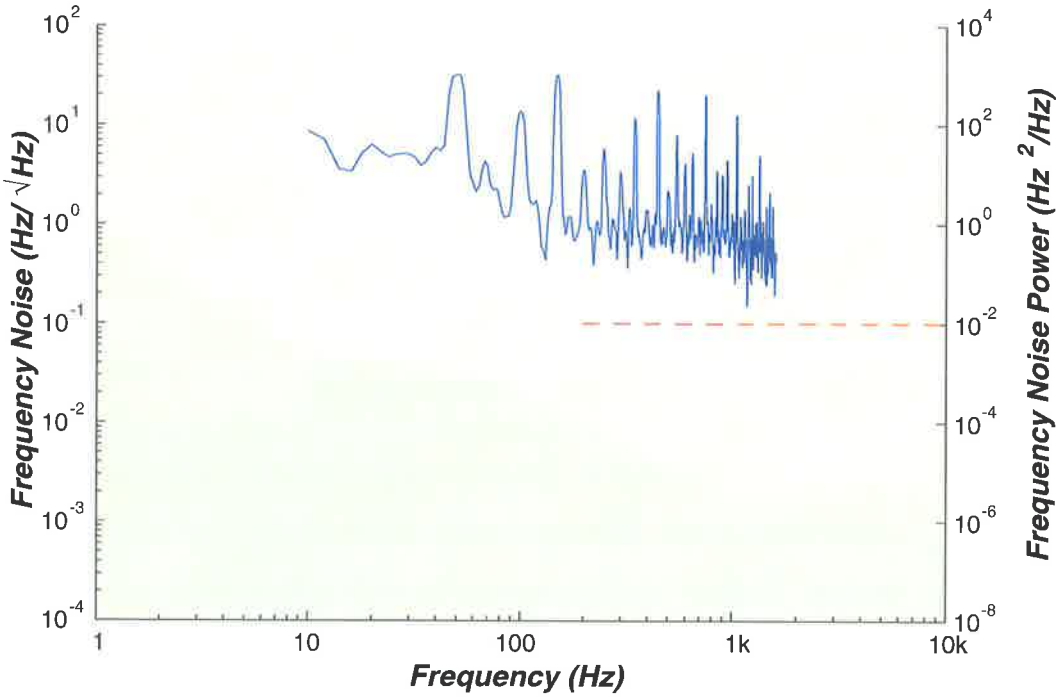

Figure 5.3: Spectrum of frequency noise taken during a drift through resonance. The red dashed line indicates the expected refractive index noise floor.

acquisition. This led to increased contamination of the measured spectrum with mains harmonics. The figure should be compared with 4.48 to see the significance of this effect.

- The measurement could only make limited use of signal averaging to reduce sensor noise. Fifteen independent spectra were averaged to form figure 5.3, compared to 100 or more used in other results. Signal averaging improves noise by a factor of $\sqrt{N}$ where N is the number of averages taken. Because the number of samples was much smaller in this case there was a significantly larger contribution from sensor noise.

When the locking system described in chapter 6 has been implemented, it will be possible to take spectra over the complete frequency range with full instrumental resolution. Sufficient signal averaging to reach the instrumental noise floor will also be possible. In this configuration the narrow band mains harmonics will be resolved and the thermoelastic noise should be visible between them. The fact that the

measurement shown here lies so close to the desired sensitivity gives us confidence that the system will indeed reach the desired sensitivity once the new systems are in place.

5.5 Conclusion

In this chapter we have presented some preliminary measurements conducted with a sapphire test cavity. The cavity finesse was found to be 920, very close to the design value of 1000. The large FSR of the test cavity made mode matching of the probe laser beam to the test cavity difficult, but adequate mode matching was achieved using thermal tuning of the laser. The large FSR of the test cavity also made lock acquisition between the reference and test cavities difficult and required thermal tuning of the reference cavity. This technique proved impractical for manual operation. Despite the lack of a reliable locking mechanism, an initial measurements of system noise was obtained and discussed. The noise level measured was within a factor of ten of the expected thermoelastic noise level, a result that is probably consistent with the degradation that would be expected due to acoustic excitation and gas damping of the system.

Chapter 6

The Next Phase

6.1 Introduction

In chapter 4 we saw that the short length of the test cavity combined with its monolithic construction led to two major experimental difficulties:

- Difficulty in mode matching the laser beam to the cavity;
- Difficulty in stably locking the reference cavity to the test cavity.

This chapter examines methods that could be used to overcome these difficulties in the next phase of the project. It begins with an examination of the mode matching problem. We will see that the use of a supplemental locking technique during mode-matching will allow the probe laser beam to be rapidly matched to the mode of the test cavity.

In section 6.3, we present a method for improving the locking of the reference cavity to the test cavity using carefully controlled thermal tuning of the reference cavity. Thermal modelling of the reference cavity is first presented and then an automatic lock acquisition system is discussed.

6.2 Mode matching

In section 5.2 we saw that aligning the probe laser beam with the test cavity was unrealistically time consuming. This was because it proved impossible to quickly

sweep the laser frequency over a test cavity FSR as it required thermal tuning of the laser. Slow sweeping of the frequency is unsuitable for mode matching because the amplitude of the higher order TEM modes must be examined as the mode matching optics are adjusted.

As we cannot rapidly sweep the laser frequency through the FSR of the test cavity we must abandon this approach and instead choose to maximise the transmitted power through the appropriate test cavity mode. We first approximately align the laser beam to the test cavity by overlapping the spot reflected from the test cavity with a HeNe laser spot transmitted backwards through the test cavity. This technique quickly aligns the probe laser beam to the cavity sufficiently well that a resonance can be established in the cavity.

Once a resonance has been achieved, the laser beam can be crudely locked to the test cavity by placing a frequency dither on the probe laser frequency and demodulating the effects of this dither on the test cavity transmitted power [116]. A schematic of this scheme is shown in figure 6.1. Note that the system requires no additional optical components and uses the same photodetector as the test cavity FSL system. The laser is assumed to be locked to the reference cavity, as described in section 4.2, so those locking components are omitted from the diagram.

To use this technique we impress a sinusoidal frequency modulation onto the probe laser frequency by placing a small sinusoidal signal on the reference cavity PZT. This frequency modulation produces an intensity variation of the power transmitted through the test cavity. Demodulation of the resultant detected power with a lock-in amplifier yields an error signal that has a different sign on the two sides of the resonance. This error signal can be amplified and fed back to the reference cavity PZT to lock the frequency of the reference cavity and thus that of the laser to the resonant frequency of the test cavity.

This stabilisation method will lock to the peak of the test cavity resonance irrespective of the transmitted power because it is a modulation technique. This property is important because we have little idea about the amount of power that is likely to be present in the TEM_{00} mode of the test cavity during the early stages of mode-matching. While this dither locking scheme is useful during the mode

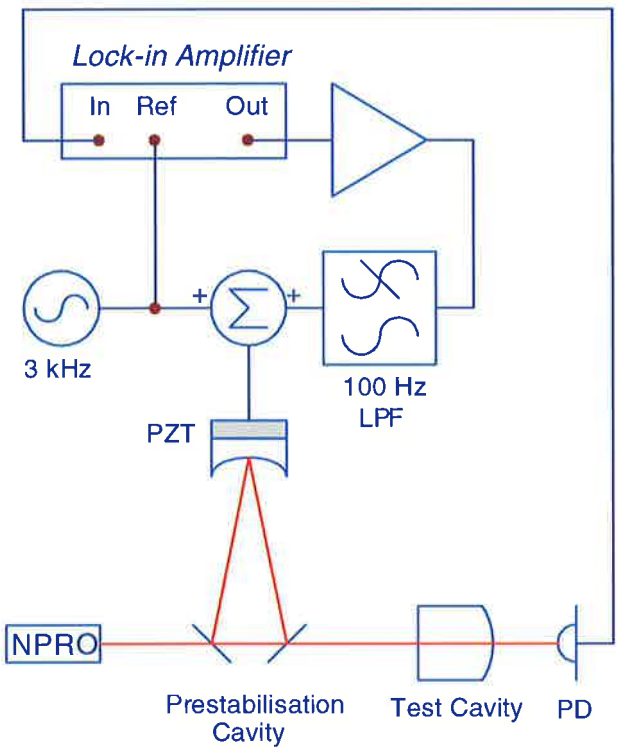

Figure 6.1: Locking system used during optimisation of test cavity mode matching. Details of the laser locking omitted from the diagram.

matching procedure, the system must be switched back to the FSL system when good noise performance is required because the modulation required by the dither would contaminate the test cavity length noise readout in the region of the modulation frequency.

During the initial stages of mode matching there will be little power in the desired TEM_{00} mode, but this power increases as the mode-matching is improved. As the power increases the sensor gain will also increase. However, the gain change will not cause stability problems because the control loop can be very conservative, as it only need keep the system in resonance (rather than achieve low noise performance).

Once the resonance has been acquired, the mode-matching optics can be adjusted to maximise the power in the mode. Maximum transmitted power will occur when the probe beam is optimally mode matched.

It should be recognised that achieving perfect mode matching using this technique will remain more difficult than the conventional technique, because the match-

ing needs to be done on a trial-and-error basis. This is because this technique does not provide the useful diagnostic information about what needs to be adjusted as can be obtained when sweeping through transverse modes (see section 4.2.2.2). Nevertheless, the moderate level of mode matching shown in section 4.3.1.2 should be readily achievable and will be adequate for the experiment.

The modulation frequency must be placed outside the 500 Hz bandwidth of the reference cavity servo loop, so that the loop does not attempt to remove the dither. A modulation frequency of 3 kHz was chosen for this system. We must also ensure that the frequency deviation introduced at the reference cavity is less than the linewidth of the test cavity so that the dither does not drive the error signal away from its linear range. The linewidth of the test cavity will be equal to the FSR of the cavity divided by the cavity finesse: for the design finesse of 1000, this yields a linewidth of approximately 4 MHz. We therefore choose a frequency deviation of 1 MHz to allow for somewhat higher cavity finesse. This 1 MHz dither can be produced by applying a 100 mV sinusoidal signal to the reference cavity PZT.

It should be noted that the lock acquisition difficulties described in chapter 4 will also apply to the dither locking scheme. These difficulties arise because the FSR of the test cavity is very high, so it takes a long time to bring the reference and test cavity resonant frequencies into coincidence. For this reason, the autolocking system described in the following section will be required before the mode-matching locking scheme can be implemented.

6.3 Lock Acquisition

The resonant frequencies of the test and reference cavities must be brought into coincidence so that the reference cavity locking loop can acquire lock. The resonant frequency of the test cavity is fixed because of its monolithic construction. Therefore the two cavities must be brought into resonance by tuning the reference cavity. One could posit a scheme where the reference cavity resonance was swept by applying a voltage to its PZT until a test cavity resonance was detected. Unfortunately there are two problems with this approach:

1. The dynamic range of the reference cavity PZT is approximately 40 MHz, while the FSR of the test cavity is approximately 4.3 GHz. It is therefore unlikely that the PZT will be able to move the reference cavity resonant frequency far enough to reach the test cavity resonant frequency.

2. The dynamic range of the NPRO PZT is only 200 MHz. The laser too is thus incapable of being rapidly tuned over the test cavity FSR. The use of the slow laser heater would thus be required to track a large frequency shift in the reference cavity resonance. Even if the reference cavity PZT were modified so that it was able to move the reference cavity resonant frequency to that of the test cavity, the laser would not be able to follow.

One possible solution for sweeping the reference cavity resonance through the required 4.3 GHz is to change the temperature of the cavity spacer. However, thermal tuning of the reference cavity will be slow because of the high thermal mass of the spacer. The high thermal inertia is problematic because it is hard to manually stop the heating once a resonance is reached. For example, if the reference cavity is heated at a rate that takes 15 minutes to sweep through a test cavity FSR, then each resonance will last only a second. A well designed compensator could lock the cavities during this second, but if the heating were not rapidly removed then the dynamic range of the loop would be quickly exhausted. This problem was encountered in the test described in chapter 5; the reference cavity length was thermally tuned until a resonance was encountered, but it proved impossible to manually remove the heating before the compensator dynamic range was exhausted.

Thermally actuated lock acquisition will only be feasible if some of the overhead can be performed by an automatic method. In particular the heater must be automatically controlled so that it can be switched off when resonance is encountered. The design of the automatic locking system thus requires thorough knowledge of the reference cavity's thermal properties. These properties are difficult to predict theoretically, so in the next section we describe finite element modelling of the cavity. Having established the required thermal properties we will present a microcontroller-based automatic locking system.

6.3.1 Reference Cavity Thermal Modeling

Measurements of thermal properties of large masses are difficult due to the long time constants involved and the inability to measure what is happening "inside" the mass. A finite element model (FEM) of the reference cavity was therefore constructed. This was used to explore the properties of the reference cavity and a simplified model was extracted and used to develop a locking scheme. All finite element modeling was carried out using the Quickfield package from Tera analysis [42]. The original engineering drawings for the cavity spacer were used as the basis of the model and the thermal properties of the aluminium were as follows:

Thermal conductivity: 238 W/Km [48].

Specific Heat: 899 $J/K.kg$ [48].

Density: 2700 kg/m^3 [48].

Two sets of boundary conditions were employed; one to model the cavity in air, the second in vacuum. The difference between the two models was that convection was removed for the vacuum case. There was assumed to be radiative losses from the cavity in both models, but no conduction in either. Parameters used in the boundary conditions were:

Convection film coefficient: 30 W/Km^2 [117].

Emissivity: 0.25 [117].

Ambient Temperature: 300 K

Emissivity of aluminium can vary over quite a wide range (0.1-0.8 [117]) and depends on the machining and oxidation of the surface. The emissivity value used above is therefore only an estimate based on surface condition of the cavity.

The inner surface of the cavity spacer was assumed lossless. Radiation losses from the inner surface of the spacer are negligible as most of the solid angle seen by a point on the inner surface of the cavity is taken up by the opposite side of the spacer. Convective losses from the inner surface were also assumed to be negligible. Although there was a small "breather" hole to allow evacuation of the reference cavity the hole was narrow (4 mm diameter) and long ($\approx$60 mm), preventing the establishment of an efficient convective cell within the breather hole.

Determination of losses from the end surfaces of the cavity was complicated due

to the presence of the glued mirrors. Because the surfaces involved were small and their thermal contact with the mirrors was poor, the losses here were assumed to be the same character as those from the outer surfaces. However, a variety of tests with the modelling packages indicated that the model was quite insensitive to difference in the boundary conditions imposed at the end faces. We therefore feel confident that this simplification did not compromise the results.

Thermal input to the cavity body is assumed to be produced by a non-contacting radiative heater located at one end of the cavity spacer. Thermal input to the spacer is introduced by specifying a value for the heat flux for the appropriate cavity surface. We assume that the thermal dynamics of this heater are fast compared to those of the spacer and model it as an ideal heat source.

A schematic of the reference cavity spacer and the corresponding finite element model is shown in the figures 6.2 and 6.3. As can be seen the mesh on which the model is calculated is finer around the heating area and along the centre of the cavity. The thermal gradients will be large in the regions around the heater, so tight spacing is required maintain the accuracy of the model. The dense mesh is also required along the central axis as this is where the length of the cavity will be measured. The analysis package assumes axial symmetry of the spacer, so only half of the spacer cross-section is shown.

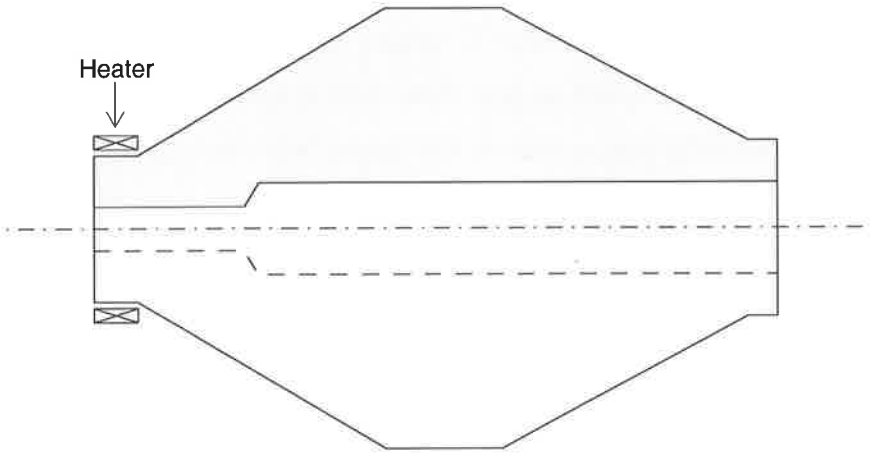

Figure 6.2: Schematic diagram of the reference cavity spacer

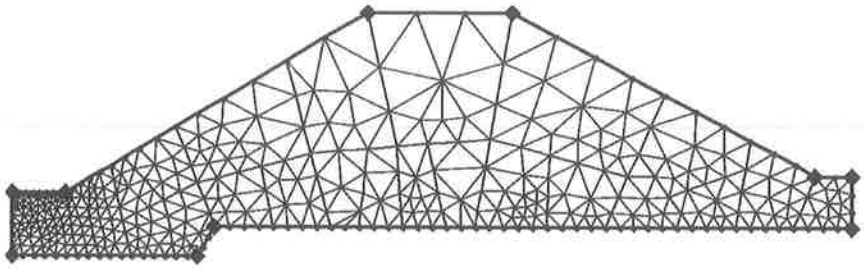

Figure 6.3: The Geometry of the Finite Element Model showing the mesh on which the field equations are solved.

6.3.1.1 Expected behaviour

We can qualitatively predict the behaviour of the cavity under thermal load. Heating of the cavity occurs in a region having small volume, so we would expect this region to heat rapidly. Heat will then diffuse into the bulk of the spacer, heating it more slowly. In regions further away from the heater we would expect transport delay to retard changes in local temperature. We would also expect thermal variations to be smoothed by passage through the intervening material. When heater power is removed we would expect the heat flow to continue until the thermal gradient within the material is removed. The spacer would then cool by losses from the surfaces.

Thus, in response to a heater input we would expect a rapid initial change in cavity length due to localised heating and the consequent expansion in the region near the heater. Once the heater is turned off we would expect the bulk of the spacer to continue to expand as heat flows into it from the heater region. However during this period the region nearest the heater will cool significantly, leading to a strong local contraction. The overall length changes immediately following turn off of the heater will therefore be dependant on the balance of competing effects and the finite element model will be required to determine the actual behaviour.

6.3.1.2 Modeling results

6.3.1.2.1 Steady state conditions Under constant thermal load the cavity temperature will rise above ambient until the losses are equal to the input power. Figure 6.4 shows the steady-state distribution of temperature throughout the spacer

material for a constant heat load of 500 mW from the heater. This simulation assumed that the spacer was in air and thus that convection was the dominant heat loss mechanism. As expected, the temperature is highest in the vicinity of the heater, although the temperature of the entire spacer is between 13 and 18 K above ambient.

The strain and displacement produced in the spacer are also shown in the figure. Displacement in the figure is along the colour gradient: as expected the cavity increased in both length and diameter.

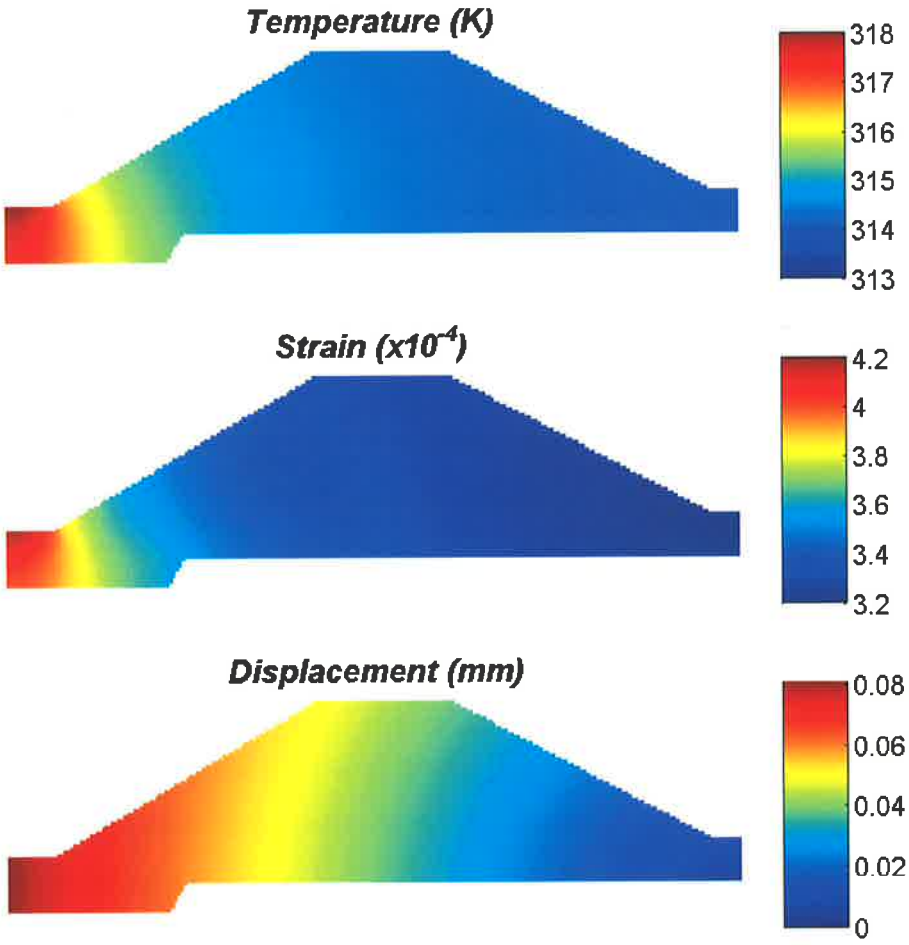

Figure 6.4: Maps of steady-state temperature, principal strain and displacement within the reference cavity spacer when heated by 500 mW in air.

The above procedure was repeated for different heater powers to determine the steady-state relationship between heater power and cavity length. The temperature,

strain and displacement maps produced in these tests are qualitatively similar to those shown in figure 6.4 and are thus not shown here. The cavity length changes from these tests are shown in figure 6.5. Analysis of the graph yields a relationship between the heater power, P, and cavity length change, L, of approximately $L = 1.7 \times 10^{-4} P \ \ m/W$.

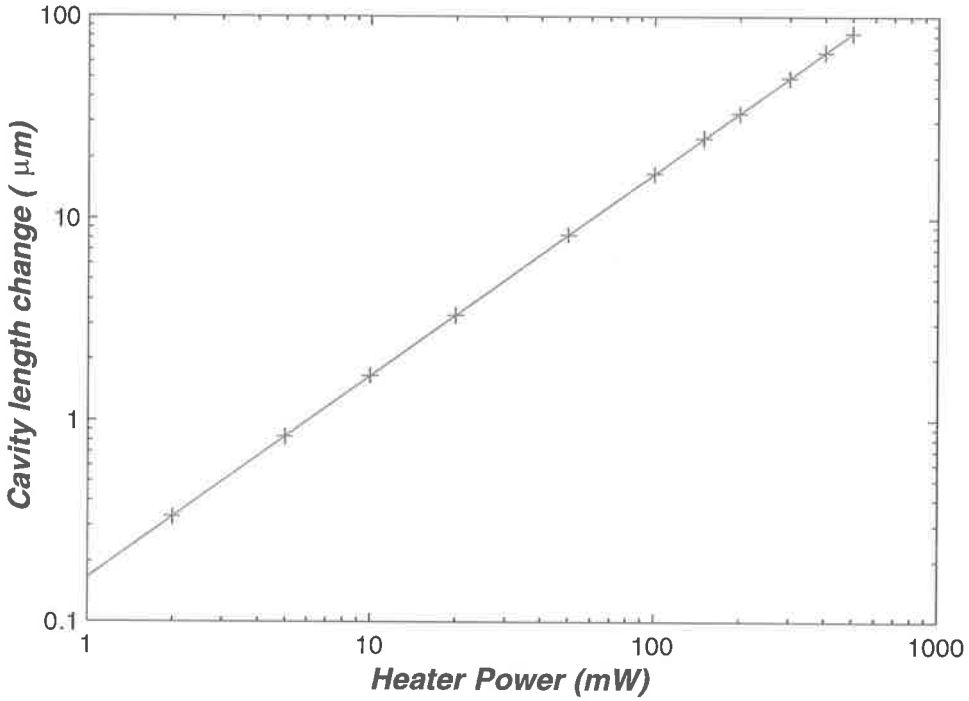

Figure 6.5: Variation of the cavity length with heater power under steady state conditions when the prestabilisation cavity is in air. The crosses are data points taken from the FEM and the line is the linear least squares fit.

These tests were repeated for the spacer in a vacuum by removing convection from the model. Figure 6.6 shows the steady state distribution of temperature, strain and displacement with 500 mW of heater power. Not surprisingly, the lack of convection leads to a significantly hotter cavity, in this case approximately 150 K above ambient. The strain and displacements produced are correspondingly higher than in the atmospheric case.

