Opto-Mechanical Noise Cancellation

Conor Mow-Lowry

A thesis submitted for the
degree of Bachelor of Science with Honours in physics
at the Australian National University

November 2002

Declaration

This thesis is an account of research undertaken between February 2002 and November 2002 at the Department of Physics and Theoretical Physics, Faculty of Science, Australian National University.

The research has been conducted under the supervision of Professor David McClelland amd Dr Malcolm Gray. Unless otherwise stated, the work presented within this thesis is my own.

Conor Mow-Lowry
November, 2002

Acknowledgments

I sincerely thank my supervisors Professor David McClelland and Dr Malcolm Gray.

David I thank for more than his supervision and scholarship this year. Without his continued guidance and support over the past three years I may never have reached honours at all. Thank you David.

I thank Mal Gray for patiently providing assistance wherever and whenever it was needed. He saved many, many hours of wasted effort with his practical knowledge of all things interferometric and electronic. Dr Positive's aureate presence has lifted spirits on many an otherwise dull night in the lab.

Dr Daniel Shaddock I thank in absentia. I feel much credit for my present knowledge of gravitational wave research is due to an unthinkable number of whiteboard sessions with Dr Dan.

For his wealth of knowledge and a plethora of helpful suggestions I thank Dr Stan Whitcomb. Without Stan's ability to solve the insoluble, I imagine that many parts of my experiment would still be a mystery to me.

For spending an unspeakable number of hours in the lab on (how many?) all night data recording runs, I thank Ben Sheard. Without Ben, I might not have hurt myself playing footy in the driveway at 5:30 in the morning. Thanks Ben.

Thanks to the other members of the physics department, who have tolerated my interruptions and provided both humour and advice. In particular Dr Ping Koy Lam for being so approachable despite his busy schedule; Dr John Close for beer, barbeques, and many fun conversations; and Nick for being the only other person mad enough to consistently work at 4 in the morning.

My other graviteer colleagues are always ready to offer humour or assistance. Especially, Bram, Glenn, Jong, Max, and Andrew - the honorary graviteer for putting up with me in their work space, but also to Dr Scott, Antony, Benedict, and the venerable Emeritus Professor Sandeman, for helping to make the gravitational waves research group *the* place to be.

The whole honours class also receives my thanks for providing a cheerful and good natured group to work with. Thanks Annabel, Anne, Aśka, Doug, Kirk, Nicolai, Phil, Simon, Taira, and Tim.

For proof reading, a strain on the brain and great sedative all in one, I greatly appreciate the efforts of Mal, Stan, David (well, I think he read it), and my Mum, Karen Mow. Any remaining mistakes are my flaunt.

To my family and friends outside the rather bizarre world of physics, I extend my sincere gratitude. Most of all, Deon, Mum, and Esther, for keeping me sane, fed, and civilised respectively.

Cheers, it's been fun.

Abstract

The quest to directly detect gravitational waves is pushing the sensitivity of interferometric measurements all the way to their quantum limits. It is expected that a decade from now, gravitational wave detectors will be limited in sensitivity by quantum fluctuations of the optical fields. They will operate at the standard quantum limit, or SQL, where shot noise in the measured light and radiation pressure noise on the mirrors cause equal errors in position measurement. One possible way to breach the SQL is to correlate the quantum fluctuations of the optical fields which cause shot noise and radiation pressure noise, and to have these two effects cancel. However, there are many steps required before the SQL is reached, let alone breached.

The experiments presented in this thesis investigate the interaction between radiation and an optical cavity, in which one mirror of the cavity is mounted on a flexure which could be moved by radiation pressure. The cavity was shown to exhibit non-linear behaviour with high input power. The radiation pressure force was shown to change the mechanical resonance frequency of the moveable mirror. Motion was induced through amplitude modulation of a high power input beam and the extent of this motion measured using the cavity control loop.

To demonstrate the way quantum correlations could be used to beat the SQL, the laser light incident on the cavity was prepared, using classical modulation techniques, with classical correlations between the quadratures that cause shot noise and radiation pressure noise. A level of modulation much higher than the quantum level was used to make the cancellation effects more visible.

The effect of radiation pressure induced motion was cancelled by the addition of correlated frequency modulation. The input amplitude was then modulated by a white noise source. The resulting noise was partially cancelled when the same white noise source was used to drive the frequency modulator with a different phase. This cancellation demonstrably improved the signal to noise ratio of a signal injected into the system.

Contents

Chapter 1

Introduction

Almost a century ago Albert Einstein revolutionised the scientific model of space and gravity with his general theory of relativity [1]. One of the many predictions made by this theory is that gravitational radiation is produced by accelerating masses [2]. It is the goal of gravitational wave detectors to, first, measure the perturbations caused by these gravitational waves, and second, use these measurements in a new form of astronomy, gravitational wave astronomy.