Plotting the power / cavity length relationship yields the figure shown in 6.7. Radiation is the only loss mechanism in vacuum and is dependent on T^4. Thus the spacer dissipates heat much more effectively when it become hot and the curve

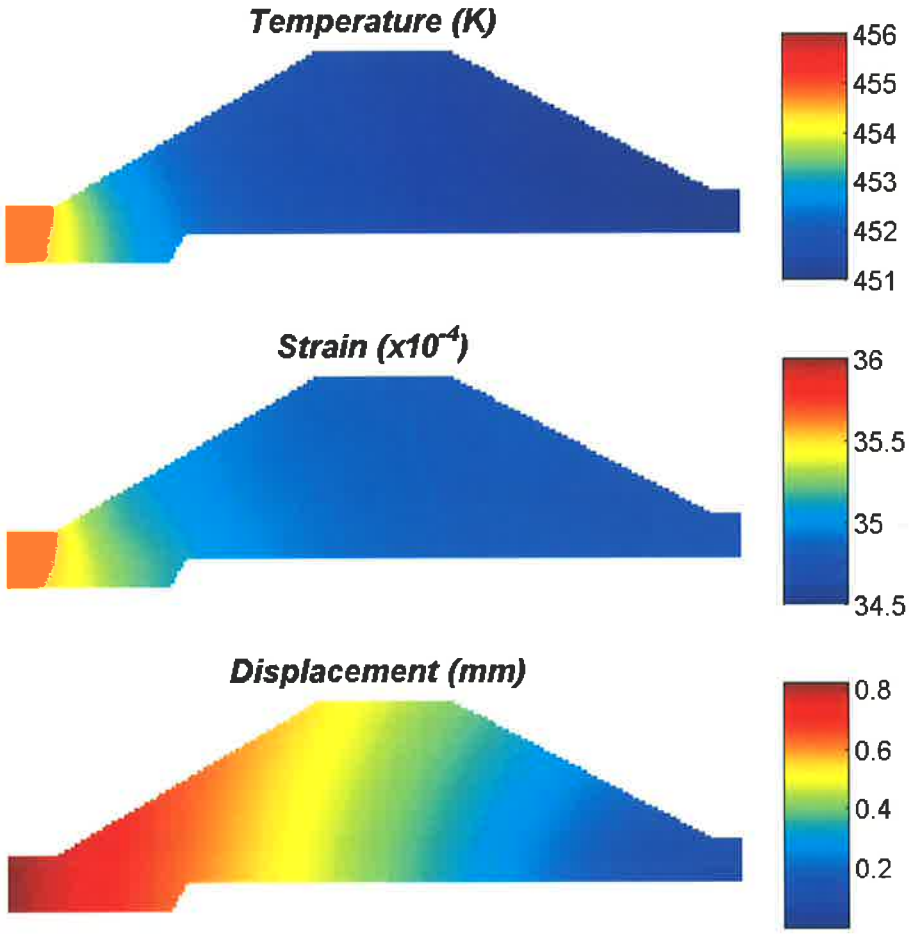

Figure 6.6: Maps of steady-state temperature, principal strain and displacement within the prestabilsation cavity spacer when heated by 500 mW in a vacuum.

deviates from a straight line for higher heater power. For modest heater power ($<$ 100 mW) the relationship between cavity length and heater power is approximately $L = 3.3 \times 10^{-3} P \ \ m/W$.

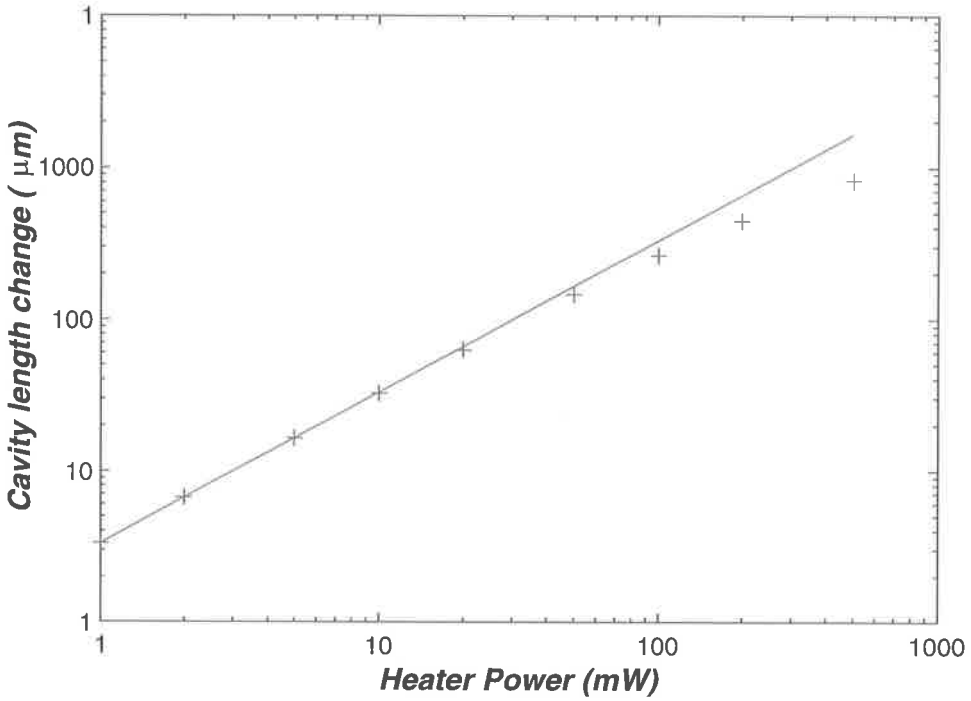

Figure 6.7: Variation of the cavity length with heater power under steady state conditions when the reference cavity is in a vacuum. The crosses are data points taken from the FEM, while the line represents a fit to the low power portion of the curve.

The mean steady-state temperature elevation above ambient at the surface of the cavity is shown in figure 6.8 as a function of heater power. As expected the spacer is considerably hotter in vacuum than in air. Figure 6.8 shows that a heater power of about 10 mW will be required to maintain the reference cavity at about 5 K above ambient in vacuum and hence maintain the length of the reference cavity after it has been heated to acquire lock.

6.3.1.2.2 Heat-Length Relationship Similar tests to those shown above were conducted with no losses present and revealed a spacer expansion of 0.096 $\mu m/J$. This thermal response allows us to determine the energy required to change the spacer length by an amount equivalent to the difference between adjacent cavity

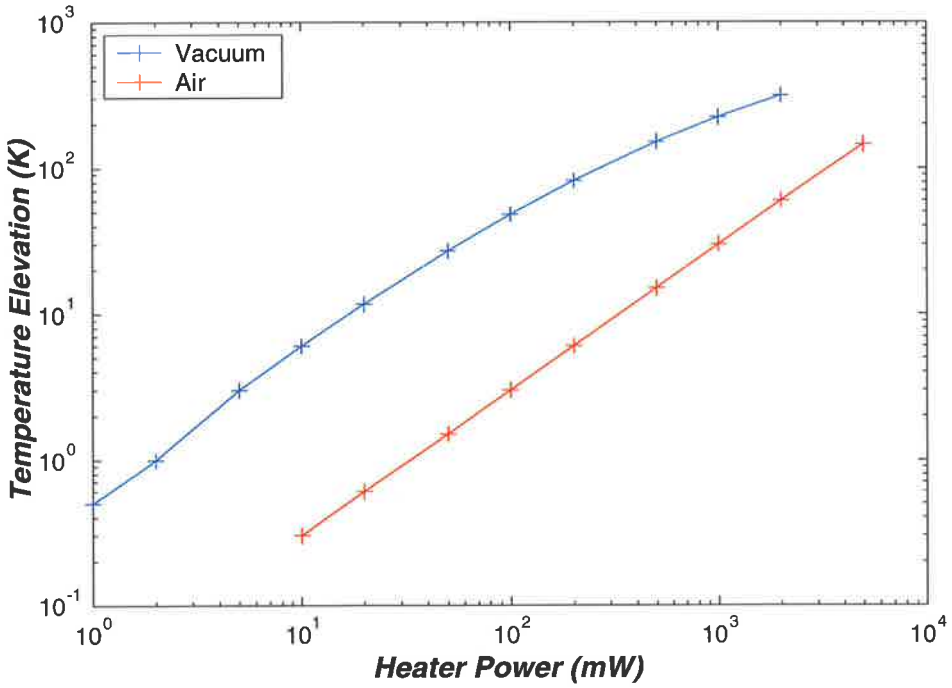

Figure 6.8: Variation of reference cavity steady state temperature as a function of heater power. The blue line shows temperature when the cavity is in vaccum, the red line the temperature in air.

modes.

We know that the relationship between changes in cavity length and resonance frequency is

$$\begin{aligned} \frac{\Delta f}{f} &= \frac{\Delta L}{L} \\ \Delta L &= L\frac{\Delta f}{f} \\ \Delta L &= \frac{\lambda L}{c}\Delta f \end{aligned} \tag{6.1}$$

Substituting in the required frequency shifts (4.27 GHz for the longitudinal mode spacing, 522 MHz for the transverse mode spacing) we find that the required optical path length changes are 8.1 μm and 0.98 μm for longitudinal and transverse modes respectively. Corresponding changes in the spacer length are approximately half this, so by using the thermal response of 0.096 $\mu m/J$ we can determine the required energies as

$$\begin{aligned} E_{LM} &= 42\ J \\ E_{TM} &= 5.1\ J \end{aligned}$$

where $E_{LM}$is the energy required to change the reference cavity resonance by one longitudinal mode of the test cavity (4.27 GHz), and E_{TM} is the energy required to change the cavity length by one transverse mode of the test cavity (522 MHz). The model yields spacer temperature changes of 0.082 K and 0.67 K corresponding to these deposited energies respectively.

Figure 6.8 shows that when the reference cavity is in vacuum, a temperature increase of 0.67 K will require a steady state power of about 2 mW, which is much less than the proposed heater power. We can thus neglect cavity losses during a heating cycle and assume that deposition of the above energies is sufficient to move between cavity modes.

6.3.1.2.3 Transient Response of the Cavity The temporal response of the reference cavity is very important to the locking process. To produce a reliable locking system we need to know the time constants of the cavity response and we also must know how the cavity behaves when the heater is turned off, which depends on competition between different effects. As these responses cannot be easily predicted, we again use the model. A variety of simulations were conducted to extract information about the time varying temperatures, strains and displacements at various points in the spacer material.

A constant heat load of 0.5 W was applied at t=0 and removed 30 minutes later. Figures 6.9 shows the evolution of temperature and displacement during the heating. As time increases the spacer temperature increases and the cavity becomes longer and wider. Figure 6.10 shows the cavity behaviour as the cavity is subsequently allowed to cool. As expected, the temperature in the cavity quickly becomes uniform after the heat load is removed.

The temporal response of the temperature at six locations was extracted from the maps shown in figures 6.9 and 6.10. These locations were at the heated end of the spacer and then at points 25, 50, 75, 150 and 225 mm along the inner surface of the cavity. The temporal behaviour at these representative locations is shown in figure 6.11. There are four distinct phases in the heating and cooling cycle. The temperature at each point rises rapidly immediately after the heater is turned on. As expected the temperatures at the locations closest to the heater rise fastest. The time constant of this heating phase is governed by heat transport times in the spacer material adjacent to the heater.

Once the appropriate temperature gradients have been established, the initial transient ceases and the temperature changes at a uniform rate throughout the spacer. The shape and time constant of this rise is governed primarily by losses from the spacer surface.

When the heater is turned off there is another transient as heat is redistributed within the spacer to remove the temperature gradients. The temperatures converge quickly and the spacer cools fairly uniformly, although the temperature in the centre of the spacer remains slightly warmer than at the surface.

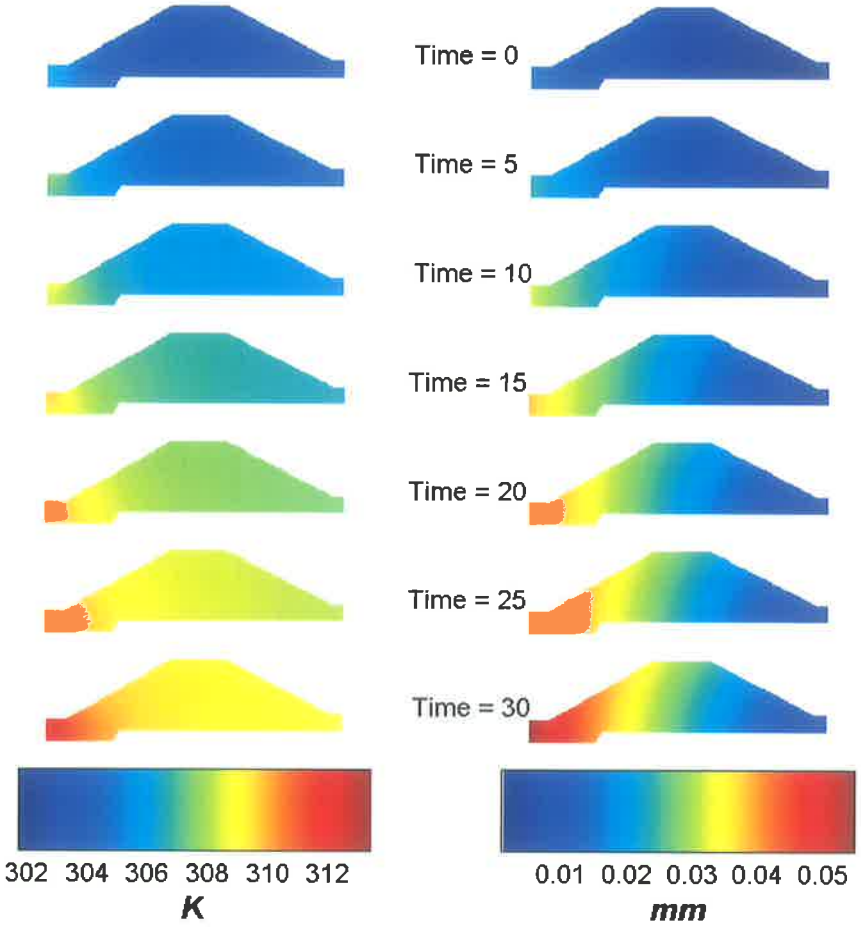

Figure 6.9: Maps of the temperature (left) and displacement (right) distributions as the cavity is heated in air. Blue indicates low temperature and displacement, red high. The time from the start of heating is shown in minutes.

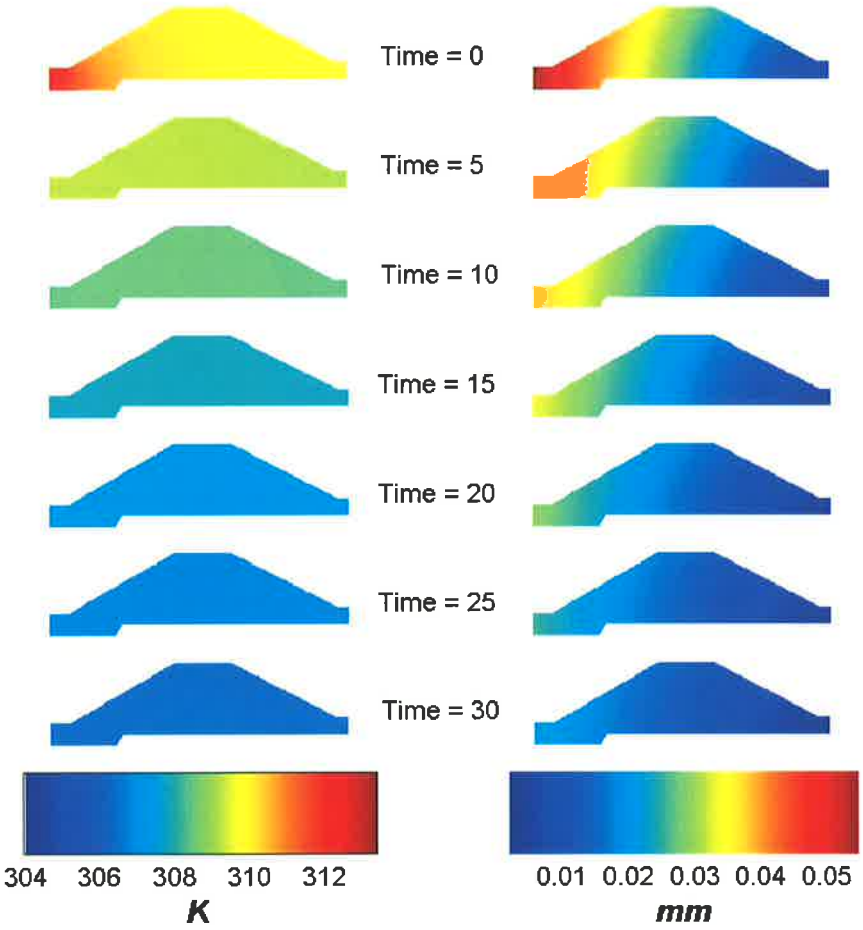

Figure 6.10: Maps of temperature (left) and displacement (right) distribution as the cavity cools in air after the heating shown in figure 6.9 above. Blue indicates low temperature and displacement, red high. The time from the end of heating is shown in minutes.

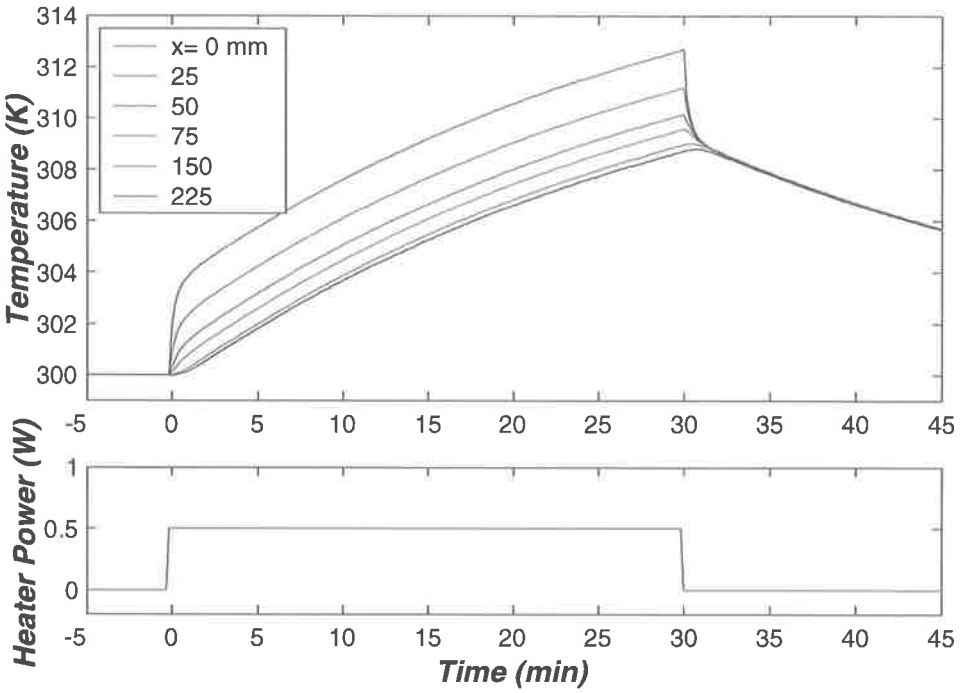

Figure 6.11: Temperature response at several locations in the spacer when subjected to the heat load shown in the bottom part of the figure. The positions of the measurement points are shown in the legend as the distance from the heater end of the spacer. All measurement points were located along the inner radius of the spacer.

Length Response: The length change of the cavity over time can be seen in figure 6.12. This figure was produced by heating the cavity with a 0.5 W heat source for ten minutes. The length response shows the same behaviour as the temperature curves above. When the heat load is changed there is a relatively fast transient and then the cavity changes length more uniformly and the change can be considered linear for the remainder of the heating cycle.

Most critically we see that the cavity contracts immediately after the heater is turned off. This is important in our application as it allows the heater to be run until a cavity resonance is reached. After the heater is switched off the cavity will not continue to expand past the desired resonance; as long as the contraction is not too strong the servo will be able to catch and hold the resonance. Additionally, if the contraction transient is too large it can be reduced by reduction of the heater power.

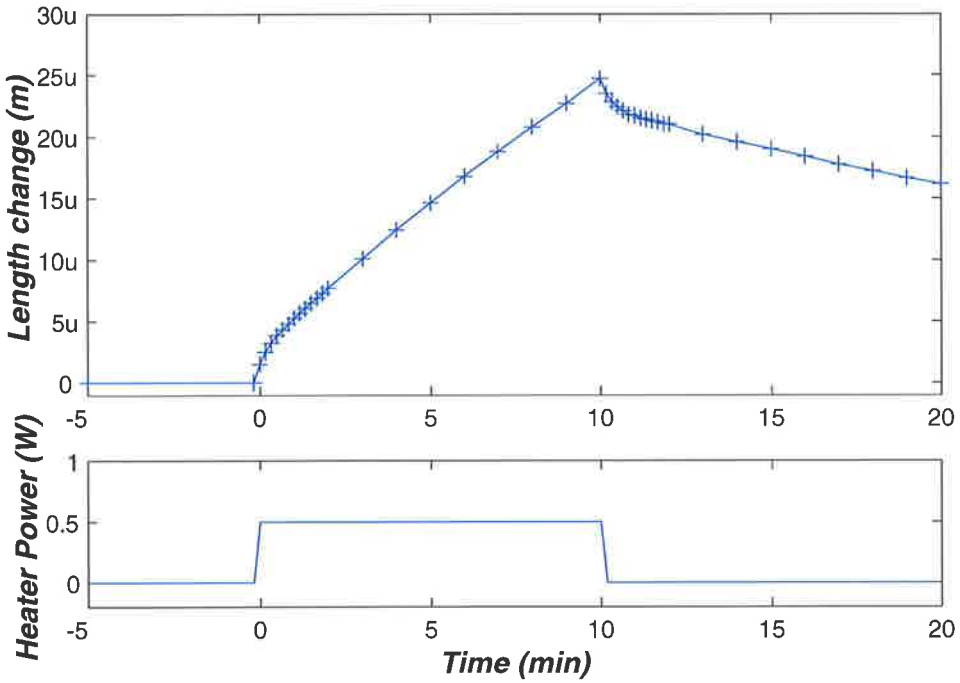

Figure 6.12: Variation of cavity length with time for a 0.5 W heat load. The spacer is in air.

Another feature of interest is the markedly different time constants of the heating and cooling parts of the cycle. This difference is even more pronounced when the same simulation is carried out in a vacuum, where heat losses can only be radiative. Unless the steady state temperature of the cavity is far above ambient, it is not possible for the spacer to lose heat quickly enough to keep the cooling time constant similar to that of the case of the cavity in air.

6.3.2 Auto-Lock Implementation

In this section we describe the design of an auto-lock system that automatically locks the frequency of the reference cavity and hence the probe laser to the TEM_{00} mode of the test cavity. This system should greatly simplify the process of locking the system.

The system starts by heating the reference cavity so that it scans through one FSR of the test cavity and finding the maximum power transmitted through the test

cavity to the FSL detector. This maximum power corresponds to the TEM_{00} mode of the test cavity and can be used to define a threshold that the auto-lock system will use to detect build up of power in the desired test cavity mode. To acquire lock the system heats the reference cavity slowly until the power threshold is detected at the FSL detector. The heater power is then immediately reduced to allow the servo system to lock the reference cavity to the test cavity mode. Once lock has been achieved the operator is asked to confirm whether the resonant mode is TEM_{00} or otherwise initiate a search for the next test cavity mode.

6.3.2.1 Finite State Machine

The auto-locking system is based on a finite state machine that handles the various stages of the locking process. Figure 6.13 shows a description of the state machine using the conventions of the Unified Modelling Language (UML) [118]. The operation of the state machine is as follows:

Laser Unlocked: This is the initial state for the system. The system returns to this state whenever there is insufficient light detected after the reference cavity, indicating that the laser is not locked to the reference cavity. Once the power at this point exceeds a preset threshold the state machine moves to the scan state.

Scan: In this state the system scans the probe laser through a test cavity FSR as rapidly as possible to measure the maximum power transmitted through the test cavity. In this mode the heater power is set to its maximum value, $P_{\max}$, so that the reference cavity expands as rapidly as possible. During the expansion the system continuously samples output of the FSL detector and records its maximum value.

The time required for the scan is determined by the rate at which reference cavity can be heated. In section 6.3.1.2.2, we saw that 42 J is required to scan the reference cavity through a FSR in the lossless case. We can therefore approximate the required time for this scan as $42/P_{max}$, where $P_{max}$is the maximum power of the heater. At the conclusion of the scan we set the FSL detector threshold at half of its maximum detected value. This will then serve as the threshold to indicate that the TEM_{00} mode is resonant in the test cavity.

Once the scan has been completed the system automatically moves to the Acquire

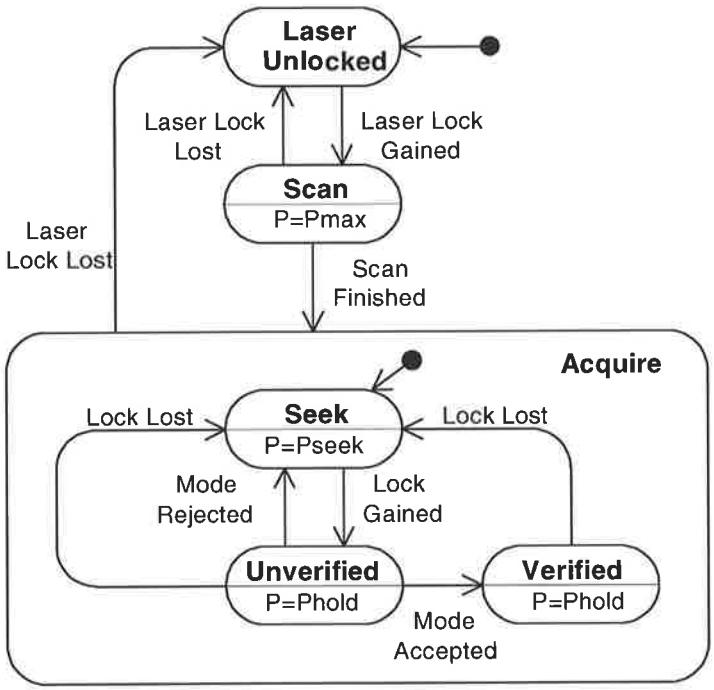

Figure 6.13: UML finite state machine diagram for the autolocking system. The rounded boxes represent the states that describe the system. Arrows indicate the possible state transitions and each is labelled with the event that causes the transition. Pmax, Pseek and Phold represent the heater powers during the Scan, Seek and Verified states.

state, specifically to its Seek substate.

Seek: The intent of this mode is to search for a test cavity mode. Thus, the heater is run at lower power than in the scan state so that the reference cavity length increases slowly.. When the power detected at the FSL detector exceeds the threshold set during the scan process, the system switches to the Unverified Lock state.

Heating during the Seek mode must be slow enough that the control system is able to lock to a resonance when it is encountered. The unity gain bandwidth of the servo system is approximately 500 Hz, so allowing a reasonable margin of safety we choose a desired resonance duration of 20 ms in the open loop case. The unity gain bandwidth of the loop will be lower when attempting to lock to a transverse mode in a well mode matched system, so the system will not lock on weak transverse modes that pass through their resonances in 20 ms.

We saw above that the energy required to move between successive longitudinal modes of the test cavity was 42 J. For a test cavity finesse of 1000 it will therefore take 42 mJ to move the reference cavity through a given resonance. We must ensure that this event takes longer than 20 ms. The locking constraint therefore imposes a maximum power of $42/20 \approx 2$ W during the Seek heating cycle.

A second constraint on maximum power is imposed by the dynamic range of the reference cavity PZT, which must take up the transient that occurs after the heater is turned off (see figure 6.12). We must ensure that the change in cavity length is sufficiently small that it can be compensated for by the reference cavity PZT. If we assume that the PZT is in the centre of its range when lock is acquired then the PZT's dynamic range is approximately 3 μm. To ensure that there is sufficient remaining dynamic range after the locking transient, we chose to limit the acceptable length change to 1.5 μm. Figure 6.14 shows modelled length versus time profiles produced when the cavity is heated in air using different powers. For the example shown, a total energy of 5 J is deposited in the reference cavity but the length transient is independent of the energy deposited. As expected, higher powers exhibit larger length transients because of the formation of a stronger hot spot in the heater region. At lower powers the hot spot temperature and consequently the

local expansion in the heater region is lower. The figure shows that a heater power of 0.5 W produces an length change of approximately 1.5 μm after turn-off. Any heater power, P_{seek}, less than 0.5 W will therefore meet the specification. Note that the dynamics in this region are dominated by the movement of heat through the cavity spacer, not by losses from the surface. Thus, we would expect that in a vacuum we would observe transients very similar to those in an atmosphere. This result is confirmed by the results of a FEM simulation of the vacuum environment, the results of which are shown in figure 6.15. As can be seen the curve is identical to that shown in figure 6.14 for the atmospheric case.

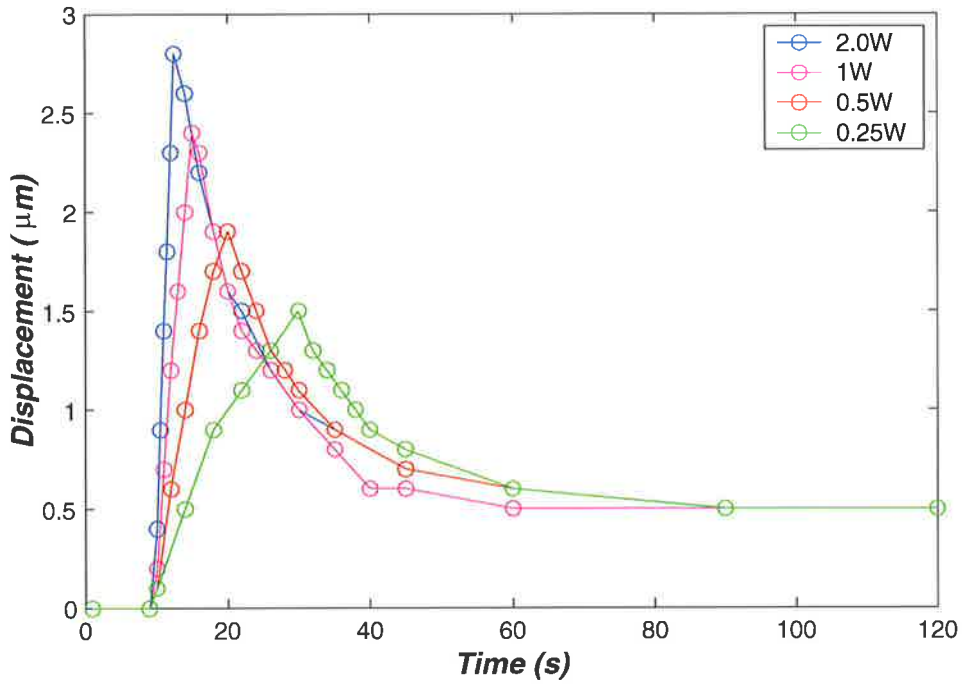

Figure 6.14: Variation of cavity length vs time as 42 mJ is deposited in the reference cavity using various heater powers. The reference cavity is in air.

Unverified Lock: In the Unverified Lock state the servo system attempts to hold the two cavities in resonance. The purpose of this state is to allow the user to confirm whether the mode of the locked test cavity is TEM_{00}. While higher order test cavity transverse modes should not lock if the threshold was set correctly in the Scan mode, this check is included to handle unexpected locking events. If the mode is TEM_{00} then the user confirms this fact and the state machine moves to the

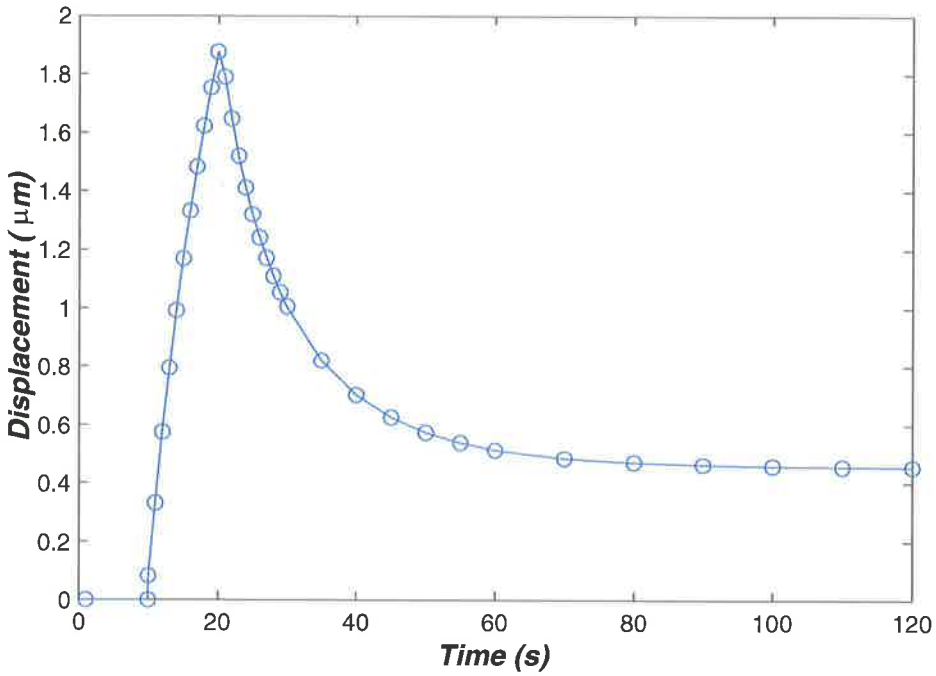

Figure 6.15: Length variation with time as 42 mJ are deposited in the reference cavity usig a heater power of 500 mW. The reference cavity is in a vacuum.

Verified Lock state. Otherwise the state machine moves back to the Seek state to search for the TEM_{00} mode.

In this mode the heater power is lowered to a level, P_{hold}, that holds the reference cavity at a constant temperature. Examination of figure 6.8 shows that to maintain the reference cavity at five kelvin above ambient in a vacuum requires a steady state heater power of order 10 mW.

Verified Lock: It is in the Verified Lock state that the test cavity length noise is measured. The system remains in this state until lock of the probe laser is lost. When lock is lost the system returns to the Seek state to reaquire the test cavity resonance.

In this Verified Lock state, the heater power is again set to a level that maintains the reference cavity at a constant temperature. The verified lock state will persist for as long as the reference cavity PZT can take up frequency drifts caused by temperature changes in the reference cavity. Eventually the dynamic range of the PZT will be exhausted and the system will lose lock. The time that the system will

stay locked will thus depend on the precision with which the heater maintains the reference cavity temperature during this period. In principle it would be possible to use a feedback control loop to maintain the system in this state, using the DC component of the PZT voltage as the error signal. However, the system is expected to be sufficiently stable that a feedback loop will not be required.

If lock is lost then the system sets the heater power to zero for a short period, thereby ensuring that the reference cavity cools, simplifying lock reacquisition during the subsequent Seek state.

6.3.2.2 Auto-Locker Design

The autolocking system uses the photodetectors already present in the optical system and described in chapter 4. In addition a high current amplifier is required to drive the heater and a microcontroller board is used to implement the finite state machine.

6.3.2.2.1 Heater Driver The heater driver is a simple amplifier which is capable of supplying a high current to its load. A simplified schematic is shown in figure 6.16; a more complete circuit description can be found in appendix E. The amplifier was designed to drive an ampere through any heater having resistance up to 24 Ω.

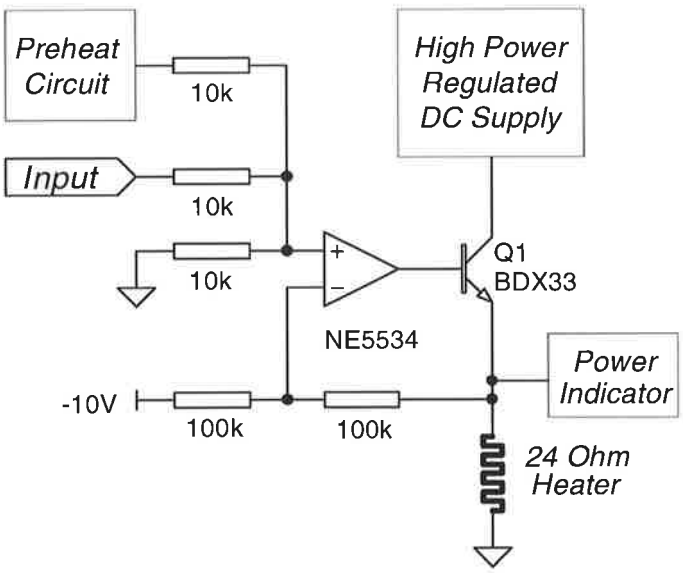

Figure 6.16: Simplified schematic of the heater driver.

6.3.2.2.2 Finite-State-Machine Microcontroller A microcontroller board will be used to implement the auto-locking finite state machine. The board uses a AT90S4433 microcontroller, a member of the AVR microcontroller family made by Atmel corporation [119]. The microcontroller is equipped with a multi-channel analog to digital converter which is used to sample the photodetectors that measure the light transmitted through the two cavities. The microcontroller is not fitted with a digital to analogue converter, so an external converter is used to drive the heater amplifier. The user interface of the board consists of two switches and a number of LEDs to indicate the state of the controller. These functions are handled by conventional input/output pins of the microcontroller.

The microcontroller board schematic is shown in figure 6.17. The microcontroller is programmed using assembly language via a programming dongle connected to the serial port of a personal computer. Pseudocode of the program used to implement the finite state machine may be found in appendix F.

6.4 Conclusion

The tests described in chapter 5 uncovered two problems that stem from the choice of a short monolithic test cavity. Conventional mode-matching procedures cannot be used with this system and the inability to tune the test cavity makes lock acquisition very difficult. This chapter presented possible solutions for these two problems.

A solution to the mode matching problem was first presented. The conventional method of sweeping the probe laser frequency through resonance with the test cavity transverse modes and minimising the power in higher order transverse modes is impossible. Instead a supplementary locking technique is used to lock the reference cavity to the test cavity. This allows the power in the TEM_{00} mode to be maximised by adjustment of the mode matching optics. This temporary locking system is too noisy for use in the system in its data collection configuration, so once mode matching is complete the system will be switched back to use fringe side locking.

A possible automatic system for locking the reference cavity to the test cavity was then described. The system tunes the reference cavity thermally so that it

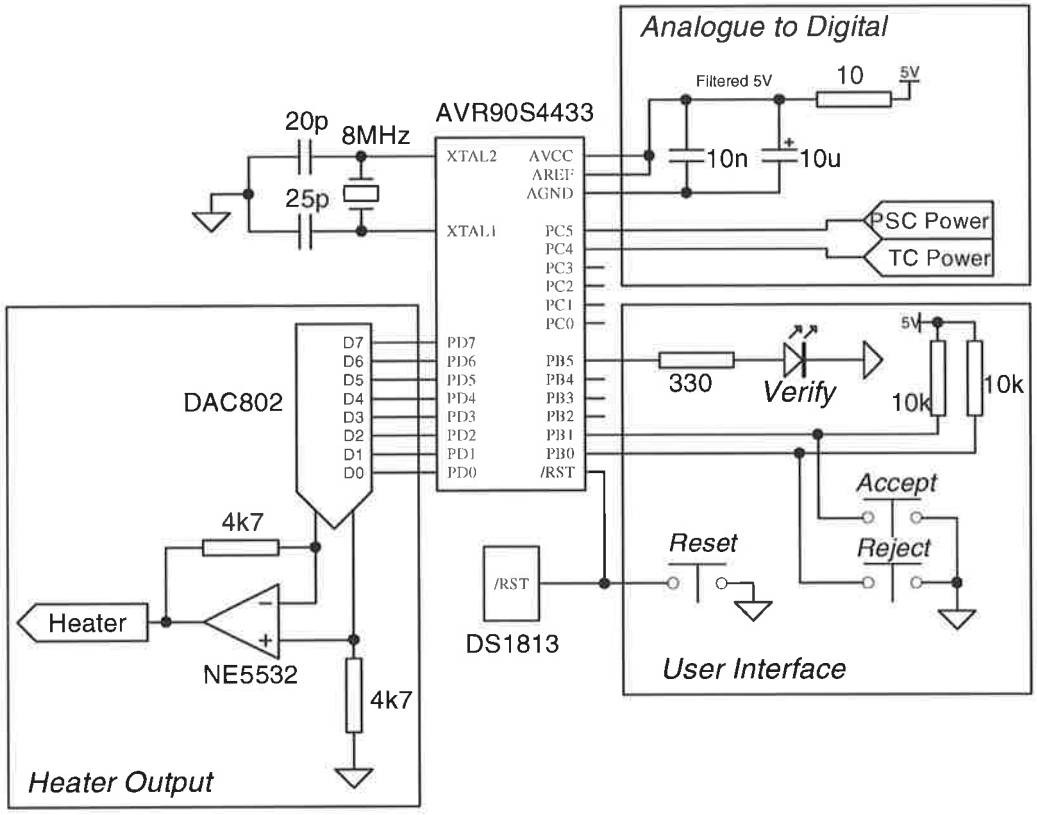

Figure 6.17: Schematic diagram of the microcontroller board used to implement the finite state machine.

comes into resonance with the test cavity. A finite element model of the reference cavity was constructed and was used to determine the thermal properties of the cavity. A detailed design of an auto-locking system design was then presented. This system automatically heat the reference cavity until a test cavity resonance occurs. When a resonance is found, the system will attempt to acquire lock. Because there is a risk of locking on a undesired transverse mode, the user may reject the cavity mode and the system will automatically find another. This semi-automatic scheme should allow locking of the test cavity TEM_{00} mode much more quickly than would be possible by hand.

Chapter 7

Conclusion

This thesis has described the design and construction of an instrument intended to measure the spectra of length noise in optical materials. It will initially be used to measure the spectrum of thermoelastic noise in sapphire, a noise source of particular significance to the next generation of interferometric gravitational wave detectors.

The instrument could also be used to characterise intrinsic length fluctuations in a variety of optical materials. This would enable the experimental confirmation of theoretical predictions for these fluctuations so that the models can be used with confidence in designing the optical elements of ultra-high precision measuring devices of the future. The system has been designed to measure length noise in the frequency band between 10 Hz and 1 kHz, where optical media intrinsic length noise dominates the instrumental noise floor of current interferometric gravitational detectors.

The system described in this thesis employs a novel monolithic test cavity that is probed by a stable laser. The use of a monolithic test cavity renders the experiment relatively insensitive to seismic noise and reduces the effects of external mechanical noise because they are common mode to the two mirrors coated onto the opposite faces of the cavity. Radiation pressure noise is also drastically reduced in a monolithic cavity because it must cause crystal strain to produce a cavity length change.

Variations in the test cavity length are inferred from their effect on the test cavity's resonant frequency. A Fringe-Side-Locking sensor is used to measure the

frequency difference between the laser and the test cavity and the resultant signal is used to provide feedback to the probe laser frequency so that the laser remains in resonance with the test cavity TEM_{00} mode.

The frequency stability of the probe laser beam must be less than the expected shift in test cavity resonant frequency caused by its thermoelastic noise. Free-running lasers do not have adequate stability for this measurement, so the laser frequency must be stabilised to a reference cavity. If the reference cavity is made sufficiently long and has sufficiently low acoustic loss, then its length noise does not unduly compromise system sensitivity. The sensitivity of the instrument is then potentially limited by two types of noise; noise intrinsic to the optical media and noise in the sensors and electronic components used to implement the experiment's various stabilisation loops.

Extensive theoretical analysis of thermoelastic noise and other intrinsic length noise was described in chapter 2. Variations in test cavity length due to thermal noise in its substrate and coatings, thermorefractive noise, photothermal noise, classical radiation pressure and quantum radiation noise were each modelled and it was shown that thermoelastic noise will dominate the intrinsic noise floor of the test cavity over the target frequency range of 10 Hz to 1 kHz and indeed somewhat beyond.

Length noise in the reference cavity was also modelled in chapter 2, where it was shown that the frequency stability of the reference cavity should be sufficient to measure thermoelastic noise in sapphire when the system is placed in a vacuum environment. However, refractive index variations in the air inside the reference cavity will limit the reference cavity frequency stability to about 100 $mHz/\sqrt{Hz}$ when the reference cavity is at atmospheric pressure. This level is insufficient to measure sapphire thermoelastic noise, so an appropriate vacuum chamber was prepared and is described in appendix C.

Although the seismic isolation required by the experiment was relatively modest, a simple system was required to isolate the test cavity from mechanical noise in the remainder of the system and to ensure an adequate acoustic Q. As described in chapter 3, a monolithic flexure mount produced at the University of Western Australia was used to suspend the test cavity from an aluminium intermediate mass,

which was in turn suspended using a double loop of stainless steel wire. Eddy current damping was applied to the intermediate mass to damp the main pendulum mode and other rocking modes of the intermediate mass. This damping was sufficient to increase the damping of the main pendulum mode from $\zeta = 3.5 \times 10^{-4}$ to about 0.4.

Chapter 4 opened with a description of the system used to lock the frequency of the NPRO laser to the reference cavity. The reference cavity is a three-mirror monolithic ring cavity, the mirrors being glued to a massive aluminium spacer using a vacuum compatible epoxy, and has a finesse of 4100 ± 500. One of the cavity mirrors was mounted on a high dynamic range piezoelectric transducer to allow tuning of the cavity's resonant frequency.

Frequency variations between the laser and the reference cavity were measured with a Pound-Drever-Hall (PDH) sensor, the design of which was discussed in detail in chapter 4. Two actuators were available for tuning of the laser frequency; a piezoelectric transducer (PZT) bonded to the laser crystal to change its resonant frequency *via* strain-induced birefringence, and a thermal control system that changed the temperature of the laser element. The response of the PZT was measured and its gain found to be constant for frequencies below its first mechanical resonance at 100 kHz. Measurements of the thermal response were also made, which confirmed that it was useful only for very low frequency control of laser frequency.

A compensator was designed to stabilise the laser frequency by driving the laser PZT element. This stabilisation loop has a unity gain frequency of 90 kHz, but a small dynamic range of order 200 MHz. This range would not be sufficient in practice, so additional feedback was applied to the thermal control system. This supplementary control loop acts to keep the PZT at the centre of its operating range. The resulting frequency stabilisation was found to be robust and measurement showed that the system achieved adequate performance for frequencies less than 700 Hz, where the servo gain became limited by the achievable bandwidth of the PZT. Unfortunately the performance of the loop was not as good as predicted by the compensator design, instead being limited by a white noise source of approximately 50 $mHz/\sqrt{Hz}$. We believe that this noise was likely caused by laser beam intensity noise at the PDH frequency, or by demodulation of pickup on the PDH

sensor line. These issues will need to be addressed in the final experiment.

Intensity stabilisation of the laser was also discussed in chapter 4. The actuators available to control the laser intensity were first characterised and as a result it was decided to use the laser's in-built noise-eater and to apply feedback through the laser control unit. This modulation input was unexpectedly found to be non-linear, which complicated the design of the compensator. A gain-scheduled compensator was designed to close the intensity stabilisation loop and allowed a loop gain five times higher than would have been achievable if the actuator were assumed to be linear. A measurement of the stabilised intensity noise was made at the in-loop sensor and was shown to meet the required levels for all frequencies below 1 kHz. However, a measurement made with an out of loop sensor showed significantly higher frequency noise at low frequencies, although the level of intensity noise was still adequate for the experiment. There are a number of possible causes for this excess out-of-loop noise; beam jitter, reference noise and thermal effects in the photodetectors being the most likely. Because several of these noise sources will be reduced in vacuum environment, the intensity noise spectrum will be remeasured when the system is in a vacuum. If appropriate, measures to reduce the effects will be undertaken at that time.

A fringe side locking detector is used to measure frequency deviation between the probe laser and the test cavity. The system is designed to feed this frequency deviation back to the reference cavity PZT so that the probe laser stays in resonance with the test cavity. However, it is difficult to obtain resonance of the probe laser in the test cavity because of the limited frequency range over which the probe laser can be tuned. The dynamic range of the reference cavity PZT is only 40 MHz, far smaller than the 4.3 GHz free spectral range of the test cavity. As a result, it was necessary to employ thermal tuning of the reference cavity to extend its dynamic range. This thermal tuning was found to be extremely difficult to manage manually and was not successfully implemented during the work described in this thesis.

A graph of the noise spectra for the three main stabilisation systems are shown in figure 7.1. As can be seen, the system sensitivity is adequate over the target frequency range up to about 700 Hz, where it is limited by noise in the reference cavity

servo. The noise in this region can be improved by increasing the bandwidth of the reference cavity servo with an additional fast actuator. As described above, the low frequency end of the spectrum is limited by intensity noise and improvements to this servo should allow significantly more margin in this frequency region. Between the two frequency extremes, the system sensitivity is good enough to measure thermoelastic noise in sapphire, and in particular should allow the direct measurement of the spectrum around the critical frequency where the thermoelastic noise spectrum changes slope.

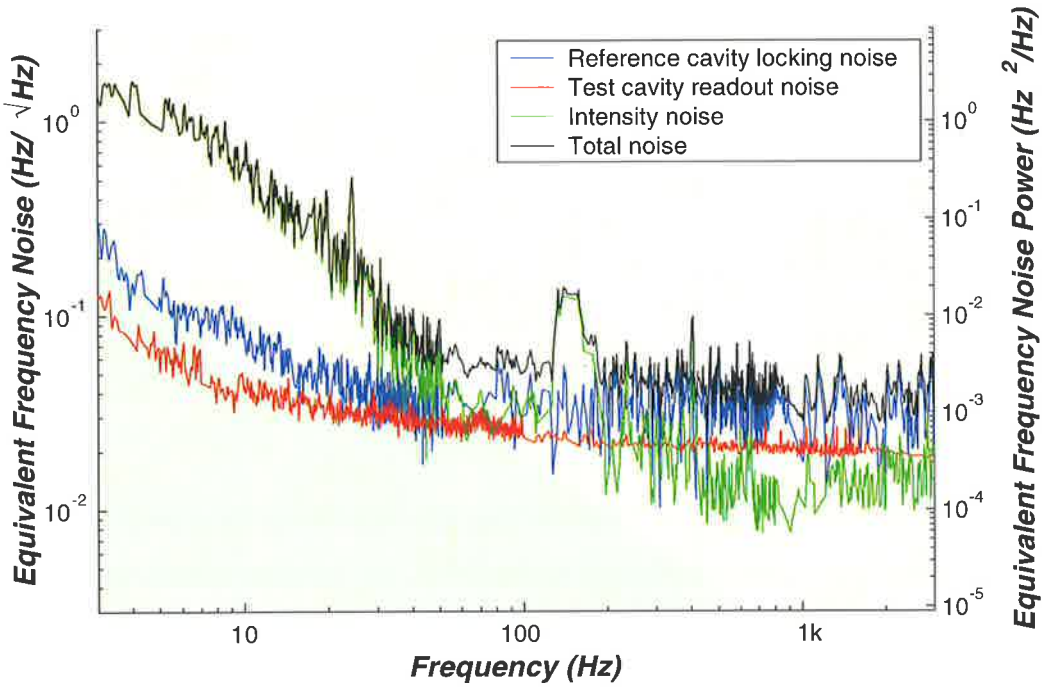

Figure 7.1: Summary of the noise levels achieved in the three main stabilisation loops used in the experiment.

Some preliminary measurements of the length noise of sapphire were presented in chapter 5. These data were taken by heating the reference cavity so that it passed through resonance with the test cavity and measuring the frequency noise while the resonance lasted. The heating during this procedure was very rapid so that a resonance could be reached in a reasonable time, but this meant that the locking servo was not able to catch and control the resonances when they occurred. As a result each measurement lasted only a few seconds, but despite this fact revealed

a spectrum that was within a factor of ten of the expected level of thermoelastic noise in sapphire. Given that the measurements were significantly compromised by the short measurement durations and that the reference cavity as at atmospheric pressure, the result gives us considerable confidence that the system is working as expected.

The design of an automatic locking system to overcome the reference cavity locking difficulties was presented in chapter 6. The autolocker automates the process of heating the reference cavity until its resonant frequency reaches that of the test cavity. The system then regulates the temperature of the reference cavity so that feedback to its PZT can control the resonant frequency while a measurement is made. Finite element modelling of the reference cavity's response to thermal loads was undertaken so that appropriate heating system parameters could be determined and code was written for a microcontroller to implement a state machine approach to the autolocker system. This system is expected to completely manage the lock acquisition process in the next phase of the project.

7.1 Future Work

The next phase in the project requires that the instrument be moved into the vacuum environment. As described in appendix C, the vacuum system has been built and tested so is ready for this step. Several aspects of the system performance are expected to change in a vacuum environment, so their response will need to be re-measured. The performance of the suspension system can be expected to improve in the absence of gas damping and reduced acoustic excitation. The out-of-loop intensity noise spectrum presented in section 4.4.4 is also expected to change significantly. There are a number of improvements that could be made to the intensity stabilisation system, such as improvements to the voltage reference noise and temperature stabilisation; whether these changes are necessary depends on the result of the updated measurements.

The automatic lock acquisition system described in chapter 6 should then be installed and the enhanced mode-matching system also described in that chapter

can also be implemented.

Initial testing of the completed instrument will confirm whether the acoustic losses in the test and reference cavities are adequate. If required, additional work could be undertaken to improve their suspension systems to further reduce acoustic losses.

The instrument will then be capable of measuring the length noise spectrum of sapphire over the frequency range between ten hertz and one kilohertz. The accessible frequency range may well extend below ten hertz, depending on the degree of improvement achieved in the intensity stabilisation system.

The instrument will then be able to begin a extensive series of tests with different test cavity materials, different test cavity geometries and different probe laser properties. Adjustment of these parameters will allow different noise mechanisms to become dominant and hence allow their measured spectra to be compared with predictions. A possible outcome of this work would be a comprehensive catalogue of length noise in various optical materials.

In the longer term, the system could be extended to form part of a quality control system in the construction of future gravitational wave interferometers. Each optical element could be tested to ensure that its noise properties was appropriate for its intended role.