1.1 Gravitational Wave Detection

Gravitational radiation bears strong similarities to electromagnetic radiation [3]. A static mass has a static gravitational field, much the same as a static charge has a static electric field. If the mass or charge is moved, the field does not instantly follow the perturbation, rather the perturbation signal propagates outwards at the speed of light. For the moving charge, the perturbation signal is in the form of electromagnetic radiation. For the moving mass, it is in the form of gravitational radiation. Unlike electric charges, which may be positive or negative, mass is always positive. For this reason the lowest mode of oscillation of a gravitational wave is quadrupole [4], whereas electromagnetic waves may have dipole radiation.

The existence of gravitational waves is of critical importance to the general theory of relativity. One of Einstein's prime motivations for developing the theory of relativity was to make gravitational systems behave causally; to have no signals propagate faster than the speed of light. The direct measurement of gravitational radiation will provide direct evidence for the causality and locality of the gravitational force. It is hoped that by measuring the exact nature of astrophysical gravitational radiation, the exact nature of the source may be determined. Gravitational wave astronomy may provide radical and otherwise unobtainable information about the universe.

To illustrate the effect of gravitational waves interacting with free masses, consider a ring of freely floating stationary test particles. In the absence of any gravitational radiation, the particles form a circle. When a gravitational wave passes through the ring, with its axis of propagation perpendicular to the plane of the particles, the ring is compressed, first in one axis, then it returns to a circle before being compressed in the orthogonal axis. This process is shown at quarter period intervals in figure 1.1.

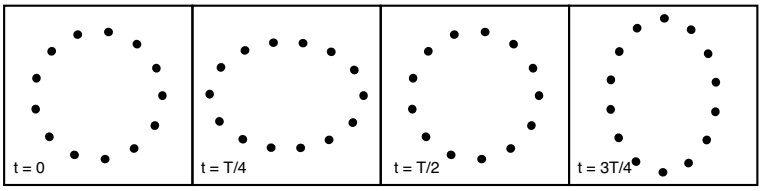

Figure 1.1: The effect of a gravitational wave passing through a ring of test particles is shown at quarter period intervals.

The most promising modern gravitational wave detectors are laser interferometers. They use a Michelson interferometer to sense the position of test masses. The test masses move under the influence of gravitational radiation in much the same way that the particles in figure 1.1 move. The interferometer then converts the distance disturbance into a power change at the output of the interferometer. Gravitational wave strength is measured in terms of strain (h), or the fractional length change, defined by:

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

Thus, it is possible to increase the signal to noise ratio of a gravitational wave measurement by increasing the length of the detector and keeping the absolute length sensitivity constant.

The sensitivity required to detect gravitational radiation is unprecedented. The largest regular events, expected to occur only a few times a year, are predicted to exert a strain of 10^{-21} to 10^{-22}. Even over a length of several kilometres this requires a measurement sensitive to length changes of 10^{-18} m, or approximately one thousandth of a protonic diameter.

Despite this seemingly impossible requirement, there are several long baseline laser interferometric gravitational wave detectors nearing completion. These detectors are enormous Michelson interferometers. The arms of the interferometers are contained in vacuum tubes several kilometres long. The LIGO project consists of a three detectors at two sites in America, a 4 km long detector in Livingstone, Louisiana, and a pair of detectors, 2 km and 4 km long, in Hanford, Washington [5]. VIRGO is a 3km long detector near Pisa, Italy [6]. GEO600 is a 600 m long detector near Hanover, Germany [7]. TAMA300 is a 300 m long detector in Tokyo, Japan [8]. Both the TAMA and LIGO projects have recorded scientific data, though neither are expected to reach target sensitivities for some time yet.

Even if the first generation of gravitational wave detectors consistently detect gravitational waves, they will not have enough signal to noise ratio to be used for astrophysical study. There are upgrades planned for this first generation which are expected to increase sensitivity by an order of magnitude.

Advanced LIGO, the second generation of the LIGO project, is expected to have a noise floor limited by the quantum fluctuations of the input electric field [9]. To extend the sensitivity of a third generation of detectors it is necessary to breach this quantum limit.

1.2 The Standard Quantum Limit

The phase and amplitude of an optical field are non-commuting variables, and as such they have an uncertainty relationship. The product of the uncertainties is constrained to a minimum level. Laser beams are very close to the coherent state in nature. The coherent state has equal uncertainty in the two quadratures of phase and amplitude. A common way to visualise this quantum noise is through the 'ball and stick' picture. The stick represents the coherent, unchanging amplitude of the laser. The ball is the quantum noise, and it is attached to the end of the stick. The net result is that the combined ball and stick varies equally in angle and length as shown in figure 1.2.

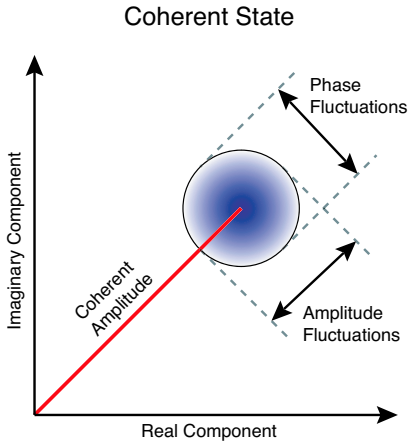

Figure 1.2: The ball and stick picture of the quantum noise of a coherent state. The area of the shaded region has a minimum size.