Appendix A

Thermoelastic Noise Correction

The spectrum of thermoelastic noise is given by

$$S_{TE}(\omega) = \frac{8(1+\sigma)^2\alpha_l^2\kappa k_B T^2}{\sqrt{2\pi}C_v^2\rho^2 r_0^3} J(\Omega)$$

where $J(\Omega)$ is derived from an integral over the test mass in three dimensions. The resultant expression for $J(\Omega)$ given as equation 21 of Cerdonio et. al. [32] has a constant multiplier of $\sqrt{2/\pi}$;

$$J(\Omega) = \sqrt{\frac{2}{\pi}} \int_0^\infty du \int_{-\infty}^\infty dv \left\{ \frac{u^3 e^{\frac{-u^2}{2}}}{(u^2+v^2)[(u^2+v^2)^2+\Omega^2]} \right\}$$

This factor was later corrected by Numata et. al. in [14] to $\sqrt{2/\pi^3}$, but no explanation for the reason for the correction was given. In this appendix the appropriate expression is rederived to show the appearance of the extra factor of $1/\pi^2$.

We begin with expression 19 of the Cerdonio treatment from which $J(\Omega)$ is derived;

$$I = \frac{a^2}{(2\pi)^3} \int_{-\infty}^\infty \int_{-\infty}^\infty \int_{-\infty}^\infty dk_x.dk_y.dk_z \left\{ \frac{k_\perp^2 e^{-k_\perp^2 r_0^2/2}}{k^2(a^4k^4+\omega^2)} \right\} \tag{A.1}$$

$$\text{where } k^2 = k_\perp^2 + k_z^2$$

Let us make a change of variables so we can solve the integral in cylindrical co-

ordinates. Following the Cerdonio treatment we use $u = r_0 k_\perp$ to describe the radial dimension, $v = r_0 k_z$ to represent the z dimension and ϕ, the angular dimension. We need to find the appropriate Jacobian to make the transformation from $dk_x.dk_y.dk_z$ to $du.dv.d\phi$ so that we can substitute back into equation A.1.

Now the Jacobian is given by

$$\begin{aligned}
Jac &= \left|\frac{\partial(k_x, k_y, k_z)}{\partial(u, v, \phi)}\right| \\
&= \begin{vmatrix} \frac{\partial k_x}{\partial u} & \frac{\partial k_x}{\partial v} & \frac{\partial k_x}{\partial \phi} \\ \frac{\partial k_y}{\partial u} & \frac{\partial k_y}{\partial v} & \frac{\partial k_y}{\partial \phi} \\ \frac{\partial k_z}{\partial u} & \frac{\partial k_z}{\partial v} & \frac{\partial k_z}{\partial \phi} \end{vmatrix} \\
&= \begin{vmatrix} \frac{1}{r_0}\cos\phi & 0 & -\frac{u}{r_0}\sin\phi \\ \frac{1}{r_0}\sin\phi & 0 & \frac{u}{r_0}\cos\phi \\ 0 & \frac{1}{r_0} & 0 \end{vmatrix} \\
&= -\frac{1}{r_0}\begin{vmatrix} -\frac{u}{r_0}\sin\phi & \frac{1}{r_0}\cos\phi \\ \frac{u}{r_0}\cos\phi & \frac{1}{r_0}\sin\phi \end{vmatrix} \\
&= -\frac{1}{r_0}\left\{-\frac{u}{r_0^2}\sin^2\phi - \frac{u}{r_0^2}\cos^2\phi\right\} \\
Jac &= \frac{u}{r_0^3}
\end{aligned}$$

we can therefore write

$$\begin{aligned}
dk_x.dk_y.dk_z &= Jac \times du.dv.d\phi \\
&= \frac{u}{r_0^3}.du.dv.d\phi
\end{aligned}$$

We will also need an expression for $k^{2;}$

$$\begin{aligned}
k^2 &= k_\perp^2 + k_z^2 \\
&= \left(\frac{u}{r_0}\right)^2 + \left(\frac{v}{r_0}\right)^2 \\
&= \frac{u^2 + v^2}{r_0^2}
\end{aligned}$$

If we make the appropriate substitutions into A.1 we obtain

$$I = \frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}\frac{1}{r_0^3}u.du.dv.d\phi\left\{\frac{\left(\frac{u}{r_0}\right)^2 e^{-u^2/2}}{\frac{1}{r_0^2}\left(u^2+v^2\right)\left[\frac{a^4}{r_0^4}\left(u^2+v^2\right)^2+\omega^2\right]}\right\} \quad \text{(A.2)}$$

following Cerdonio we substitute for ω, using

$$\begin{aligned}\omega &= \Omega\omega_c \\ &= \Omega\frac{a^2}{r_0^2}\end{aligned}$$

$$\begin{aligned}\Rightarrow I &= \frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}\frac{1}{r_0^3}u.du.dv.d\phi\left\{\frac{\left(\frac{u}{r_0}\right)^2 e^{-u^2/2}}{\frac{1}{r_0^2}\left(u^2+v^2\right)\left[\frac{a^4}{r_0^4}\left(u^2+v^2\right)^2+\Omega^2\frac{a^4}{r_0^4}\right]}\right\} \\ &= \frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}\frac{1}{r_0^3}du.dv.d\phi\left\{\frac{u^3e^{-u^2/2}}{\left(u^2+v^2\right)\left[\frac{a^4}{r_0^4}\left(u^2+v^2\right)^2+\Omega^2\frac{a^4}{r_0^4}\right]}\right\} \\ &= \frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}du.dv.d\phi\left\{\frac{\left(\frac{1}{r_0^3}\right)}{\left(\frac{a^4}{r_0^4}\right)}\frac{u^3e^{\frac{-u^2}{2}}}{\left(u^2+v^2\right)\left[\left(u^2+v^2\right)+\Omega^2\right]}\right\} \\ &= \frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}du.dv.d\phi\left\{\frac{1}{\left(\frac{a^4}{r_0}\right)}\frac{u^3e^{\frac{-u^2}{2}}}{\left(u^2+v^2\right)\left[\left(u^2+v^2\right)+\Omega^2\right]}\right\} \\ &= \frac{r_0}{a^4}\frac{a^2}{(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}du.dv.d\phi\left\{\frac{u^3e^{\frac{-u^2}{2}}}{\left(u^2+v^2\right)\left[\left(u^2+v^2\right)+\Omega^2\right]}\right\} \\ &= \frac{r_0}{a^2(2\pi)^3}\int_0^{2\pi}\int_{-\infty}^{\infty}\int_0^{\infty}du.dv.d\phi\left\{\frac{u^3e^{\frac{-u^2}{2}}}{\left(u^2+v^2\right)\left[\left(u^2+v^2\right)+\Omega^2\right]}\right\}\end{aligned}$$

We now complete the integration with respect to ϕ

$$\Rightarrow I = \frac{r_0}{a^2(2\pi)^3}2\pi\int_{-\infty}^{\infty}\int_o^{\infty}du.dv\left\{\frac{u^3e^{\frac{-u^2}{2}}}{\left(u^2+v^2\right)\left[\left(u^2+v^2\right)+\Omega^2\right]}\right\}$$

Examination of the Cerdonio equations 18 and 20 allows us to write;

$$32\frac{r_0}{a^2\left(2\pi\right)^2} = \frac{8}{\sqrt{2\pi}}\frac{r_0}{a^2}J_0$$

where J_0 is the appropriate constant in the expression for J.

$$\begin{aligned} \Rightarrow \frac{8r_0}{a^2\pi^2} &= \frac{8}{\sqrt{2\pi}}\frac{r_0}{a^2}J_0 \\ J_0 &= \frac{\sqrt{2\pi}}{\pi^2} \\ J_0 &= \frac{\sqrt{2}}{\pi^{\frac{3}{2}}} \end{aligned}$$

This result is in agreement with the correction of Numata et. al. [14]

Appendix B

A Review of Control Theory

B.1 Improving System Behavior *via* Feedback

The inclusion of feedback in a system can modify the system behavior so that it is better behaved than the corresponding open loop system. Common uses for feedback include:

- Modifying plant behavior (eg. improving plant stability, resonance frequency and/or damping);
- Reducing the sensitivity of the plant to disturbances and noise;
- Ensuring that a plant follows a command signal with appropriate behavior.

In this work the most important motivation for the use of feedback control systems was the reduction of noise in various parts of the experiment. Consider a plant $G(s)$ which is affected by some noise process having a spectrum $G_n(s)$ as shown in figure B.1. Here s represents the Laplace transform variable. The output of this plant will be given by $Y(s) = G(s)R(s) + G_n(s)$. Notice that the output noise source $G_n(s)$ is appears directly at the output and can cause problems. The ability of this system to follow an input command $R(s)$ may not be adequate as it will be determined solely by the nature of the plant transfer function, $G(s)$.

Let us now enclose the plant $G(s)$ in a feedback loop as shown in figure B.2. In doing this we add a so-called *compensator*, having transfer function $C(s)$.

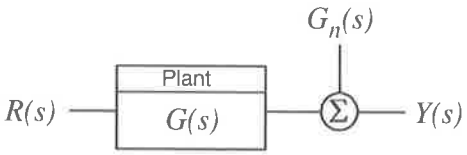

Figure B.1: An open-loop plant subject to a disturbance

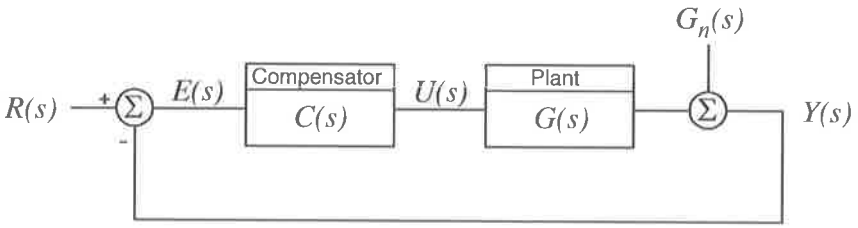

Figure B.2: A closed-loop plant subject to a disturbance

where
- $G(s)$ is the transfer function of the plant
- $C(s)$ is the compensator transfer function
- $G_n(s)$ is the plant disturbance
- $R(s)$ is the reference input
- $E(s)$ is the error signal
- $U(s)$ is the control signal
- $Y(s)$ is the plant output

We can now calculate the plant output $Y(s)$ in the closed loop case:

$$\begin{aligned}
Y(s) &= G(s)U(s) + G_n(s) \\
&= C(s)G(s)E(s) + G_n(s) \\
&= C(s)G(s)\{R(s) - Y(s)\} + G_n(s) \\
&= G(s)C(s)R(s) - G(s)C(s)Y(s) + G_n(s) \\
&= \frac{1}{1 + C(s)G(s)}\{C(s)G(s)R(s) + G_n(s)\} \\
&= \frac{C(s)G(s)}{1 + C(s)G(s)}R(s) + \frac{1}{1 + C(s)G(s)}G_n(s) \\
&= T(s)R(s) + S(s)G_n(s) \qquad \text{(B.1)}
\end{aligned}$$

The transfer functions $S(s)$ and $T(s)$ given in equation B.1 are given special names; the *Sensitivity* and *Complementary Sensitivity* functions respectively. They are defined as shown in equations B.2 and B.3 below. It should be noticed that the two transfer functions sum to unity, thus if we wish to lower $S(s)$ we must raise $T(s)$ and *vice versa*. This is the source of the term "complementary" in the name for $T(s)$.

$$S(s) = \frac{1}{1 + C(s)G(s)} = \frac{Y(s)}{G_n(s)} \tag{B.2}$$

$$T(s) = \frac{C(s)G(s)}{1 + C(s)G(s)} = \frac{Y(s)}{R(s)} \tag{B.3}$$

$$S(s) + T(s) = 1 \tag{B.4}$$

The sensitivity function describes the effectiveness of a system in rejecting disturbances to the plant. Ideally $S(s)$ would be zero at all frequencies, but this is not practical because of stability concerns (see section B.2). We must instead aim to make $S(s)$ small in the frequency band where $G_n(s)$ has significant energy. We can make $S(s)$ small by making the product $C(s)G(s)$ large at a given frequency. As we cannot change $G(s)$, the sensitivity function must be managed by careful design of the compensator transfer function, $C(s)$. This requires that $C(s)$ be made large over the frequency range where we need significant plant noise reduction. Good disturbance rejection thus places a lower bound on the required system bandwidth.

The complementary sensitivity function describes the relationship between the system output $Y(s)$ and its input $R(s)$ - it is simply the closed loop transfer function of the loop. As the complementary sensitivity function describes the system's ability to track an input command, we would like it to be unity at all frequencies where $R(s)$ has important behavior. This too will set a lower limit on the system bandwidth.

Because $S(s)$ and $T(s)$ sum to one it is clear that we cannot design $C(s)$ with only one of them in mind; $C(s)$ must be designed to optimise both command tracking and disturbance rejection performance. This can be problematic as there are competing

considerations in our choice of $C(s)$. Because we have only one block to tailor system performance, the control scheme shown in figure B.2 is known as a *one-degree-of-freedom* controller. We can, however, extend the structure to obtain a so-called *two-degree-of-freedom* (2-dof) controller. There are several possible topologies for a 2-dof controller [120], [96] but all have the defining characteristic that they contain different controllers in the path from $Y(s)$ to $E(s)$ and from $R(s)$ to $E(s)$. As a representative example we consider the topology shown in figure B.3 where we have extended the topology of figure B.2 by adding an additional transfer function $H(s)$ in the feedback path.

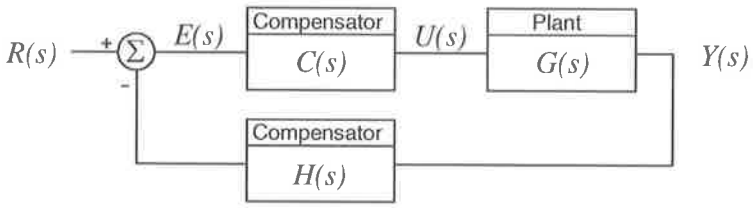

Figure B.3: Two-degree-of-freedom controller topology

It is straightforward to show that this topology yields a system whose behavior is governed by the equation:

$$Y(s) = \frac{C(s)G(s)}{1 + C(s)G(s)H(s)}R(s) + \frac{1}{1 + C(s)G(s)H(s)}G_n(s) \tag{B.5}$$

In this configuration we can independently tailor $S(s)$ and $T(s)$ because we have two transfer functions, $C(s)$ and $H(s)$, available to do the shaping. If we again consider the sensitivity and complementary sensitivity functions we see that in this case, we see that

$$H(s)T(s) + S(s) = 1 \tag{B.6}$$

The presence of $H(s)$ in this equation reflects the additional design freedom in this controller.

B.1.1 Noise Performance of a General Control Loop

In general, every component in the control loop contributes noise to the system output. Each noise source in the loop will make a contribution with a distinct transfer function that depends on its point of entry into the system. We expand the model discussed above to include additional noise components as shown in figure B.4 and examine the propagation of the noise through the system.

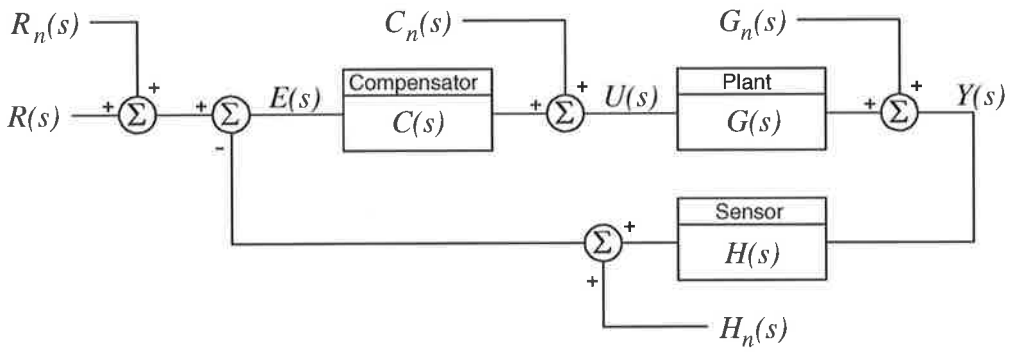

Figure B.4: A control system with all noise sources shown

where
- $G_n(s)$ is the plant noise
- $C_n(s)$ is the compensator noise
- $H_n(s)$ is the sensor noise
- $R_n(s)$ is the reference noise

All noise sources are referred to the output of their corresponding block.

$$
\begin{aligned}
Y(s) &= G(s)U(s) + G_n(s) \\
&= C(s)G(s)E(s) + G(s)C_n(s) + G_n(s) \\
&= C(s)G(s)\{R(s) - H(s)Y(s) + R_n(s) - H_n(s)\} + G(s)C_n(s) + G_n(s) \\
&= C(s)G(s)R(s) - C(s)G(s)H(s)Y(s) \\
&\quad + C(s)G(s)\{R_n(s) - H_n(s)\} + G(s)C_n(s) + G_n(s) \\
&= \frac{C(s)G(s)}{1 + C(s)G(s)H(s)}R(s) + \frac{C(s)G(s)}{1 + C(s)G(s)H(s)}\{R_n(s) - H_n(s)\} \\
&\quad + \frac{G(s)}{1 + C(s)G(s)H(s)}C_n(s) + \frac{1}{1 + C(s)G(s)H(s)}G_n(s) \\
&= T(s)R(s) + T(s)R_n(s) + T(s)H_n(s) + S_i(s)C_n(s) + S(s)G_n(s) \qquad \text{(B.7)}
\end{aligned}
$$

Here we have introduced a new special transfer function $S_i(s)$ which is known as the *Input Sensitivity* function. The input sensitivity function describes the relationship between the output and the plant input noise and is defined by

$$S_i(s) = \frac{G(s)}{1 + C(s)G(s)H(s)} \qquad \text{(B.8)}$$

If we consider the frequency region where $C(s)G(s)H(s)$ is much larger than one, we can simplify equation B.7 to

$$Y(s) = \frac{1}{H(s)}R(s) + \frac{1}{H(s)}\{R_n(s) - H_n(s)\} + \frac{1}{C(s)H(s)}C_n(s) + \frac{1}{C(s)G(s)H(s)}G_n(S) \qquad \text{(B.9)}$$

This equation shows that the noise sources $R_n(s)$ and$H_n(s)$ are treated in exactly the same way as the reference signal $R(s)$, which is as we would expect from their positions in diagram B.4. In particular we note that making the compensator gain large ($|C(s)|$ large) does not reduce the effect of these noise sources as it does for $G_n(s)$ and $C_n(s)$. Thus for the best noise performance of the overall system we need to take greatest care in the design of the reference source and in the sensor.

Equation B.9 suggests that the optimal method for minimising system noise is to make $C(s)$ and $H(s)$ as large as possible at all frequencies. However, in terms of noise reduction, there is little to be gained from making $C(s)$ and $H(s)$ large over a bandwidth greater than that in which $G_n(s)$ has significant energy. Keeping the bandwidth of the system small has the advantage of attenuating the high frequency components of sensor and reference noise, which is useful in avoiding overloading and saturation effects in the servo system. The design goal from a noise perspective is therefore to make the gain as high as possible over the range of frequencies where loop disturbances have significant amplitude. This is equivalent to making the sensitivity small as possible only where there is noise to be suppressed.

B.2 Control Loop Stability

Although we have seen that the application of negative feedback to a plant can improve the properties of the plant, the feedback cannot be applied arbitrarily. In particular, phase shifts in the control loop can reach a point where the desired negative feedback becomes positive feedback. This usually leads to instability in the closed loop system. Stability can be maintained by ensuring that the phase delay in a system does not reach 180°while the system still has gain greater than unity. (For simplicity, we ignore marginally stable system for the purpose of this discussion.)

Various graphical tools have been developed that allow control systems to be easy visualised and thus let instability be predicted. The better known displays of this kind are the Bode plot, Nyquist and Nichols plots and the root locus. Bode plots and Nichols plots are used in this thesis and will therefore be discussed in more detail below. The treatment here is necessarily brief; additional information may be found in any introductory control theory text [97, 121, 122].

B.2.1 Bode plots

Bode plots are a graphical technique for displaying the frequency response of a control system. The Bode plot consists of a plot of system gain (in dB) against frequency and a plot of system phase against frequency. Visual inspection of the plots

allows the stability of the system to be determined. If the gain is more than unity when the phase curve crosses -180° then the system will be closed loop unstable.

As an example, consider the two systems described by the similar transfer functions

$$G_1(s) = \frac{10}{(s+1)(s^2+0.2S+1)}$$

and

$$G_2(s) = \frac{1}{(s+1)(s^2+0.2S+1)}$$

The two systems have the same dynamics, but the dc gain of G_1 is higher by a factor of ten.

Bode plots of the two systems can be seen in figure B.5. The frequency at which the phase response is -180° is indicated by the dotted vertical line. At this frequency, G_1 is seen to have a gain greater than zero, indicating an unstable closed loop response. The curve for G_2 has a gain lower than zero at this frequency and will therefore be closed loop stable.

B.2.2 Nichols plot

A Nichols plot directly shows the gain/phase relationship of a system, but does not explicitly display the frequency dependence. The Nichols chart is very similar to the more common Nyquist plot except that the Nichols chart uses a Cartesian axes and a logarithmic gain scale, which allows a larger gain range to be visualised. Nichols plots for stable and unstable systems are shown in figure B.5. As can be seen the stability of a system can be easily determined from the plots and the gain and phase margins extracted. Stability requires that the curve on the Nichols chart have no downwards crossings of the dashed line shown in the figure.

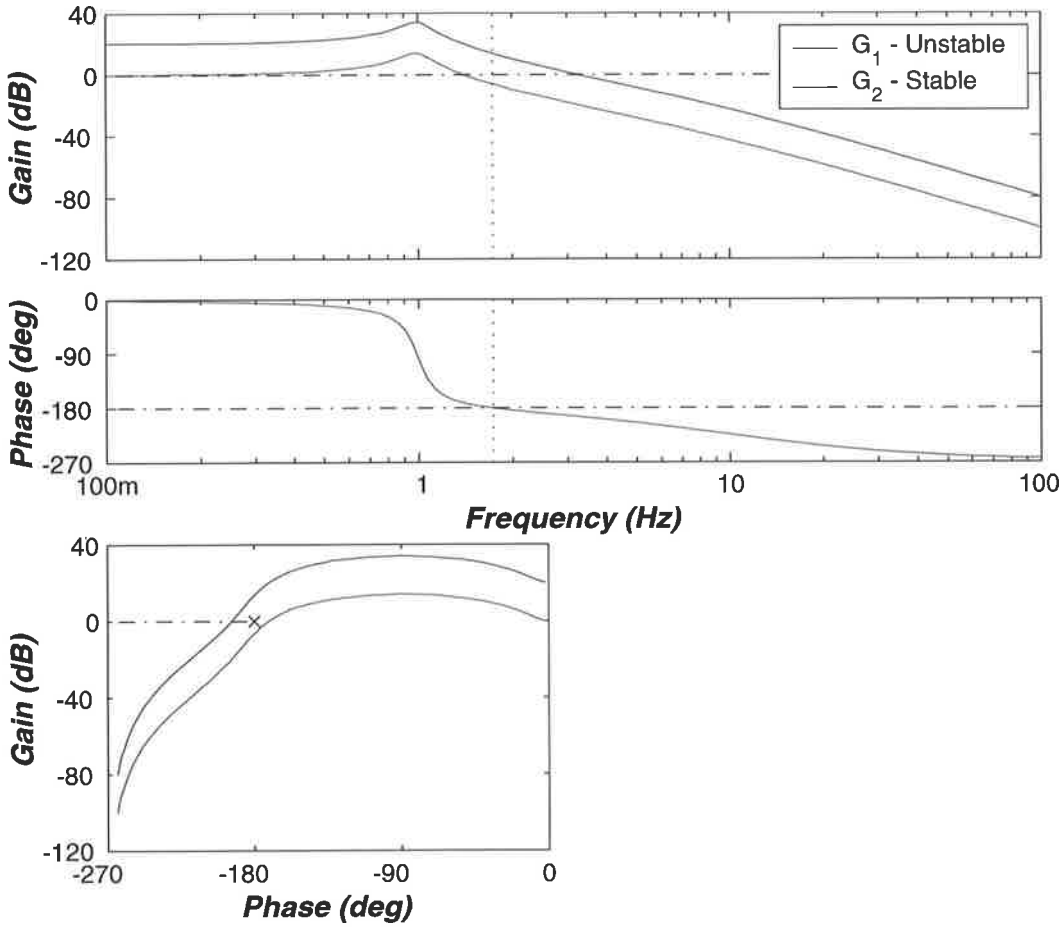

Figure B.5: Bode and Nichols plots of two transfer functions, a stable system shown in green and an unstable system in blue.

B.2.3 Stabilisation of an Unstable System

The addition of a compensator to a system allows the loop transfer function to be modified before it is enclosed in a feedback loop. If analysis shows that a plant will be closed loop unstable, then the addition of the compensator transfer function allows the open loop transfer function to be changed to a more appropriate form. In the simplest case, such as illustrated in figure B.5, this can be as easy as modifying the dc gain. In more general cases though we must tailor the transfer function of the compensated system to improve its behaviour in the vicinity of its unity gain point.

The requirement to tailor the system response at its unity gain frequency is somewhat in conflict with producing optimal performance in a control loop. As we saw in section B.1 above, we would like to keep compensator gain high at all frequencies to produce good closed loop performance. However, stability requirements require that we reduce compensator gain in a controlled manner so that it behaves well near the unity gain frequency. In contrast to the performance requirements discussed in section B.1, stability requirements place an upper limit on system bandwidth. Producing adequate closed loop performance while preserving system stability is the trade-off at the heart of every servo system design problem.

B.3 Conclusion

In this appendix we have discussed the uses of control engineering in improving the performance of a plant. In the first section we analysed a general servo system and found that good noise rejection requires a high compensator gain at frequencies where there is significant system noise. We then briefly discussed the requirements for maintaining stability in a servo system and found that the behaviour of a servo system transfer function must be carefully controlled in the region around its unity gain point. We thus showed that there is a trade-off between designing a system to be stable and achieving optimal performance.

Appendix C

The Vacuum System

C.1 Introduction

Gas pressure fluctuations and acoustic vibration can cause significant difficulties in high precision optical experiments (see chapter 2). As described in section 2.3.4.6, the frequency noise resulting from gas pressure fluctuations is sufficient to prevent the detection of the very small length noise sources important to gravitational wave detectors. A vacuum system suitable for the final experiment was therefore constructed as part of the project. At the time of writing the vacuum system was ready for installation of the experiment.

As was described in section 2.3.4.6, the ultimate pressure requirements for the vacuum system are relatively modest. The system far exceeded the one torr requirement and achieved an ultimate pressure of $3x10^{-6}$ torr. Cleanliness of the vacuum system was important so that optical components placed inside the chamber were not contaminated with vacuum oils; this requirement was critical because the finesse of high finesse cavities can be significantly degraded by even a small amount of surface contamination.

The production of an appropriate vacuum environment was simplified by reusing components of a pre-existing vacuum monochromator system that once occupied the laboratory [49]. Nevertheless, the salvaged system was badly degraded at the time this thesis was begun and many features of the previous system were inappropriate for the new experiment.

C.2 Vacuum Chamber

A large, cylindrical mild steel vacuum chamber which formed part of an old monochromator system was used to hold the experiment. The chamber was already fitted with a large number of electrical feedthroughs, two mechanical feedthroughs to the preisolator slab described in section 3.2 and an easily removable end cap that allowed access. The tank was originally equipped with two glass viewing ports; one of these was modified to serve as the pumping port in our experiment. The chamber had a length of 1370 mm and a diameter of 910 mm, a total volume of approximately 3.6 m^3. The working height of the vacuum system was significantly reduced by the design of the preisolator feedthroughs and the need to place additional structures on top of these to support an optical table. The final working height of the vacuum system was about 700 mm, a constraint that limited the possible height of the seismic isolation system.

A photograph of the inside of the chamber is shown in figure C.1. The optical table was not installed in the tank in this photograph so that the feedthrough structures can be seen. Each of the feedthroughs is fitted with a large stainless steel pier which was designed to support the optical table. Actual mounting of the table is through three ball-mounted brass pads which locate in the piers. To prevent disturbing the table it will be tied to the pier by two large bolts when work is being conducted inside the chamber. The electrical feedthrough port can be seen in the upper right of the picture (slightly rust coloured in this photograph). This port provides 214 electrical feedthroughs into the vacuum system. The photograph also shows two other ports half way down the length of the tank; the left port is a viewing port, the right port is used to attach the pumps.

C.3 Pumping System

A schematic of the vacuum system is shown in figure C.2. All valves in the system are electrically controlled, pneumatically actuated valves that are controlled from a central control rack. The large gate valves, V3 and V4, are large reconditioned gate valves from the monochromator system; all other valves were new to the system,

Figure C.1: Photograph of the inside of the vacuum chamber. Isolator feedthroughs can be seen on the floor of the tank. Pumping is through the right port, midway up the sie ofthe tank. The external optical table and suspension tower are visible in the background.

some having been reconditioned from other obsolete experiments.

The roughing pump, P1, is a large mechanical pump that was salvaged from the monochromator system. This pump had a very high throughput, being able to pump the system to 100 $mTorr$ in about 10 minutes. The pump is water cooled and housed in a lead-lined room in order to reduce acoustic noise and vibration. The smaller backing pump, P2, is a smaller rotary pump which is used to back the 6" oil diffusion pump (P3) when the chamber is at operating pressure. Also shown in the schematic is a two-stage oil trap that prevents back-streaming from the diffusion pump. The trap includes a water cooled optical baffle and a liquid air trap.

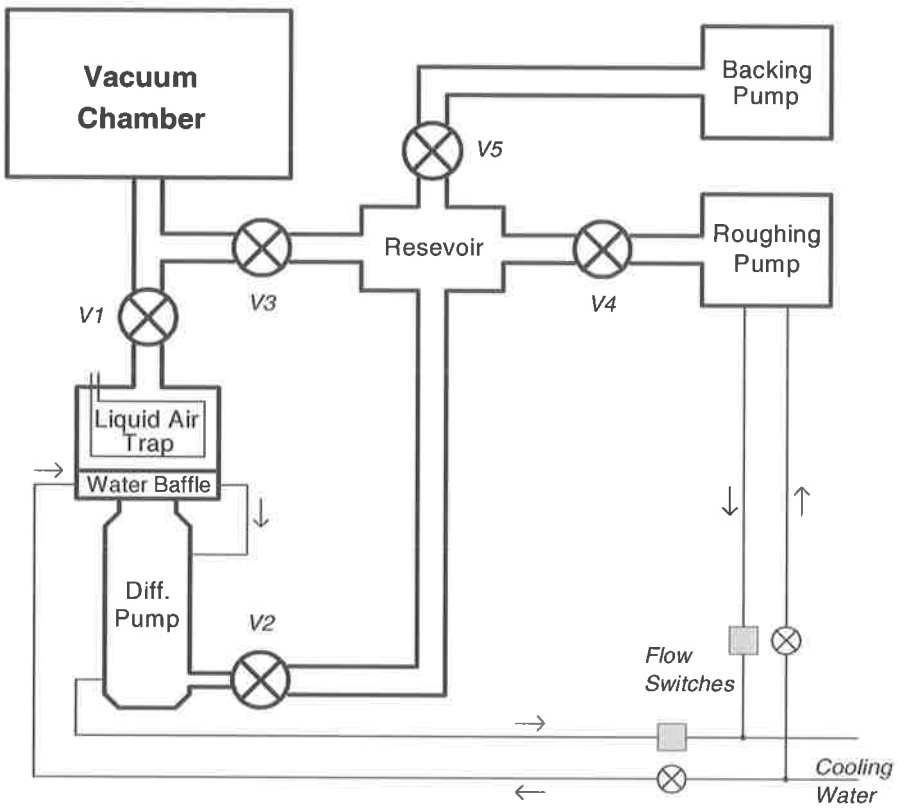

Figure C.2: Schematic of the vacuum system.

Oil molecules striking the walls of either trap are cooled through conduction and are likely to stick to the surfaces, preventing further progress towards the vacuum chamber.

Photographs of the three pumps can be seen in figures C.3, C.4 and C.5.

While the preisolator provides some measure of decoupling between the experiment and mechanical vibration of the floor, it was thought prudent to minimise any vibrations produced by mechanical pumps. The large, noisy roughing pump, P1, is used only for initial roughing of the chamber. The second mechanical pump, P2, is used when the system is at operating pressure and the diffusion pump is operational. Vibrations from the smaller pump are reduced by mounting it on a triple stack of concrete blocks separated by rubber sheeting. All noise spectra measured for this thesis were unaffected by running the small backing pump running. For the ultra-quiet experiments which will be conducted in the next phase of this experiment, it

Figure C.3: Photograph of the roughing pump. Pumping is through the green striped black rubber pipe at the top of the picture. The pump exhaust is vented through the large black pipe on the left of the photo to a vapour trap and then expelled from the pump-room by an exhaust fan.

is possible to run the system for a limited time with no mechanical pumps. In this mode the diffusion pump exhaust is simply stored in the reservoir formed by the system's large diameter roughing line. The system can remain in this mode until the pressure in the reservoir rises to the stall pressure of the diffusion pump. In practice this provides approximately ten minutes of ultra quiet operation, a time which is expected to be adequate for data acquisition.

Figure C.4: Photo showing the small backing pump n the trench at the bottom-left of the picture.

Figure C.5: Photograph of the vacuum chamber, oil trap and diffusion pump. The large silver object in the centre of the frame is the oil trap. The diffusion pump sits beneath it in the grey frame. Both the backing and roughing pumps are attached to the large copper pipe visible in the right of the frame.

Appendix D

The Pound Drever Hall Error Signal

D.1 Derivation of the General PDH Response

Consider the field of the input laser beam,

$$\begin{aligned} E_{inc} &= Ae^{j(\omega_c t+\phi)} \\ &= Ae^{j\omega_c t} \times e^{j\phi} \end{aligned}$$

We now include the effect the phase modulation $\phi = \beta e^{j\omega_m t}$, getting

$$E_{inc} = A \times e^{j\omega_c t} \times e^{j\beta e^{j\omega_m t}}$$

We now extract the generalised phasor $\tilde{E}_{inc}$, where $E_{inc} = \tilde{E}_{inc} e^{j\omega_m t}$

$$\begin{aligned} \tilde{E}_{inc} &= Ae^{j\beta e^{j\omega_m t}} \\ &\approx A\{J_o(\beta) + J_1(\beta)e^{j\omega_m} - J_1(\beta)e^{-j\omega_m t}\} \end{aligned}$$

If we consider the reflection of a laser beam which is offset from the cavity centre

by ψ, then we can write the resultant reflected field as

$$\begin{aligned}\tilde{E}_{refl} &= \Gamma\tilde{E}_{inc} \\ &= \Gamma(\psi)AJ_o(\beta) + \Gamma(\psi+\omega_m)AJ_1(\beta)e^{j\omega_m t} - \Gamma(\psi-\omega_m)AJ_1(\beta)e^{-j\omega_m t}\end{aligned}$$

Next we examine this field with a photodetector. A photodetector is sensitive to incident power, not electric field. The detected signal will thus be given by

$$V_{pd} = G_{pd}\tilde{E}_{refl}\tilde{E}^*_{refl}$$

where G_{pd} is the photodetector gain and * represents complex conjugation. Substituting the expression for $\tilde{E}_{refl}$ we obtain,

$$\begin{aligned}V_{pd} = A^2G_{pd} \times [\quad & \Gamma(\psi)J_o(\beta) \\ + \quad & \Gamma(\psi+\omega_m)J_1(\beta)e^{j\omega_m t} - \Gamma(\psi-\omega_m)J_1(\beta)e^{-j\omega_m t}] \\ \times[\quad & \Gamma(\psi)J_o(\beta) + \\ & \Gamma(\psi+\omega_m)J_1(\beta)e^{j\omega_m t} - \Gamma(\psi-\omega_m)J_1(\beta)e^{-j\omega_m t}]^* \\ = A^2G_{pd}\times \quad & [\Gamma(\psi)J_o(\beta) \\ + \quad & \Gamma(\psi+\omega_m)J_1(\beta)e^{j\omega_m t} - \Gamma(\psi-\omega_m)J_1(\beta)e^{-j\omega_m t}] \\ \times \quad & [\Gamma^*(\psi)J_o(\beta) \\ + \quad & \Gamma^*(\psi+\omega_m)J_1(\beta)e^{-j\omega_m t} - \Gamma^*(\psi-\omega_m)J_1(\beta)e^{j\omega_m t}]\end{aligned}$$

Performing the multiplication we obtain;

$$
\begin{aligned}
V_{pd} = A^2 G_{pd} \times \{ \quad & \Gamma(\psi)\Gamma^*(\psi)J_o(\beta)^2 \\
+ \quad & \Gamma(\psi)\Gamma^*(\psi+\omega_m)J_o(\beta)J_1(\beta)e^{-j\omega_m t} \\
- \quad & \Gamma^*(\psi)\Gamma(\psi-\omega_m)J_o(\beta)J_1(\beta)e^{-j\omega_m t} \\
- \quad & \Gamma(\psi)\Gamma^*(\psi-\omega_m)J_o(\beta)J_1(\beta)e^{j\omega_m t} \\
+ \quad & \Gamma^*(\psi)\Gamma(\psi+\omega_m)J_o(\beta)J_1(\beta)e^{j\omega_m t} \\
- \quad & \Gamma(\psi+\omega_m)\Gamma^*(\psi-\omega_m)J_1(\beta)^2 e^{j2\omega_m t} \\
- \quad & \Gamma(\psi-\omega_m)\Gamma^*(\psi+\omega_m)J_1(\beta)^2 e^{-j2\omega_m t} \\
+ \quad & \Gamma(\psi+\omega_m)\Gamma^*(\psi+\omega_m)J_1(\beta)^2 \\
+ \quad & \Gamma(\psi-\omega_m)\Gamma^*(\psi-\omega_m)J_1(\beta)^2 \quad \}
\end{aligned}
$$

The first and last two lines of the above expression contain dc terms. These terms arise from mixing of the carrier wave with itself and of each sideband with itself. Because there is no frequency difference between a field and itself the terms appear at zero frequency. The second through fifth terms correspond to mixing between the carrier and the sidebands and therefore vary at the modulation frequency ω_m. The sixth and seventh terms correspond to terms derived from interaction between the two sidebands and therefore vary at twice the modulation frequency, $2\omega_m$.

We can simplify this expression if we consider that the power in the carrier is $P_c = A^2 J_o(\beta)^2$ and the power in each sideband is given by $P_s = A^2 J_1(\beta)^2$. In addition we can write $\Gamma(x)\Gamma^*(x) = |\Gamma(x)|^2$ and thus obtain;

$$
\begin{aligned}
V_{pd} = G_{pd} \quad & \times\{P_c|\Gamma(\psi)|^2 + P_s|\Gamma(\psi+\omega_m)|^2 + P_s|\Gamma(\psi-\omega_m)|^2 \\
& +\sqrt{P_c P_s}[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m)]e^{-j\omega_m t} \\
& +\sqrt{P_c P_s}[\Gamma^*(\psi)\Gamma(\psi+\omega_m) - \Gamma(\psi)\Gamma^*(\psi-\omega_m)]e^{j\omega_m t} \\
& -P_s\Gamma(\psi+\omega_m)\Gamma^*(\psi-\omega_m)e^{j2\omega_m t} \\
& -P_s\Gamma(\psi-\omega_m)\Gamma^*(\psi+\omega_m)e^{-j2\omega_m t}\}
\end{aligned}
$$

Let us extract the part of the photodetector signal which varies at frequency $\pm\omega_m$ and expand the complex exponentials to $e^{\pm j\omega_m t} = cos(\omega_m t) \pm j sin(\omega_m t)$. This procedure allows us to extract the real part of the photodetector signal at the modulation frequency.

$$\begin{aligned} V_{pd}|_{\pm\omega_m} &= G_{pd}\sqrt{P_c P_s}\{ \\ &\Re e[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m) \\ &+\Gamma^*(\psi)\Gamma(\psi+\omega_m) - \Gamma(\psi)\Gamma^*(\psi-\omega_m)]cos(\omega_m t) \\ &-\Im m[-\Gamma(\psi)\Gamma^*(\psi+\omega_m) + \Gamma^*(\psi)\Gamma(\psi-\omega_m) \\ &+\Gamma^*(\psi)\Gamma(\psi+\omega_m) - \Gamma(\psi)\Gamma^*(\psi-\omega_m)]sin(\omega_m t)\} \end{aligned}$$

$$\begin{aligned} V_{pd}|_{\pm\omega_m} &= G_{pd}\sqrt{P_c P_s}\{ \\ &\Re e[\Gamma(\psi)\Gamma^*(\psi+\omega_m) + \Gamma^*(\psi)\Gamma(\psi+\omega_m) \\ &-\Gamma^*(\psi)\Gamma(\psi-\omega_m) - \Gamma(\psi)\Gamma^*(\psi-\omega_m)]cos(\omega_m t) \\ &-\Im m[-\Gamma(\psi)\Gamma^*(\psi+\omega_m) + \Gamma^*(\psi)\Gamma(\psi+\omega_m) \\ &+\Gamma^*(\psi)\Gamma(\psi-\omega_m)) - \Gamma(\psi)\Gamma^*(\psi-\omega_m)]sin(\omega_m t)\} \end{aligned}$$

But, $\Gamma(A)\Gamma^*(B) = \Gamma^*(A)\Gamma(B)$, so

$$\begin{aligned} V_{pd}|_{\pm\omega_m} &= 2G_{pd}\sqrt{P_c P_s}\{ \\ &\Re e[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m)]cos(\omega_m t) \\ &-\Im m[-\Gamma(\psi)\Gamma^*(\psi+\omega_m) + \Gamma^*(\psi)\Gamma(\psi-\omega_m))]sin(\omega_m t)\} \end{aligned}$$

$$\begin{aligned} &= 2G_{pd}\sqrt{P_c P_s}\{ \\ &\Re e[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m)]cos(\omega_m t) \\ &+\Im m[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m))]sin(\omega_m t)\} \end{aligned}$$

When we feed this signal into a mixer we extract the component which is proportional to $sin(\omega_m t)$. The mixer is followed with a low pass filter which removes the mixer products at ω_m and $2\omega_m$ and we finally obtain the general PDH error signal:

$$V_{mixer} = 2G_{mix}G_{pd}\sqrt{P_c P_s} \times \Im m[\Gamma(\psi)\Gamma^*(\psi+\omega_m) - \Gamma^*(\psi)\Gamma(\psi-\omega_m)] \tag{D.1}$$

D.2 The Linear Region

Near resonance we can consider the sidebands to be completely reflected from the input coupler, ie. $\Gamma(\psi-\omega_m) = \Gamma(\psi+\omega_m) = -1$. Therefore we can write

$$V_{mixer} = -4G_{mix}G_{pd}\sqrt{P_c P_s}\Im m[\Gamma(\psi)] \tag{D.2}$$

It can be shown [88] that Γ can be given by the expression

$$\Gamma(\phi) = \frac{\sqrt{R}(e^{j\phi}-1)}{1-Re^{j\phi}}$$

where $\begin{cases} R \text{ is the power reflection coefficient of each mirror} \\ \phi \text{ is the phase shift in one cavity round trip} \end{cases}$

Thus

$$\Gamma(\phi) = \frac{\sqrt{R}[\cos(\phi)+j\sin(\phi)-1]}{1-R[\cos(\phi)+j\sin(\phi)]}$$

For small ϕ we can write $cos(\phi) = 1$ and $sin(\phi) = \phi$.

$$\begin{aligned} \Rightarrow \Gamma(\phi) &= \frac{\sqrt{R}[1+j\phi-1]}{1-R[1+\phi]} \\ &= \frac{\sqrt{R}}{1-R}j\phi \end{aligned} \tag{D.3}$$

We now need to find an expression for ϕ. We know that the phase shift picked up in a single round trip around the cavity can be expressed as;

$$\begin{aligned}\phi &= 2\pi\frac{2L}{\lambda} \\ &= 4\pi\frac{L}{\lambda}\end{aligned}$$

If we consider a small variation in cavity length δx,we can write

$$\phi = 4\pi\frac{L+\delta x}{\lambda}$$

However near resonance we can assume that L is close to a multiple of $\lambda/2$ and thus

$$\phi = n2\pi + 4\pi\frac{\delta x}{\lambda}$$

for some integer n.

Substituting into equation D.3 we obtain an estimate for Γ :

$$\Gamma(\phi) = \frac{\sqrt{R}}{1-R}j\left(4\pi\frac{\delta x}{\lambda}\right)$$

But $\frac{\sqrt{R}}{1-R}$ is defined to be the cavity finesse F. We make the appropriate substitution;

$$\begin{aligned}\Gamma(\phi) &= F\left(j4\pi\frac{\delta x}{\lambda}\right) \\ &= j4F\pi\frac{\delta x}{\lambda} \qquad \text{(D.4)}\end{aligned}$$

Substituting back into D.2 we obtain the response in the linear region near resonance.

$$\begin{aligned} V_{mixer} &= -4G_{mix}G_{pd}\sqrt{P_cP_s}\Im m[j4F\pi\frac{\delta x}{\lambda}] \\ &= -4G_{mix}G_{pd}\sqrt{P_cP_s} \times 4F\pi\frac{\delta x}{\lambda} \\ &= -16\pi G_{mix}G_{pd}\sqrt{P_cP_s}F\frac{\delta x}{\lambda} \end{aligned}$$

Dividing through by the length change we obtain the slope of the error signal in the linear region.

$$\begin{aligned} D &= \frac{V_{mixer}}{x} \\ &= -16\pi G_{mix}G_{pd}\sqrt{P_cP_s}\frac{F}{\lambda} \\ D &= D_o\sqrt{P_cP_s}\frac{F}{\lambda}, \end{aligned} \tag{D.5}$$

where D_o describes the contributions of all of the constants in the equation. We can now convert to an equivalent slope of voltage over frequency.

$$H = \frac{V_{mixer}}{\delta f}$$

Now

$$\begin{aligned} \delta f &= \frac{c}{\lambda L}x \\ \Longrightarrow H &= \frac{\lambda L}{c}D \\ &= \frac{\lambda L}{c}D_o\sqrt{P_cP_s}\frac{F}{\lambda} \\ &= \frac{D_o\sqrt{P_cP_s}FL}{c} \\ H &= H_o\sqrt{P_cP_s}FL, \end{aligned} \tag{D.6}$$

where H_o describes the contributions of all of the constants in the equation.

D.3 Shot Noise

Let us consider the variation in optical power falling upon the detector. The variation produced by shot noise in the detector is given by

$$\begin{aligned} S_{P,sn} &= \sqrt{2}\sqrt{2E_{photon}P} \\ &= \sqrt{4h\nu P} \end{aligned}$$

Near resonance only the two sidebands fall on the detector. Hence,

$$\begin{aligned} S_{P,sn} &= \sqrt{4h\nu 2P_s} \\ &= \sqrt{8h\nu P_s} \end{aligned}$$

we know that the voltage output resulting from this signal will be given by

$$\begin{aligned} S_{V,sn} &= G_{mix}G_{pd}S_{P,sn} \\ &= G_{mix}G_{pd}\sqrt{8h\nu P_s} \end{aligned}$$

Dividing by the PDH sensor gain we can find the equivalent frequency noise of this level.

$$\begin{aligned} S_{f,sn} &= \frac{G_{mix}G_{pd}}{H_o\sqrt{P_cP_s}FL}\sqrt{8h\nu P_s} \\ &= \frac{1}{H_oFL\sqrt{P_c}}\sqrt{8h\nu} \end{aligned} \tag{D.7}$$

But

$$\begin{aligned} H_o &= \frac{D_o}{c} \\ &= -\frac{16G_{mix}G_{pd}}{c} \end{aligned}$$

Substituting back into D.7 we find

$$\begin{aligned} S_{f,sn} &= \frac{G_{mix}G_{pd}}{16G_{mix}G_{pd}FL\sqrt{P_c}}\sqrt{8h\nu} \\ &= \frac{c}{FL}\sqrt{\frac{h\nu}{8P_c}} \\ S_{sn} &= \frac{1}{FL}\sqrt{\frac{hc^3}{8\lambda P_c}}. \end{aligned} \tag{D.8}$$

Equation D.8 gives the best noise performance that a PDH frequency discriminator can produce.

This equation can be compared with the equivalent limit for fringe side locking;

$$S_{sn} = \frac{1}{FL}\sqrt{\frac{2hc^3}{\lambda P}} \tag{D.9}$$

The ratio between these two techniques is thus given by

$$\begin{aligned} \frac{FSL}{PDH} &= \sqrt{\frac{2P}{8P_c}} \\ &= \sqrt{\frac{1}{4}\frac{P}{J_o^2(\beta)P}} \end{aligned}$$

A PDH discriminator will therefore have a shot noise limit which is $\sqrt{4J_o^2(\beta)}$ times better than a fringe side locked system detecting the same amount of power.

Appendix E

Electronic Schematics

E.1 Introduction

This appendix discusses the many circuits that have appeared throughout this thesis. These circuits were all designed and built by the author and, except for where noted in the photodetector section, represent original work.

The chapter opens with a general discussion of the physical construction of the electronic subsystems, including their construction and location in the laboratory. The operation of each circuit is then discussed in detail. First, the design of subrack power supplies is discussed in section E.6. The photodetectors used throughout the experiment are then presented, along with a high power variant used in fringe side locking (E.4). These photodetectors produce differential signals which are then transmitted from the optics table back to the electronics rack. Once there, the differential signals are converted back to single-ended signals by the "input card", which is discussed in section E.5. Section E.7 describes the general design of the three compensator boards used for servo loop implementation in the experiment. All three compensators used the same basic circuit, which were customised for each loop by appropriate component choice. A high current amplifier used to drive heaters is then discussed and the appendix concludes with discussion of two high voltage amplifier designs.

E.2 Overview

The electronic systems used in this experiment were all built on standard 100 mm $\times$ 160 mm double-sided Eurocards. The majority of the circuits are then housed in two 3U subracks in a standard 17" rack. The photograph shown in figure E.1 shows the physical layout of the various electronic components in the system. The top rack in the photograph holds the low power circuits, namely the input board (see sectionE.5), PDH processing board (section E.6) and various compensators (sectionE.7) along with a card which extracted signals from the subrack backplane for monitoring. The low power subrack includes its own mains-powered power supply (section E.3), located on the extreme right of the subrack. These power supply rails are supplied to each card through the wiring chassis at the rear of the subrack.

Mains-powered circuits are located in the bottom sub-rack. Mains power is supplied to each of these cards through wiring at the rear of the subrack. The high power subrack holds the RF amplifier (section E.6) and heater driver (section E.8) circuits. Like the top subrack, it too has a power supply to supply low voltage supply rails to each card through connectors at the rear of each card (section E.3).

A "patch-board" was built and placed immediately above the top subrack. This board allows signals to be easily connected through to the optical section of the experiment. Eight 8-pin DIN connectors were available for the connection of photodetectors, although only four were used for this experiment. The eight connections were divided into two sets and each set was bussed into a single 40-core cable for the 8 m run to the optical table. At the other end of the cable the lines were again split into individual 8 pin DIN sockets for connection to the photodetectors. The patch board also included eight BNC connectors for the connection of high frequency signals, including the ac output signals from the various photodetectors E.4 and the drive signal for the EOM E.6.

E.3 Power Supplies

Each subrack is equipped with a mains-powered linear power supply. These power supplies each produce ± 15 V rails, which are connected into the backplane of the

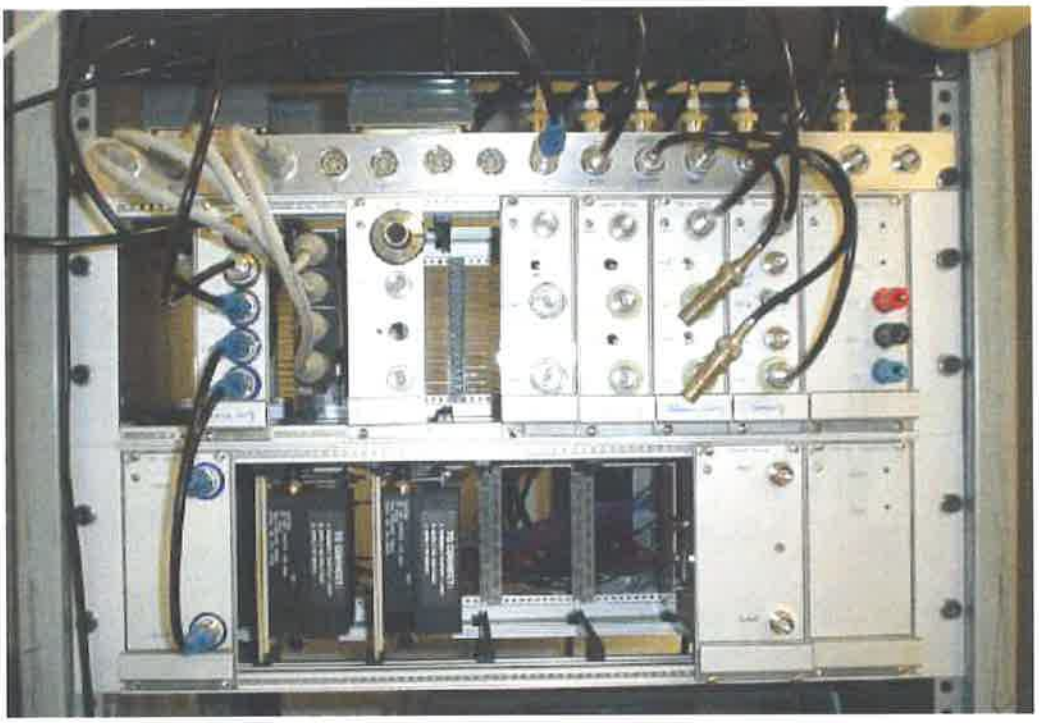

Figure E.1: The electonic components of the experiment. The top subrack houses the RF processing board, the input board, various compensators and a monitoring board. The bottom subrack holds the various mains-powered high power boards, including RF and high current amplifiers. The strip of connectors immediately above the top subrack provide easy connection to the experiment.

corresponding subrack. The schematic of the power supplies are shown in figure E.2. The power supplies are each capable of supplying two amperes on each output rail.

To ensure local power supply stability and to further reduce ripple, each board that used the $\pm 15\ V$ rails included a local power circuit to further regulate the power rails down to $\pm 12\ V$. These power supplies were simple linear power supplies using 78L12 and 79L12 regulators and appropriate filter capacitors to produce the local power rails.

E.4 Photodetector

The photodetectors used in the experiment were based on the "high power, large dynamic range detector" design given in Gray et al. [108], although the design has been modified somewhat. A schematic of the detector as shown in figure E.3, consists

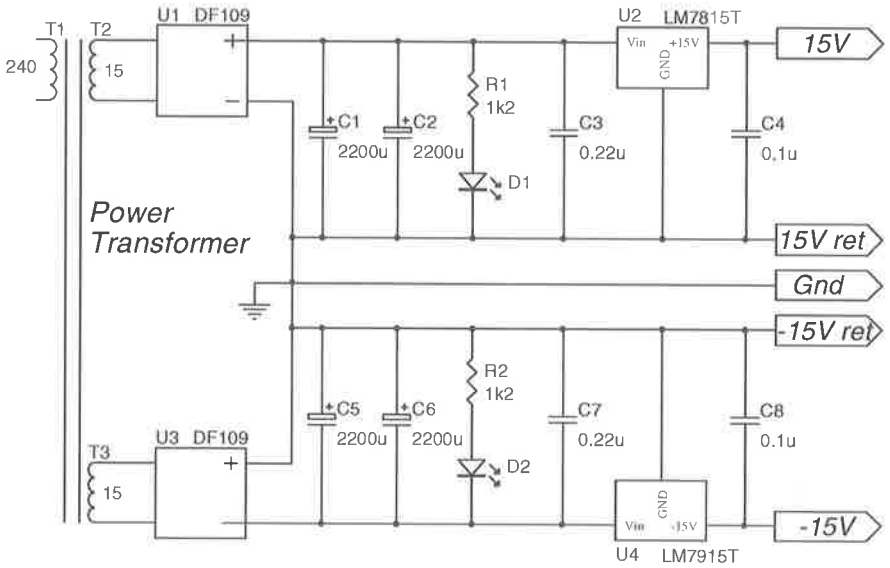

Figure E.2: Schematic of the sub-rack power supplies.

of two parts; the upper part is a simple transimpedance amplifier that is used to process the high frequency component of the photocurrent. The lower part amplifies the low frequency component of the photocurrent.

Photocurrent is produced when laser light is incident on the C30641 high power photodiode. This photodiode is an InGaAs PIN device, with a responsivity of approximately 0.8 A/W at 1064 nm [123]. The C30641 is an ideal photodiode for high power applications, having a maximum allowable photocurrent of 100 mA and a maximum dissipation of 100 mW. The C30641 is very similar to the ETX1000 high power photodiode used in the original circuit [124]. The photodiode is biased with a five volt rail that is produced by an on-board 78L05 voltage regulator and filtered by a 100 nF, surface mount, monolithic ceramic capacitor placed directly next to the photodiode. The bias voltage across the photodiode is a function of detected photocurrent, though the resultant change in photodiode capacitance and dark current are small for the relatively small bias change [123]. No problems with this varying bias are reported in [108], nor were any encountered in this work.

The photocurrent is split into its ac and dc components by a passive LC circuit. This filter has a corner frequency at 80 kHz and was realised with the inductor Lf and the capacitor Cf. High frequency components of the photocurrent flow

through the capacitor and then through the 1 $k\Omega$ feedback resistor (Rt) of the transimpedance amplifier to produce an output voltage. The use of the CLC409 current-feedback operational amplifier allows a bandwidth of about 60 MHz for the ac output [108]. This voltage was available on a BNC connector for connection to the PDH board via RG58 coaxial cable. The ac component was also connected to the input board (see below) *via* standard RG45 twisted pair cable.

Components of the photocurrent with frequency below 80 kHz flow through Lf and generate a voltage across the 15 Ω load resistor, $R1$. This voltage is then amplified by the amplifier made from the OP27 operational amplifier and its associated passive components. Two NE5532-based amplifiers are then used to convert the signal to a differential signal for transmission to the electronics rack. The differential transmission method was employed to reduce interference along the 8 m cable run. The section on the Input Board E.8 below describes how the differential signal was converted back to a single-ended signal for use in the various monitoring and servo systems.

E.4.1 High Power Photodetector

The photodetector used in the FSL system was required to measure high optical power so that the shot noise level could be lowered (see section 4.3.2). However, the detector circuit described above can not produce photocurrents above 21 mA without violating the 100 mW power limit of the photodiode. A variant of the standard photodetector was thus built that could handle up to 100 mW of continuous laser power. The resultant circuit is mostly identical to the circuit shown above, but the photodiode is bootstrapped with a voltage reference to fix the photodiode bias voltage to 1.25 V irrespective of the amount of light incident on the detector. This in turn limits the maximum dissipation of the photodiode to 100 mW over the expected operating range of the detector. Figure E.4 shows the modification made to the photodiode section of the standard photodetector circuit.

The physical construction of the photodetector is shown in figure E.5. The ac part of the circuit is made using surface mount components and is located in the right part of the figure. The photodiode is not visible in the picture as it is mounted

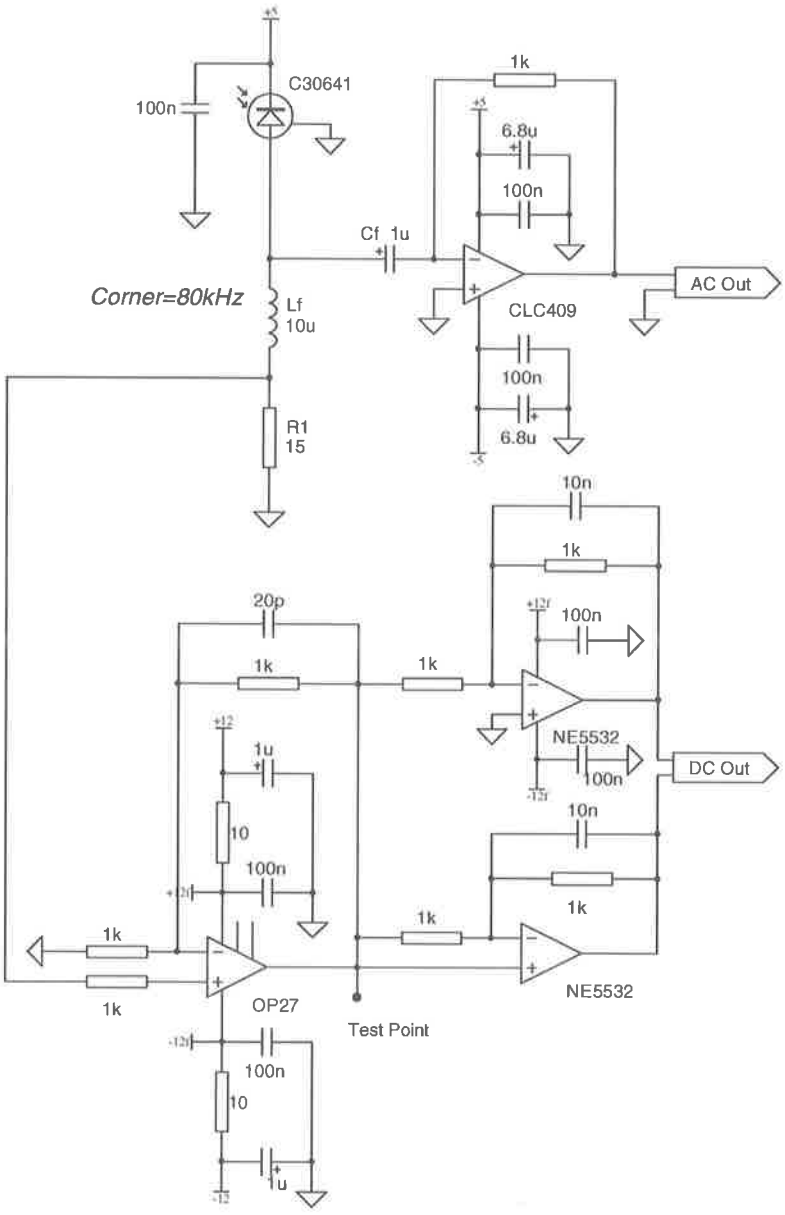

Figure E.3: Schematic diagram for the photodetector circuit.

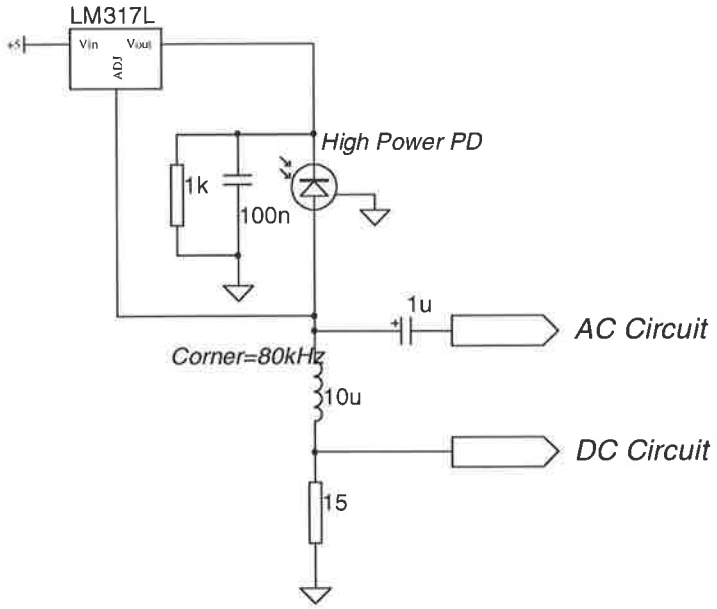

Figure E.4: Schematic of the photodiode section of the high power photodetector.

on the rear of the board, where it looks out through a hole drilled in the aluminium case. The dc part of the circuit is on the left of the diagram; the dc amplifier is at the top and the single-ended to differential output driver is at the bottom.

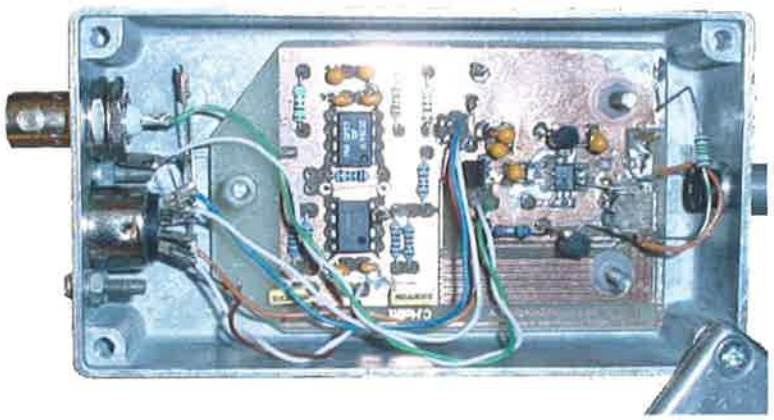

Figure E.5: The high power photodetector.

E.5 Input Board

The input board takes the differential signals produced by the detectors and converts them back to a single-ended signals. The board then connects the resulting signals onto the subrack backplane so that they can be accessed by the various servo boards. The input board also includes a linear power supply for the photodetectors. These power rails, along with ground and the two differential output signals, are passed to the photodetectors through an 8 core cable with 8-pin male DIN audio connectors at each end.

The input board consists of four independent channels for processing the outputs from four photodetectors. Each channel includes two amplifier stages, one each for processing the dc and ac components of the photodetector output. Figure E.6 shows the schematic for a single channel of the input board. As can be seen, the dc and ac signal processing in each channel are identical.

The circuit is an instrumentation amplifier made from three operational amplifiers and passive components. This arrangement resulted in a better noise performance than using commercially available instrumentation amplifiers, which are optimised for CMRR rather than noise performance. Adjustment of $R2$ and $R4$ allows tuning of the amplifier gain. In practice $R4$ was chosen to coarsely set the desired gain and then $R2$ was used for calibration of the photodetector. Optimisation of the amplifier common mode rejection was performed by applying a common-mode, ac signal to the two inputs of each channel and adjusting $R10$ until here was no detectable signal at the output of the amplifier.

Not shown in the schematic is the linear power regulation circuit which regulated the backplane power supplies down to $\pm 12\ V$. Each operational amplifier was decoupled with 100 nF decoupling capacitors and each channel had an additional 10 μF tantalum capacitor.

Figure E.7 shows the physical layout of the input board. The four 8-pin DIN connectors are located on the left of the figure. They carry the power and ground lines out to the photodetectors and bring back the differential ac and dc signals. The four identical modules of the board are clearly visible in the figure. The large piece of aluminium extrusion on the right of the figure acts as a heatsink for the

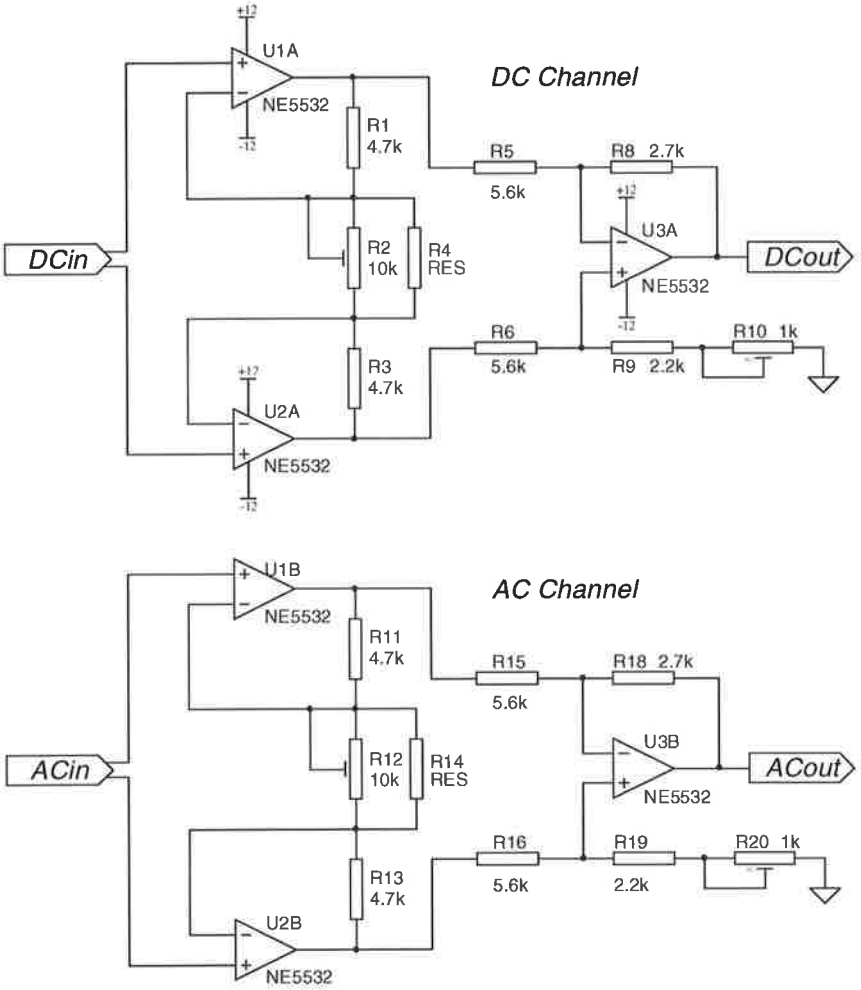

Figure E.6: Schematic of the input board.

power regulators that produce the photodiode supply rails.

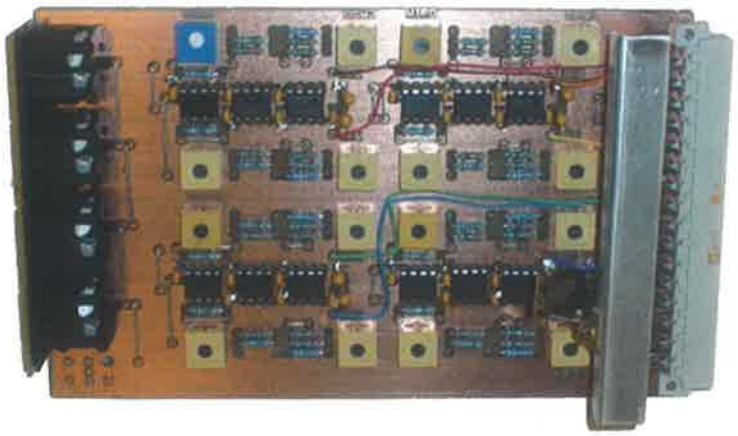

Figure E.7: The input board.

E.6 Pound-Drever-Hall Circuit

Figure E.8 shows the layout of the circuit used to implement the Pound-Drever-Hall sensor used to lock the laser to the reference cavity. The voltage controlled oscillator (VCO), $U2$, provides a stable 17 MHz signal. The VCO tuning voltage is supplied by a stable reference diode, $U1$. Adjustment of $R1$ allows the reference's output voltage and hence the VCO frequency to be varied. The signal is then amplified using a high power amplifier, capable of approximately 1 W of output power. At 17 MHz, the POS-25 VCO produces 9.5 dBm of RF power. This power is then boosted to approximately 26.6 dBm by the 16 dB gain of the ZHL-1A power amplifier. After amplification the signal is sampled using the directional coupler, $U4$, to provide the local oscillator signal (LO). The coupling of the directional coupler is 20 dB, so the resultant LO has an RF power of 6.6 dB. The mainline loss in the coupler is, at worst, 0.5 dB, so approximately 26 dB of RF power remains to drive the EOM.

In the signal path, the input signal first encounters $U5$, a low pass filter that removes the second harmonic of the PDH modulation signal. The attenuator then ensures that the input signal does not overload the level 7 mixer, $U7$, which combines the input signal with the local oscillator. The high frequency mixer products are

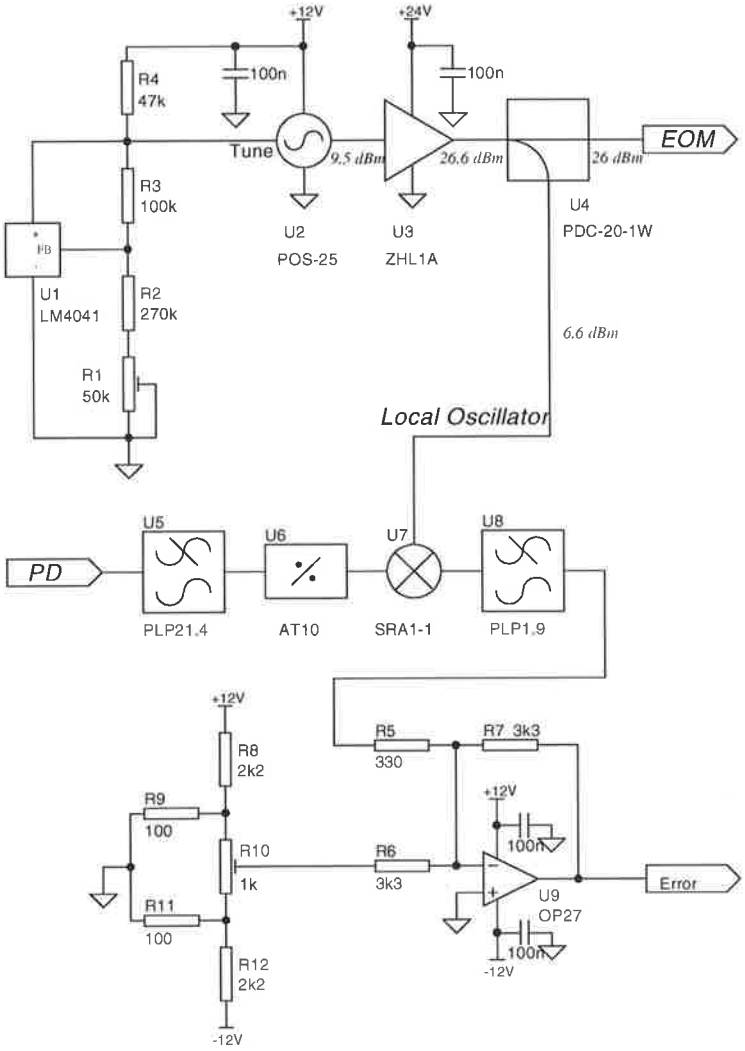

Figure E.8: Pound-Drever-Hall circuit.

removed using the low pass filter, $U8$, which has a corner frequency at 1.9 MHz. The operational amplifier $U9$ and its associated circuitry amplifies the error signal and buffers it before sending it to the output.

Adjustment of $R10$ allows any offset in the circuitry to be removed, ensuring that the output voltage is zero when the laser is at cavity line centre. This prevents any frequency offset between the laser frequency and the cavity resonance once the system is locked.

The RF processing card includes an on-board linear power supply. This supply takes the ± 15 V rails produced by the subrack power supply card and regulates

them down to $\pm 12\ V$. Local decoupling at each active device is supplied by 100 nF capacitors, which were placed as close to the devices as possible and were connected directly to the ground plane.

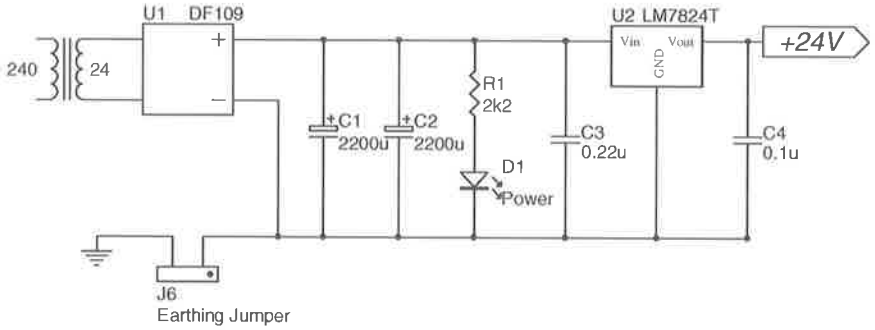

Figure E.9: Schematic of the RF amplifier's power supply.

Figure E.9 shows the amplifier's power supply. This power supply had a direct mains electrical connection and was therefore completely independent of all other circuitry in the experiment. This was important due to the very high level of RF power in this component; if connected to a common power rail it would have caused significant RF interference.

The RF circuitry was built on two standard Eurocards. One card was devoted to the high power amplifier and its dedicated power supply, the second card housed the remainder of the circuit. The RF section of the circuit used through-hole components manufactured by Minicircuits. These components were connected using 50 Ω stripline techniques on the PCB and appropriate 50 Ω cable for off-board connections. A well grounded copper shield was placed over the RF circuitry to prevent contamination of other electronics in the sub-rack with 17 MHz interference.

E.7 Compensator

The circuit used to implement the various compensators used in the experiment is shown in figure E.10. Different compensator designs were configured by appropriate selection of the different components. Table E.1 shows the components used to implement the fast and slow laser frequency servos, the servo to lock the reference

cavity to the test cavity and the servo to stabilise the laser intensity.

$U1$, $U2A$ and the associated passive components form a differential to single-ended low noise preamplifier. In this circuit $U1$ buffer the inputs if required and resistors $R1 - R7$ define the dc-gain of the amplifier. Adjustment of $R3$ allows easy tuning of the overall compensator gain. The capacitor, $C1$ combined with $R6$ produce a high frequency pole in the preamplifier transfer function. This pole was generally located beyond the final servo unity gain point in order to provide attenuation of high frequency sensor noise. Maintenance of a good common mode rejection ratio (CMRR) in the circuit required good matching of $R4$ and $R5$ and of $R6$ and $R7$. In addition the time constants $R6 \times C1$ and $R7 \times C2$ needed to be matched to maintain the CMRR at higher frequencies. High precision components were therefore used in this part of the circuit. The section of the circuit implemented with $R9 - R13$ provides a small DC voltage that can be tuned to cancel the circuit's input offset voltage. This voltage was tuned by grounding the input of the amplifier and then adjusting the offset voltage until the preamplifier output was zero.

The amplifier sections implemented with $U2B$ and $U3$ were used to shape the compensator transfer function. As described in the body of the thesis, amplifier $U2B$ was generally used to implement the high frequency part of the transfer function. $U3$ was then used to provide additional gain at low frequencies. When required this additional gain could be switched in and out of the circuit by closing switch $S1$. The desired transfer function was formed by an appropriate choice of the passive components surrounding $U2B$ and $U3$. The ratio $R15/R14$ sets the dc gain of the high frequency amplifier, while the addition of $C3$ forms a pole in the transfer function at $1/(R15 \times C3)$. Similarly, the ratio $\{R18//R22\}/R17$ sets the dc gain of the low frequency boost amplifier and there is a pole at $1/(\{R18//R22\} \times C6)$. This circuit also includes a zero at $1/(R17 \times C5)$. $R19$ and $R20$ form a voltage divider to determine how much of the low frequency boost is passed through to the compensator output. This feature was exploited in the intensity stabilisation compensator where $R20$ was made variable so that the gain could be smoothly adjusted. Finally, $R16$ and $C4$ provide a path for a calibration signal to be injected into the compensator output. Setting $R16 = R15$ and $C4 = C3$ ensures that the

calibration signal is impressed directly onto the output without being shaped by the output amplifier response.

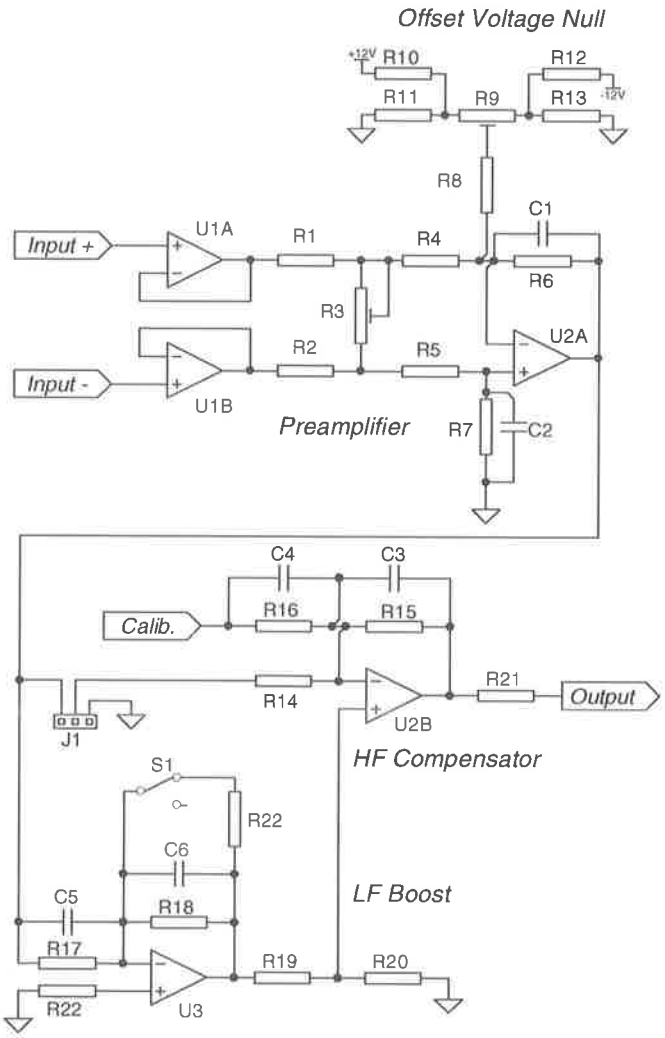

Figure E.10: Schematic diagram of the basic compensator circuit board.

The compensator circuits were built on a standard 100×160 mm Eurocard using a double layer process. Two complete compensator circuits were included on each card. Input signals could be accessed through either front panel connectors or through the sub-rack backplane as necessary. In addition to the amplifier circuits, the card also included a simple linear power supply to provide local regulation of the ± 15 V power rails down to ± 12 V. All active components were bypassed at their power terminals with 100 nF monolithic ceramic capacitors and a pair of 10

Component	Laser Freq (fast) (See fig 4.24)	Laser Freq (slow) (See fig 4.32)	PSC Freq (See fig 4.58)	Intensity (See fig 4.68)
U1	AD798	-	-	NE5532
U2	AD708	LF353	AD708	LF353
U3	AD708	LF353	AD708	LF353
R1,R2	1l	10k	1k	1k
R3	1k	10k	1k	10k
R4,R5	1k	1k	1k	1k
R6,R7	1k	1k	1k	150k
C1,C2	220p	-	-	22n
J1	SC	OC	SC	OC
R14	10k	-	220	-
R15	475	1k	3.3k	18k
C3	22n	-	2.2n	-
R16	475	-	3.3k	909+909
C4	22n	-	2.2n	-
R19	10k	10k	100k	100
R20	10k//100n	0	2k pot	1k pot
R21	100	0	0	47
R17	4.7k	470k	3.3k	2.2k
C5	-	-	-	220n
R18	475k	-	-	22k
R22	100	4.7M	3.3k	-
C6	1n	470n	470n	220n
S1	Included	Included	Included	OC
R22	OC	4.7M	0	SC

Table E.1: Component types and values used in the various compensators

μF tantalum capacitors were dedicated to each compensator section. One of the compensator cards can be seen in figure E.11

E.8 Heater Driver

The heater driver can provide up to an ampere to a heater that has a maximum (hot) resistance of 24 ohms. With no input voltage applied, the driver runs at approximately half power and applies 10 volts to the load. This feature is used to drive the steady state temperature of the load to above ambient temperature. The load can then be cooled by lowering the applied voltage, thus producing a heater that can be employed in linear control loops.

At the heart of the driver is a NE5534 operational amplifier, $U1$, with a Darling-

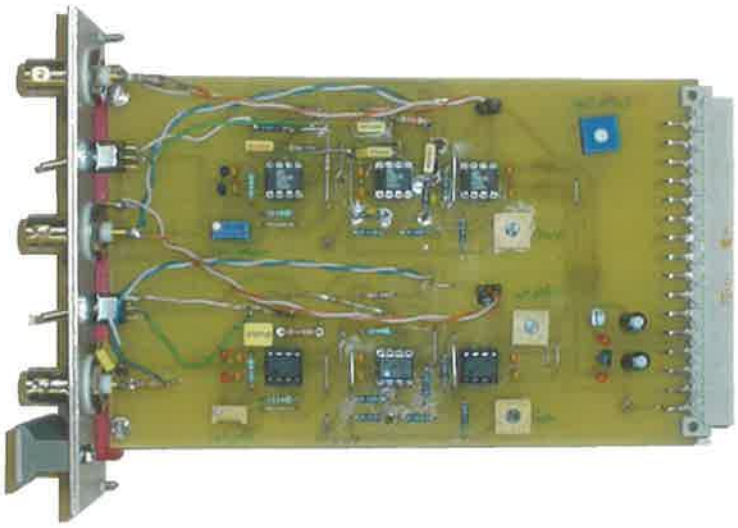

Figure E.11: The compensator board.

ton power transistor, $Q3$, placed inside its feedback loop. The feedback loop acts to remove any non-linearities produced by the transistor, so the $U1/Q3$ pair acts like a single operational amplifier with the input characteristics of the NE5534 and the high output drive of the Darlington transistor. The voltage appearing at the output is indicated at all times by a tri-colour indicator LED. The LED appears red when the heater is running at full power, yellow at half power and green as the output voltage drops below half-way. The colour change is a smooth function of output voltage, so the applied heat can be easily assessed at any time.

The driver is also equipped with a preheat circuit. This circuit applies full power to the heater at switch on so that the load is driven quickly up from ambient to its steady state temperature. The combination of Ctc and Rtc form a large time constant, leading to a very slowly falling voltage at the input to the buffer amplifier $U2$. During this time the slowly falling voltage holds $Q2$ off, producing in 24 volts at the output of the preheat circuit. Eventually the voltage falls sufficiently far that the two base voltages of the transistors, $Q1$ and $Q2$, are equal. Beyond this point $Q2$ turns on and the driver enters its normal operating mode. Adjustment of $R4$ allows the behaviour of this preheat circuit to be tailored to the particular application.

A photograph of the heater can be seen in figure E.13. The toroidal power

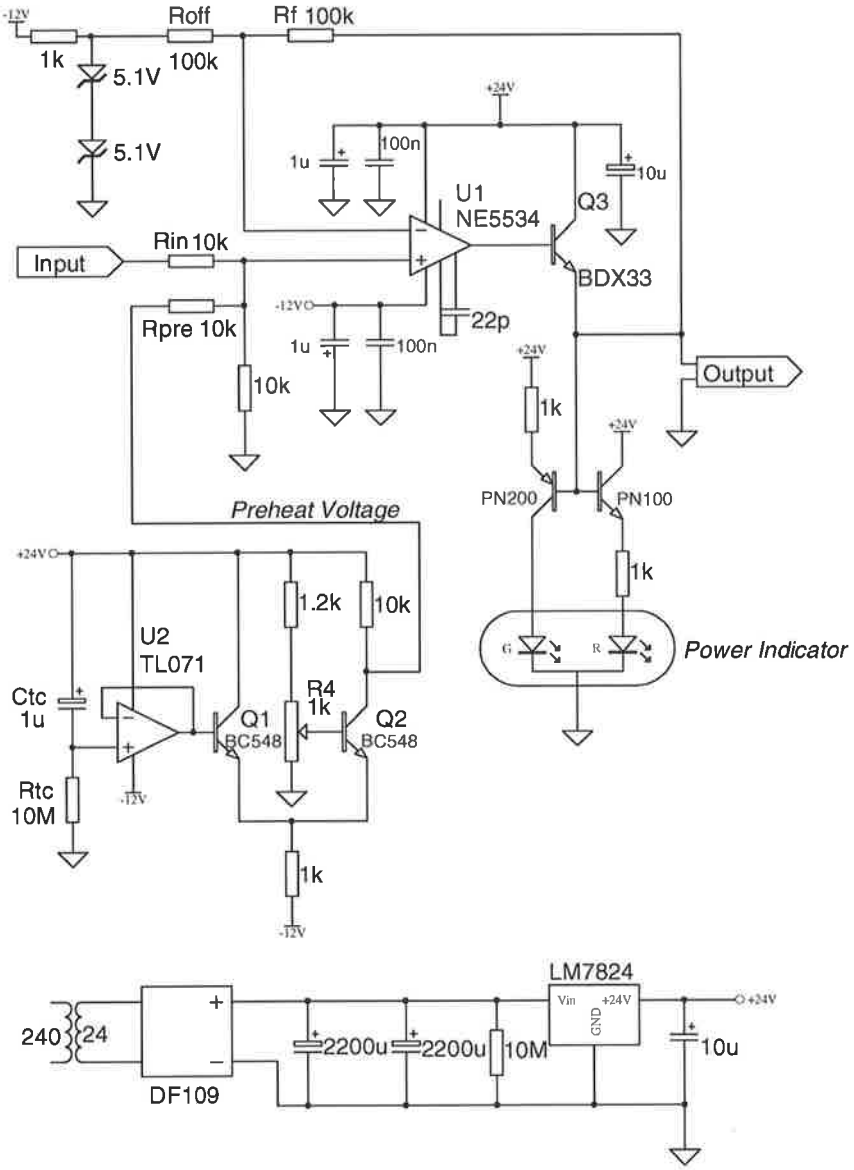

Figure E.12: Schematic diagram of the heater driver.

transformer dominates the circuit, which in this case was made using a single layer PCB. Unetched copper on the back side of the PCB was used as a heatsink for the power transistor and the voltage regulator.

Figure E.13: The heater driver.

E.9 Low-Noise High-Voltage Amplifier

The low noise, high voltage amplifier was based around the PA85 hybrid operational amplifier made by APEX semiconductors. The amplifier gain of 100 is set by the ratio of the feedback resistors, Rf, to the input resistors, Rin. The full 400 volts of the amplifier output could be placed across the feedback resistor: as this voltage is considerably in excess of a normal $1/4\ W$ resistor's voltage rating, the feedback resistor was built from five $1\ W$ resistors placed in series. An offset circuit made with $R1 - R3$ and buffered by $U2$ is used to tune the dc voltage at the amplifier output.

The amplifier was observed to be significantly noisier for output voltages less than about $100\ V$ or greater than approximately $300\ V$. For this reason the amplifier was always operated within this range. It is unknown whether this behaviour is typical of the PA85 or is the result of some damage to the chip.

The amplifier's switch on inrush current is limited by a 10 $k\Omega$ resistor that is placed into the supply line. This limiting prevents tripping the foldback current limiting present on many high voltage supplies. Once the amplifier's high voltage decoupling capacitor has been charged the inrush limiting resistor is switched out of the circuit.

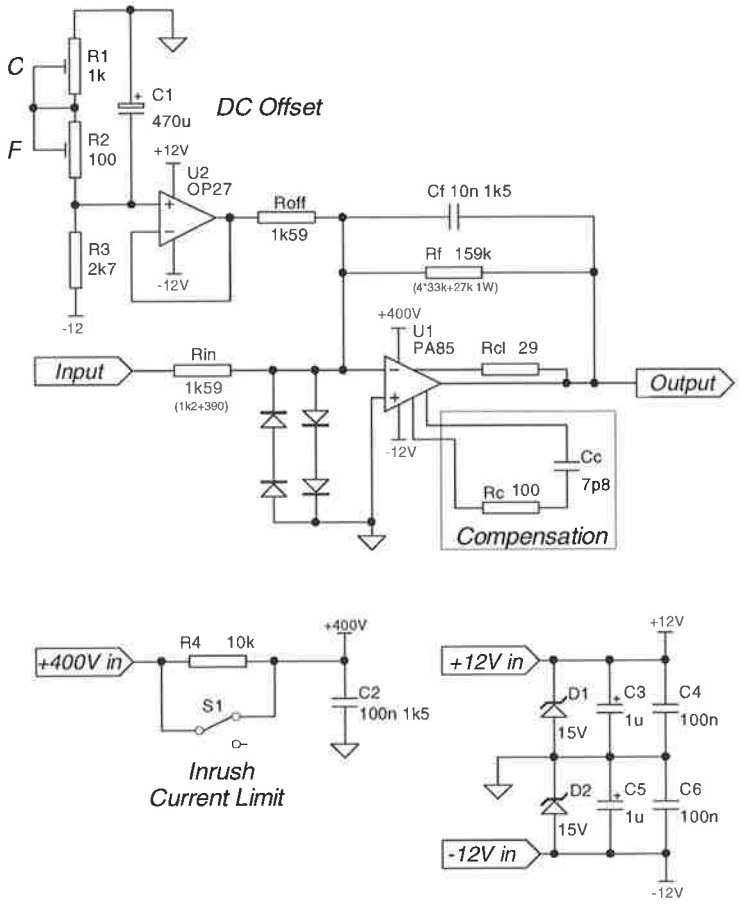

Figure E.14: Schematic of the low noise, high voltage amplifier.

E.10 Extended Range High-Voltage Amplifier

The amplifier shown in figure E.15 is a high voltage amplifier with a output dynamic range of 1000 volts. The amplifier uses an APEX hybrid PA89 operational amplifier that is configured to have a gain of 100. This gain is set by the ratio of feedback resistors ($R12-R19$) to the input resistors ($R6$, $R7$ and $R8$, $R9$). The offset circuitry

produced by $R1 - R3$ and $U1$ allows the dc level at the output to be changed. A low voltage monitoring point is produced by $R21 - R25$ and is then buffered by $U3$. This part of the circuit reduces the output voltage by a factor of 100, allowing the amplifier's output voltage to be monitored with conventional low voltage equipment.

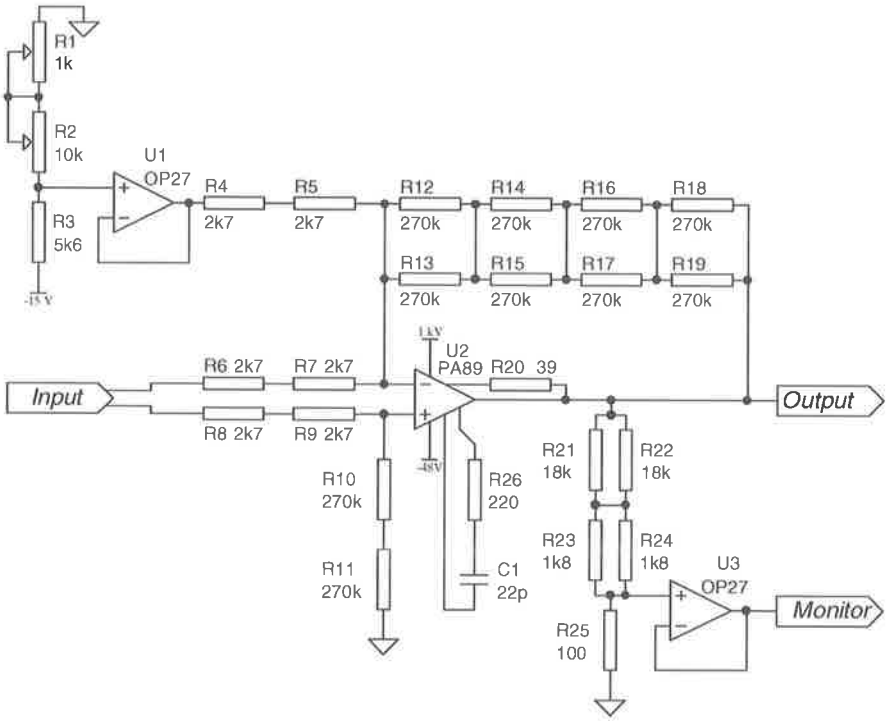

Figure E.15: Schematic of the high voltage amplifier having 1000 V dynamic range.

Appendix F

Autolocking Pseudocode

```
//================================================================
   // Pseudocode
   // Autolocking of the prestabilistion cavity to the test cavity
   //================================================================

   //Global Constants-----------------------------------------------
   P_SEEK= 0.5 //Heater power in seek state
   P_TM= 0.5 //Heater power in transverse mode dump
   P_LM= 5 //Heater power in longitudional mode dump
   T_VERIFY= 40 //Delay to see whether mode has locked (ms)
   T_TM_Dump= 6 //Length of transverse mode dump
   T_TM_Wait= 30 //Length of wait after TM dump
   T_LM_Dump= 7 //Length of longitudinal mode dump
   T_LM_Wait= 40 //Length of wait after LM dump
   OutThreshold=2.5 //Threshold to specify that cavities are locked

   //Initialisation code--------------------------------------------
   init_chip //Initialise the chip
   init_timer //set up the timer
   init_comparator //set up the voltage comparator
   init_adc //set up the analogue to digital converter
```

```
//Finite State Machine Code--------------------------------------

mode= &Seek
 //Mode holds the address of the current substate
 // start the FSM in the Seek substate

indicatorUser=OFF
 //An indicator that the user needs to check which
 // mode has been locked.
 // The indicator starts in the off state

//-----------------------------------------------------------------
//FSM_Main implements the Laser_Locked superstate.
// If the laser is locked the code jumps to the
// appropriate substate.  If not then the code waits
// for laser lock to be achieved
FSM_Main:
 if laserLocked
 goto mode
 else
 mode= &Seek //Reinitialise the FSM substate
 goto FSM_Main
//-----------------------------------------------------------------
//Laser Locked Substates

Seek:
 outPower=P_SEEK //Set output power to appropriate level
 indicatorUser=OFF
 if TCresonant
 then mode= &Unverified_Entry
```

```
 goto FSM_Main

Unverified_Entry:
 outPower=0
 Timer(T_VERIFY)
 if ADCInput > threshold for test cavity power
 mode= &Verified
 else
 mode= &TM_Dump
 goto FSM_Main
Unverified:
 indicatorUser=ON
 if ~TCresonant
 then mode= &Seek
 if mode accepted by operator
 then mode= &Verified
 if mode rejected by operator
 then mode= &TM_Dump
 goto FSM_Main

Verified:
 indicatorUser=OFF
 outPower=0
 if ~TCresonant
 then mode= &LM_Dump
 goto FSM_Main

TM_Dump:
 indicatorUser=OFF
 outPower= P_TM
 Timer(T_TM_DUMP)
```

```
 mode= &TM_Wait
 goto FSM_Main

TM_Wait:
 Timer(T_TM_WAIT)
 mode= &TM_Wait
 outPower= 0
 goto FSM_Main

LM_Dump:
 outPower= P_LM
 Timer(T_LM_DUMP)
 mode= &LM_Wait
 goto FSM_Main

LM_Wait:
 outPower=0
 Timer(T_LM_WAIT)
 mode= &LM_Wait
 goto FSM_Main

//Timer management-----------------------------------------------

subroutine Timer(Sleep_time)
//This subroutine puts the uC to sleep for a specified time
//The timer interrupt handler wakes the uC up when the time
// has expired.
 timer_variable=Sleep_time
 sleep
end
```

```
//Timer Interrupt Handler
interrupt_Timer:
 return

//Laser locked detector-----------------------------------------
interrupt_CompHI:
// This interrupt routine is triggered whenever the photodiode
// exceed its threshold.  If this situation presists for a second
// then the laser is assumed to have gained lock.
 Timer(1)
 if comparator output is high
 laserLocked=1;
interrupt_CompHI:
// This interrupt routine is triggered whenever the photodiode
// output drops below its threshold.
// If this situation presists for a second then the laser is
// assumed to have lost lock.
 Timer(1)
 if comparator output is low
 laserLocked=0;

//END=============================================================
```

Appendix G

Line Removal

Many of the graphs of experimental spectra presented in this thesis show strong, narrowband signals. Most frequently these are due to the 50 Hz mains frequency and its harmonics. Because these lines can make interpretation of the data difficult they have been removed from some graphs (as indicated in the text). The algorithm works as follows:

1. Find a local maximum, Max;
2. Find the local minima immediately to the left and right of Max;
3. Calculate the mean of the local minima;
4. If the maximum is greater than a specified multiple of the mean of the local minima then the local maximum is assumed to be due to a narrowband source. Otherwise, the local maximum is assumed to be natural variation in the spectrum.
5. If the local maximum is deemed to be due to a narrowband source then the frequency region between the local minima is replaced by interpolation between the values of the local minima.

The MATLAB function used to implement the line removal algorithm is shown below. The threshold value was set to two for all graphs presented in this thesis.

```
function [Out]= removeLines(In,Threshold)
```

```
%remove lines from spectra
% Threshold sets the level below which peaks are not removed
% (ie if the peak is not a factor of threshold above the
%    background then it is left alone).
if nargin<2
 Threshold = 1
end
InLength=length(In);
NaNPointer=isnan(In);
Out=In;
for ii = 2:InLength-1
    if (In(ii)>In(ii-1))&(In(ii)>In(ii+1)) %find peak
        LeftPointer=ii;
        RightPointer=ii;
        %find left local minimum
        while ( In(LeftPointer)>In(LeftPointer-1) )
              | ( In(LeftPointer)>In(LeftPointer+1) )
                if LeftPointer < 3, break, end;
                LeftPointer=LeftPointer-1;
         end

        %find right local minimum
        while ( In(RightPointer)>In(RightPointer-1) )
          | ( In(RightPointer)>In(RightPointer+1) )
            if RightPointer > (InLength-2), break, end;
                RightPointer=RightPointer+1;
         end
        if In(ii) > Threshold*mean([In(LeftPointer),In(RightPointer)])
            %interpolate across peak
            FreqSpan=RightPointer-LeftPointer;
            Increment=(In(RightPointer)-In(LeftPointer))/FreqSpan;
            jj=LeftPointer;
```

```
            for jj=1:(RightPointer-LeftPointer)
                Out(LeftPointer+jj)=In(LeftPointer)+jj*Increment;
            end
        end
    end
end
Out(1)=In(1);
Out(InLength)=In(end);
Out(NaNPointer)=nan;
if size(In)~=size(Out)
    Out=Out';
end
```

Bibliography

[1] N. Dadhich and J. Narlikar, eds., *Astrophysical Sources of Gravitational Radiation and Prospects for Their Detection*, Inter-University Centre for Astronomy and Astrophysics, 1998.

[2] C. Cutler and K.S.Thorne, "An overview of gravitational-wave sources," *ArXiv General Relativity and Quantum Cosmology e-prints*, April 2002.

[3] C.W.Misner, K. Thorne, and J. Wheeler, *Gravitation.* San Francisco: W. H. Freeman, 1973.

[4] G. M. Harry, H. Armandula, E. Black, D. R. Crooks, G. Cagnoli, M. M. Fejer, J. Hough, S. D. Penn, S. Rowan, R. Route, and P. Sneddon, "Optical coatings for gravitational wave detection," *Proceedings of SPIE.*, vol. Vol. SPIE-5527, pp. 33–40, 2004.

[5] J. Johnson, "Thermal agitation of electricity in conduction," *Physical Review*, vol. 32, pp. 97–109, 1928.

[6] H. Nyquist, "Thermal agitation of electric charge in conductors," *Physical Review*, vol. 32, pp. 110–113, 1928.

[7] Y. Levin, "Internal thermal noise in the LIGO test masses: A direct approach," *Physical Review D*, vol. 57, pp. 659–663, 1998.

[8] P. Michelson, J. Price, and R. Taber, "Resonant-mass detectors of gravitational radiation," *Science*, vol. 237, p. 150, 1987.

[9] P. R. Saulson, "Thermal noise in mechanical experiments," *Physical Review D*, vol. 42, pp. 2437–2445, 1990.

[10] A. Heidmann, P. Cohadon, and M. Pinard, "Thermal noise of a plano-convex mirror," *Physics Letter A*, vol. 263, pp. 27–32, 1999.

[11] V. Kinra and K. Milligan, "A second-law analysis of thermoelastic damping," *Jouranl of Applied Mechanics*, vol. 61, pp. 71–76, 1994.

[12] V. Braginsky, M. Gorodetsky, and S. Vyatchanin, "Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae," *Physics Letters A*, vol. 264, pp. 1–10, 1999.

[13] E. Gustafson, D.Shoemaker, K. Strain, and R. Weiss, "LSC whitepaper on detector research and development," Tech. Rep. LIGO-T990080-00-D, LSC, 1999.

[14] K. Numata, M. Ando, K. Tamamoto, S. Otsuka, and K. Tsubono, "Wide-band direct measurment of thermal fluctuations in an interferometer," *Physical Review Letters*, vol. 91, pp. 260602–1 – 260602–4, 2003.

[15] M. Kajima, N. Kusumi, S. Moriwaki, and N. Mio, "Wide-band measurement of mechnical thermal noise using a laser interferometer," *Physics Letters A*, vol. 263, pp. 21–26, 1999.

[16] E. D. Black, A. Villar, K. Barbary, A. Bushmaker, J. Heefner, S. Kawamura, F. Kawazoe, L. Matone, S. Meidt, S. R. Rao, K. Schulz, M. Zhang, and K. G. Libbrecht, "Direct observation of broadband coating thermal noise in a suspended interferometer," *Physics Letters A*, vol. 328, pp. 1–5, 2004.

[17] E. D. Black, A. Villar, and K. G. Libbrecht, "Thermoelastic-damping noise from sapphire mirrors in a fundamental-noise-limited interferometer," *Physics Review Letters*, vol. 93, pp. 241101–1 – 241101–4, 2004.

[18] 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*, vol. 45, pp. 1569–1574, 2006.

[19] G. M. Harry, M. R. Abernathy, A. E. B.-T. A. Armandula, E. Black, K. Dooley, M. E. A. Nwabugwu, A. Villar, D. R. M. C. A. Cagnoli, J. Hough, C. R. How, I. M. A. Murray, S. Reid, S. Rowan, P. H. S. A. M. Fejer, R. Route, S. D. Penn, P. G. A.-M. Mackowski, C. Michel, L. Pinard, and A. Remillieux, "Titania-doped tantala/silica coatings for gravitational-wave detection," *Classical and Quantum Gravity*, vol. 24, pp. 405–415, 2007.

[20] E. Black, I. Grudinin, A. Villar, and K. G. Libbrecht, "Thermal noise interferometer update and status," Tech. Rep. LIGO-G040156-00-R, LIGO, 2004.

[21] E. Black, "A possible thermal noise measurement using the 40-meter," Tech. Rep. LIGO-T000051-00-R, LIGO, 2000.

[22] A. Yariv, *Optical Electronics.* Saunders College Publishing, 1991.

[23] A. D. Virgilio, A. Giassa, and P. LaPenna Tech. Rep. VIR-NOT-PIS-1390-061, VIRGO, 1996.

[24] V. Leonhardt, L. Ribichini, P. Klöverkorn, B.Wilke, H. Lück, and K. Danzmann, "Towards measuring the off-resonant thermal noise of a pendulum mirror," *Classical and Quantum Gravity*, vol. 19, pp. 1717–1721, 2002.

[25] V. Leonhardt, L. Ribichini, H. Luck, and K. Danzmann, "The hannover thermal noise experiment," *Classical and Quantum Gravity*, vol. 21, pp. S1127–S1131, 2004.

[26] B. Slagmolen, *Direct Measurement of the Spectral Distribution of Thermal Noise.* PhD thesis, Australian National University, 2005.

[27] C. M. Mow-Lowry, S. Goßler, B. J. J. Slagmolen, J. Cumpston, M. B. Gray, and D. E. McClelland, "Towards the SQL: Status of the direct thermal-noise measurements at the ANU," *Journal of Physics: Conference Series*, vol. 32, pp. 362–367, 2006.

[28] A. Di Virgilio, S. Bigotta, L. Barsotti, S. Braccini, C. Bradaschia, G. Cella, V. Dattilo, M. Del Prete, I. Ferrante, F. Fidecaro, I. Fiori, F. Frasconi, A. Gennai, A. Giazotto, P. La Penna, G. Losurdo, E. Majorana, M. Mantovani,

F. Paoletti, R. Passaquieti, D. Passuello, F. Piergiovanni, A. Porzio, P. Puppo, F. Raffaelli, P. Rapagnani, F. Ricci, S. Solimeno, G. Vajente, and F. Vetrano, "Displacement power spectrum measurement of a macroscopic optomechanical system at thermal equilibrium," *ArXiv General Relativity and Quantum Cosmology e-prints*, Dec. 2006.

[29] H. B. Callen and T. A. Welton, "Irreversibility and generalised noise," *Physical Review*, vol. 83, pp. 34–40, 1951.

[30] L. Landau and E. Lifschitz, *Statistical Physics.* Butterworth-Heinemann, 3rd ed., 1984.

[31] Y. T. Liu and K. S. Thorne, "Thermoelastic nose and homogeneous thermal noise in finite sized gravitational-wave test masses," *Physical Review D*, vol. 62, pp. 122002–1 – 122002–10, 2000.

[32] M. Cerdonio, L. Conti, A. Heidmann, and M. Pinard, "Thermoelastic effects at low temperatures and quantum limits in displacement measurements," *Physical Review D*, vol. 63, pp. 082003–1 – 082003–9, 2001.

[33] T. Uchiyama, K. Kuroda, M. Ohashi, S. Miyoki, H. I. A. Yamamoto, H. Hayakawa, K. Kasahara, M.-K. F. A. Kawamura, R. Takahashi, T. Yamazaki, K. Arai, D. T. A. Ueda, M. Fukushima, S. Sato, Y. Tsunesada, Z.-H. Z. A. Shintomi, A. Yamamoto, T. Suzuki, Y. Saito, T. Haruyama, N. S. A. Higashi, T. Tomaru, K. Tsubono, M. Ando, K. Numata, Y. A. A.-I. Ueda, H. Yoneda, K. Nakagawa, M. Musha, N. Mio, S. M. A. Somiya, A. Araya, A. Takamori, N. Kanda, S. Telada, H. T. A. Nakamura, M. Sasaki, T. Tanaka, H. T. A. N. K-I Ohara and, O. Miyakawa, and M. E. Tobar, "Present status of large-scale cryogenic gravitational wave telescope," *Classical and Quantum Gravity*, vol. 21, pp. S1161–S1172, 2004.

[34] S. D. Penn, P. H. Sneddon, H. Armandula, J. C. Betzwieser, G. Cagnoli, J. Camp, D. R. M. Crooks, M. M. Fejer, A. M. Gretarsson, G. M. Harry, J. Hough, S. E. Kittelberger, M. J. Mortonson, R. Route, S. Rowan, and

C. C. Vassiliou, "Mechanical loss in tantala/silica dielectric mirrorcoatings," *Classical and Quantum Gravity*, vol. 20, pp. 2917–2928, 2003.

[35] D. R. M. Crooks, P. Sneddon, G. Cagnoli, J. Hough, S. Rowan, M. M. Fejer, E. Gustafson, R. Route, N. Nakagawa, D. Coyne, G. M. Harry, , and A. M. Gretarsson, "Excess mechanical loss ossociated with dielectric mirror coatings on test masses in interferometric gravitational wave detectors," *Classical and Quantum Gravity*, vol. 19, pp. 883–896, 2002.

[36] 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, JimHough, and N. Nakagawa, "Thermal noise in interferometric gravitational wavedetectors due to dielectric optical coatings," *Classical and Quantum Gravity*, vol. 19, pp. 897–917, 2002.

[37] W. H. Glenn, "Noise in interferometic optical systems: An optical nyquist theorem," *IEEE Journal of Quantum Electronics*, vol. 25, pp. 1218–1224, 1989.

[38] K. H. Wanser, "Fundamental phase noise limit in optical fibres due to temperature fluctuations," *Electronics Letters*, vol. 28, pp. 53,54, 1992.

[39] V. Braginsky, M. Gorodetsky, and S. Vyatchanin, "Thermo-refractive noise in gravitational wave antennae," *Physics Letters A*, vol. 271, pp. 303–307, 2000.

[40] E. Gustafson, B. Kells, G. Cagnoli, J. Hough, N. Nakagawa, and M. Fejer, "Thermooptic phase noise induced in light beams trasmitted through bulk optics.".

[41] E. Rønnekleiv, "Frequency and intensity noise of single fiber grating lasers," *Optical Fiber Technology*, vol. 7, pp. 206–235, 2001.

[42] K. R. Foster, "Quick and easy field analysis," *IEEE Spectrum*, vol. 36, pp. 87–88, 1999.

[43] Newport, *The Newport Resource*, 2004.

[44] A. Gillespie and F. Raab, "Thermally excited vibrations of the mirror of laser interferometer gravitational-wave detectors," *Phyiscal Review D*, vol. 52, pp. 577–585, 1995.

[45] D. Shoemaker, R.Shilling, L. Schnupp, W.Winkler, K. Maischberger, and A. Rudiger, "Noise behaviour of the garching 20-metre gravitational-wave detector," *Physical Review D*, vol. 38, pp. 423–432, 1988.

[46] K. Danzmann, J. Chen, P. Nelson, T. Niebauer, A. Rudiger, R. Schilling, L. Schnupp, K. Strain, H. Walther, W.Winkler, J. Hough, A. Campbell, C. Cantley, J. Logan, B. Meers, E. Morrison, G. Newton, D. Robertson, N. Robertson, S. Rowan, K. Skeldon, P. Veitch, H.Ward, H.Welling, P. Aufmuth, I. Kropke, D. Ristau, J. Hall, J. Bennet, I. Corbett, B. Edwards, R. Elsey, R. Greenhalgh, B. Schultz, D. Nicholson, J. Shuttleworth, J. Ehlers, P. Kafka, G. Schafer, H. Braun, and V. Kose, *Mechanical Aspects in Interferometric Gravitational Wave Antenna.*, p. 251. Springer, 1992.

[47] D. C. Giancoli, *Physics for Scientists and Engineers with Modern Physics.* Englewood CLiffs, New Jersey 07632: Prentice-Hall, 1988.

[48] H. Ellis, ed., *Book of Data.* Burnt Mill, Harlow, Essex CM20 2JE, England: Longman Group, 1990.

[49] J. H. Carver, G. N. Haddad, T. I. Hobbs, B. R. Lewis, and D. G. McCoy, "Vacuum ultraviolet 6-m monochromator," *Applied Optics*, vol. 17, p. 420, 1978.

[50] J. Winterflood, D. Blair, R. Schilling, and M. Notcutt, "Position control for suspended masses in laser interferometer gravitational wave detectors," *Review of Scientific Instruments*, vol. 66, pp. 2763–2776, 1995.

[51] M. Plissi, C. Torrie, M. Husman, N. Robertson, K. Strain, H. Ward, H. Lück, and J.Hough, "GEO 600 triple pendulum suspension system: Seismic isolation and control," *Review of Scientific Instruments*, vol. 71, pp. 2539–2545, 2000.

[52] G. Ballardin, L. Bracci, S. Braccini, C. Bradaschia, C. Casciano, G. C. A. Cavalieri, R. Cecchi, G. Cella, E. Cuoco, E. DAmbrosio, V. Dattilo, A. D. Virgilio, L. Fabbroni, F. Fidecaro, F. Frasconi, A. Gaddi, A. Gennai, G. Gennaro, A. Giazotto, G. Losurdo, L. Holloway, P. L. Penna, F. Lelli, E. M. M. Mazzoni, F. Paoletti, M. Pasotti, A. Pasqualetti, R. Passaquieti, D. Passuello, R. Poggiani, P. Puppo, F. Raffaelli, P. Rapagnani, F. Ricci, P. Ruggi, R. Stanga, R. Taddei, F. Vetrano, A. Vicere', and Z. Zhang, "Measurement of the VIRGO superattenuator performance for seismic noise suppression," *Review of Scientific Instruments*, vol. 72, pp. 3643–3652, 2001.

[53] A. Takamori, M. Ando, A. Bertolini, G. Cella, R. DeSalvo, M. Fukushima, Y. Iida, F. Jacquier, S. Kawamura, S. Márka, Y. Nishi, K. Numata, V. Sannibale, K. Somiya, R. Takahashi, H. Tariq, K. Tsubono, J. Ugas, N. Viboud, H. Yamamoto, T. Yoda, and C. Wang, "Mirror suspension system for the TAMA SAS," *Classical and Quatum Gravity*, vol. 19, pp. 1615–1621, 2002.

[54] T. V. C. P. by G. Losurdo), "The inertial damping of the VIRGO superattenuator and the residual motion of the mirror," *Classical and Quatum Gravity*, vol. 19, p. 16311637, 2002.

[55] T. V. C. P. by S Braccini), "The VIRGO suspensions," *Classical and Quatum Gravity*, vol. 19, pp. 1623–1629, 2002.

[56] N. A. Robertson, B. Abbott, R. Abbott, R. Adhikari, G. Allen, H. Armandula, S. Aston, A. Baglino, M. Barton, B. Bland, R. Bork, J. Bogenstahl, G. Cagnoli, C. Campbell, C. A. Cantley, K. Carter, D. Cook, D. Coyne, D. Crooks, E. Daw, D. DeBra, E. Elliffe, J. Faludi, P. Fritschel, A. Ganguli, J. Giaime, S. Gossler, A. Grant, J. Greenhalgh, M. Hammond, J. Hanson, C. Hardham, G. Harry, A. Heptonstall, J. Heefner, J. Hough, D. Hoyland, W. Hua, L. Jones, R. Jones, J. Kern, J. LaCour, B. Lantz, K. Lilienkamp, N. Lockerbie, H. Lück, M. MacInnis, K. Mailand, K. Mason, R. Mittleman, S. Nayfeh, J. Nichol, D. J. Ottaway, H. Overmier, M. Perreur-Lloyd, J. Phinney, M. Plissi, W. Rankin, D. Robertson, J. Romie, S. Rowan, R. Scheffler, D. H. Shoemaker, P. Sarin, P. Sneddon, C. Speake, O. Spjeld, and G. Stapfe,

"Seismic isolation and suspension systems for advanced LIGO," *Proceedings of SPIE,*, vol. 5500, pp. 81–91, 2004.

[57] N. A. Robertson, M. Barton, G. Cagnoli, C. A. Cantley, D. Coyne, D. Crooks, E. Elliffe, P. Fritschel, S. Gossler, A. Grant, A. Heptonstall, J. Hough, R. Jones, H. Lueck, R. Mittleman, M. Perreur-Lloyd, M. Plissi, D. Robertson, J. Romie, S. Rowan, D. Shoemaker, P. Sneddon, K. Strain, C. Torrie, H. Ward, and P. Willems, "Advanced LIGO suspension system conceptual design," Tech. Rep. T010103-04-D, 2005.

[58] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni, and M. Punturo, "Full scale prototype of high q pendulum for interferometric gravitational wave detectors," *Review of Scientific Instruments*, vol. 71, pp. 2206–2210, 2000.

[59] Y. Huang and P. R. Saulson, "Dissipation mechanisms in pendulums and their implication for gravitational wave interferometers," *Review of Scientific Instruments*, vol. 69, pp. 544–553, 1998.

[60] K. Tsubono, A. Arayo, K. Kawabe, S. Moriwaki, and N. Mio, "Triple-pendulum vibration isolation system for a laser interferometer," *Review of Scientific Instruments*, vol. 64, pp. 2237–2240, 1993.

[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 multi-stage pendulum suspensions for use in interferometric gravitational wave detectors," *Review of Scientific Instruments*, vol. 75, pp. 4516–4522, 2004.

[62] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni, and M. Punturo, "Eddy current damping of high q pendulums in gravitational wave detection experiments," *Review of Scientific Instruments*, vol. 69, pp. 2777–2780, 1998.

[63] L. E. Holloway, C. Bradaschia, E. Calloni, M. Cobal, R. D. Fabro, A. D. Virgilio, A. Giazotto, H. Kautzky, V. Montelatici, M. Morganti, D. Passuello, S. Braccini, R. Flamino, L. D. Fiore, and W. Velloso, "A coil system

for VIRGO providing a uniform magnetif field gradient," *Physics Letters A*, vol. 171, pp. 162–166, 1992.

[64] K. Somiy, H. Tariq, O. Miyakawa, G. Heinzel, N. Mio, and S. Kawamura, "Induced current damping for the suspension system of a gravitational-wave detector," *Review of Scientific Instruments*, vol. 73, pp. 3942–3945, 2002.

[65] S. M. Aston and D. M. Hoyland, "OSEM preliminary design document test report," Tech. Rep. LIGO-T050111-00-K, LIGO, 2005.

[66] J. Romie, "Sensor/Actuator head assembly," Tech. Rep. D960011-B, LIGO, 1996.

[67] M. Plissi, K.A.Strain, C.I.Torrie, N. Robertson, S.Killbourn, S.Rowan, S. Twyford, H. Ward, K.D.Skeldon, and J.Hough, "Aspects of the suspension system for GEO 600," *Review of Scientific Instruments*, vol. 69, pp. 3055–3061, 1998.

[68] D. J. Dawid and S. Kawamura, "Investigation of violin mode q for wires of various materials," *Review of Scientif Instruments*, vol. 68, pp. 4600–4603, 1997.

[69] I.A.Bilenko, V.B.Braginsky, and N. Y. Markova, "Thermal and excess noise in suspension fibres," *Classical and Quantum Gravity*, vol. 19, pp. 2035–2038, 2002.

[70] J. Kovalik and P. R. Saulson, "Mehcanical loss in fibres for low noise pendulums," *Review of Scientic Instruments*, vol. 64, pp. 2942–2946, 1993.

[71] V. Braginsky, V. Mitrofanov, and S. Vyatchanin, "Isolation of test masses inthe advance laser interferometric gravitational-wave antennae," *Review of Scientific Instruments*, vol. 65, pp. 3771–3774, 1994.

[72] P. Amico, L. Bosi, L. Carbone, L. Gammaitoni, M. Punturo, F. Travasso, and H. Vocca, "Fused silica suspension for the VIRGO optics: Status and perspectives," *Classical and Quatum Gravity*, vol. 19, p. 16691674, 2002.

[73] L. Ju, D. Blair, I. Bilenko, and D. Paget, "Low loss niobium flexure suspension systems," *Classical and Quantum Gravity*, vol. 19, pp. 1703–1708, 2002.

[74] I. Bilenko, L. Ju, D. Paget, and D. Blair, "Niobium flexure suspension design for high q sapphire test masses for future gravitational wave detectors," *Measurement Science and Technology*, vol. 13, pp. 1173–1177, 2002.

[75] L. Ju, D.G.Blair, M. Taniwaki, and R. Andrew, "The quality factor of niobium flexure pendulums," *Physics Letters A*, vol. 254, pp. 239–244, 1999.

[76] L. Ju, M. Baker, and D. Blair, "High-q niobium membrane flexure pendulum," *Physics Letters A*, vol. 280, pp. 182–184, 2001.

[77] P. R, W, J.L.Hall, F.V.Kowalski, J.Hough, G.M.Ford, A.J.Munley, and H.Ward, "Laser phase and frequency stabilisation using an optical resonator," *Applied Physics B*, vol. 31, pp. 97–105, 1983.

[78] C. Salomon, D.Hills, and J.L.Hall, "Laser stabilization at the millihertz level," *Journal of the Optical Society of America*, vol. 5, pp. 1576–1587, 1988.

[79] M. Roberts, P. Taylor, and P. Gill, "Laser linewidths at the sub-hertz level," Tech. Rep. CLM 8, Centre for Length Metrology, National Physicsl Laboratory, 1999.

[80] A. E. Siegman, *Lasers*. Sausalito, CA 94965: University Science Books, 1986.

[81] M. Meade, *Lock-in Amplifiers: Principles and Applications*. Savoy Place, London: Peter Peregrinus, 1983.

[82] G. Galzerano, M. Marano, S. Taccheo, and P. Laporta, "2.1-um lasers frequency stabilized against CO2 lines: Comparison between fringe-side and frequency modulatin locking methods," *Optics Letters*, vol. 28, pp. 248–250, 2003.

[83] J. Hough, D. Hils, M. Rayman, M. L.-S., L.Hollberg, and J. Hall, "Dye-laser frequency stabilization using optical resonators," *Applied Physics B*, vol. 33, pp. 179–185, 1984.

[84] D. Shoemaker, A. Brillet, C. N. Man, O. Cregut, and G. Kerr, "Frequency-stabilized laser-diode-pumped nd:YAG laser," *Optics Letters*, vol. 14, pp. 609–611, 1989.

[85] T. Day, E. K. Gustafson, and R. L. Byer, "Sub-hertz relative frequency stabilization of two-diode laser-pumped nd:YAG lasers locked to a fabry-perot interferometer," *IEEE Journal of Quantum Electronics*, vol. 28, pp. 1106–1115, 1992.

[86] G. Ruoso, R. Storz, S. Seel, S. Schiller, and J. Mlynek, "Nd:YAG laser frequency stabilization to a supercavity a the 0.1 hz level," *Optics Communications*, vol. 133, pp. 259–262, 1997.

[87] E. Bava and F. Massari, "Phase sensitive detection of light reflected from a fabry-perot interferometer," *Review of Scientific Instruments*, vol. 67, pp. 1714–1720, 1996.

[88] E. D. Black, "An introduction to pound-drever-hall laser frequency stabilization," *American Journal of Physics*, vol. 69, pp. 79–87, 2001.

[89] J. Camp, L. Sievers, R. Bork, and J. Heefner, "Guided lock acquisition in a suspended fabry-perot cavity," *Optics Letters*, vol. 20, pp. 2463–2465, 1995.

[90] M.J.Lawrence, B.Wilke, M. Gustafson, and R. Byer, "Dynamic response of a fabry-perot interferometer," *Journal of the Optical Society of America*, vol. 16, pp. 523–532, 1999.

[91] D. Ottaway, P. Veitch, M. Hamilton, C. Hollitt, D. Mudge, and J. Munch, "A compact injection-locked nd:YAG laser for gravitaitonal wave detection," *IEEE Jouranl of Quantum Electronics*, vol. 34, pp. 2006–2009, 1998.

[92] D. Ottaway, P. Veitch, C. Hollitt, D. Mudge, M. Hamilton, and J. Munch, "Frequency and intensity noise of an injection-locked nd:YAG ring laser," *Applied Physics B*, vol. 71, pp. 163–168, 2000.

[93] New Focus Inc., *Model 4003 Electro-Optic Modulator Manual*, 1999.

[94] B.Wilke, S.Brozek, K.Danzmann, V.Quetschke, and S.Gossler, "Frequency stabilization of a monolithic Nd:YAG ring laser by controlling the power of the laser-diode pump source," *Optics Letters*, vol. 25, pp. 1019–1021, 2000.

[95] InnoLight GmbH, *Mephisto Series Users Manual*, 3.1 ed., 1999.

[96] G. C. Goodwin, S. F. Graebe, and M. E. Salgado, *Control System Design.* Upper Sadle River, New Jersey: Prentice Hall, 2001.

[97] G. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems.* Addison Wesley, 1994.

[98] J. Helmcke, S. A. Lee, and J. L. Hall, "Dye laser spectrometer for ultra-high spectral resolution: Design and performance," *Applied Optics*, vol. 21, pp. 1686–1694, 1982.

[99] C. Motchenbacher and J. Connelly, *Low-Noise Electronic System Design.* Wiley-Interscience, 1993.

[100] E. Calloni, F. Barone, L. D. Fiore, A. Grado, P. L. Penna, and L. Milano, "Effects of misalignment and beam jitter in fabry-perot laser stabilization," *Optics Communication*, vol. 142, pp. 50–54, 1997.

[101] R. C. Weast, ed., *Handbook of Chemsirty and Physics.* Cranwood Parkway, Cleveland Ohio: The Chemical Rubber Company, 1968.

[102] S. Kawamura, A. Abramovici, and M. Zucker, "Improved multistage wide band laser frequency stabilization," *Review of Scientific Instruments*, vol. 68, pp. 223–229, 1997.

[103] D. Shaddock, M. Gray, and D. McClelland, "Frequency locking to an optical cavity using spatial mode interference," *Optics Letters*, vol. 24, pp. 1499–1501, 1999.

[104] E. Bava, G. Galzerano, and C. Svelto, "Frequency-noise sensitivity and amplitude-noise immunity of discriminators based on fringe-side fabry-perot cavities," *IEEE Transactions of Ultrasonics, Pyroelectrics and Frequency Control*, vol. 49, pp. 1150–1159, 2002.

[105] P. C. D. Hobbs, "Ultrasensitive laser measurement without tears," *Applied Optics*, vol. 36, pp. 903–920, 1997.

[106] M. Houssin, M. Jardino, and M. Desaintfuscien, "Comparison of the calculated transient responses of a FabryPerot used in reflection and in transmission," *Review of Scientific Instruments*, vol. 61, pp. 3348–3352, 1990.

[107] P. Tremblay and R. Ouellet, "Frequency response of a fabry-perot interferometer used as a frequency discriminator," *IEEE transactions on intrumentation and measurement*, vol. 40, pp. 204–207, 1991.

[108] M. B. Gray, D. A. Shaddock, C. C. Harb, and H.-A. Bachor, "Photodetector designs for experiments in quantum optics," *Review of Scientific Instruments*, vol. 69, pp. 3755–3764, 1998.

[109] M. Tröbs, *Laser Development and Stabilization for the Spaceborne Interferometric Gravitational Wave Detector LISA*. PhD thesis, Univsersity of Hannover, 2005.

[110] F. Seifert, P. Kwee, M. Heurs, B. Willke, and K. Danzmann, "Laser power stabilization for second-generation gravitational wave detectors," *Optics Letter*, vol. 31, pp. 2000–2002, 2006.

[111] R. S. Abbott and P. J. Kinga, "Diode-pumped nd:YAG laser intensity noise suppression using a current shunt," *Review of Scientific Instruments*, vol. 72, pp. 1346–1349, 2001.

[112] I. Zawischa, M. Brendel, K. Danzmann, C. Fallnich, M. Heurs, S. Nagano, V. Quetschke, H. Welling, and B. Willke, "The GEO 600 laser system," *Classical and Quantum Gravity*, vol. 19, pp. 175–1781, 2002.

[113] J. Rollins, D. Ottaway, M. Zucker, R. Weiss, and R. Abbott, "Solid-state laser intensity stabilization at the 10-8 level," *Optics Letters*, vol. 29, pp. 1876–1878, 2004.

[114] M. Heurs, V. Quetschke, B. Wilke, K. Danzmann, and I. Freitag, "Simultaneously suppressing frequency and intensity noise in a nd: YAG nonplanar

ring oscillator by means of the current-lock technique," *Optics Letters*, vol. 29, pp. 2148–2150, 2004.

[115] J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Laser stabilization to a single ion," in *Proc. Int. School of Physics "Enrico Fermi", Frontiers in Laser Spectroscopy* (T. W. Hänsch and M. Inguscio, eds.), (Amsterdam, The Netherlands), pp. 359–376, 1994.

[116] M. W. Hamilton, "An introduction to stabilized lasers," *Contemporary Physics*, vol. 30, pp. 21–33, 1989.

[117] Y. A. engel, *Introduction to Thermodynamics and Heat Transfer.* McGraw-Hill, 1997.

[118] Object Management Group Incorporated, *OMG United Modeling Language Specification*, 1.5 (03-03-01) ed., 2003.

[119] Atmel Corporation, *8-Bit Microcontroller with 4K Bytes of In-System Programmable Flash - AT90S4433 AT90LS4433*, 2003.

[120] S. Shinners, *Advanced Modern Control Theory.* Wiley Interscience, 1st ed., 1998.

[121] K. Ogata, *Modern Control Engineering.* Prentice Hall, 2nd ed., 1990.

[122] R. Dorf and R. Bishop, *Modern Control Systems.* Prentice Hall, 9th ed., 2000.

[123] EGG Optoelectornics, *C30619, C30641, C30642, C30665 Large-Area InGaAs Photodiodes.*

[124] Epitaxx, *ETX1000T Dataheet.*