Laser interferometers, in their standard form, measure distance using the phase of a coherent state. Thus, the quantum phase fluctuations of the electric field will add some noise to the measured signal. The noise induced by phase fluctuations is called *shot noise*. The amplitude fluctuations also couple into the signal. As the amplitude varies, the radiation pressure force on the end mirrors of the interferometer changes, causing the mirrors to move. The noisy motion of the mirrors induced by quantum amplitude fluctuations is called *radiation pressure noise*.

The signal to noise ratio of a shot noise limited measurement is proportional to the square root of the laser power. However, the signal to noise ratio of a radiation pressure noise (or RP noise) limited measurement is inversely proportional to the laser power:

$$\begin{aligned} SNR_{shot\ noise} &\propto \sqrt{P} \\ SNR_{RP\ noise} &\propto \frac{1}{\sqrt{P}} \end{aligned} \tag{1.2}$$

where SNR is the signal to noise ratio for the given noise source and P is the optical

power. Since these two noise sources scale in opposite directions as a function of power, they will be equal for some optical power. At this power, the total quantum noise limited sensitivity is optimal. This limiting sensitivity is known as the *standard quantum limit*, or SQL [10, 11]. The SQL is not a true limit. As early as 1980, a means of breaching the standard quantum limit was proposed [12].

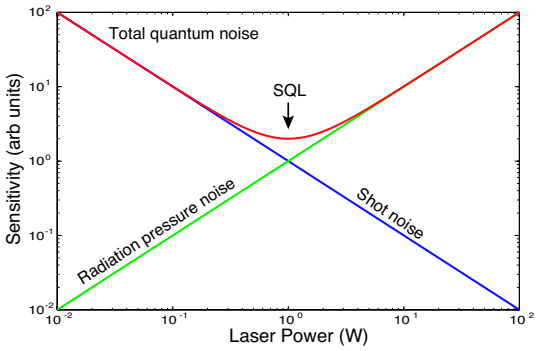

Figure 1.3: The SQL is measured when quantum noise is the limiting noise source and when shot noise and radiation pressure noise are equal in magnitude.

One method for breaching the SQL in a laser interferometer involves squeezing. The quantum fluctuations of the coherent state are the same in both quadratures. It is possible to use non-linear methods to 'squeeze' the ball of quantum noise into an ellipse. Amplitude and phase squeezing are shown in figure 1.4.

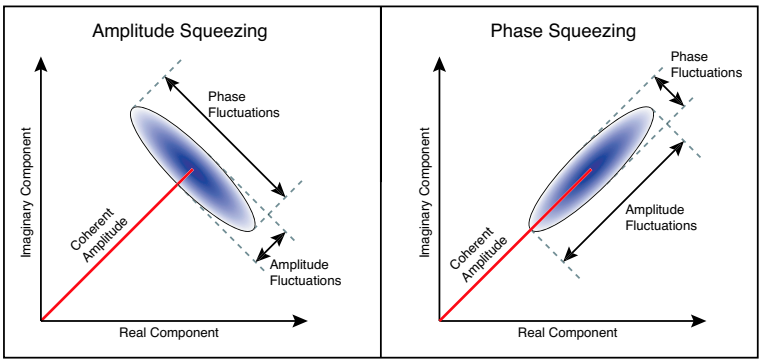

Figure 1.4: Quantum noise on a squeezed beam. The noise in one quadrature is reduced, at the expense of the other.

The injection of either amplitude or phase squeezing into an interferometer does not reduce the SQL, since it decreases one noise source at the expense of the other. The limiting sensitivity is the same, though the power required to reach that sensitivity is different. However, if squeezing is injected with a 45 degree angle (relative to the phase and amplitude axes), it serves to correlate the fluctuations in the two quadratures, as shown in figure 1.5. If the fluctuations are correlated, then their

effect on the output can destructively interfere. Since shot noise and radiation pressure noise are equal at the SQL, the cancellation should be limited only by the degree of correlation.

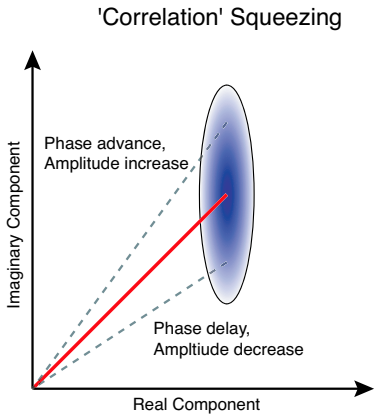

Figure 1.5: When the squeezing angle is at 45 degrees, fluctuations in the two quadratures are correlated.

It is through the injection of this 45 degree squeezing into a system at the SQL that the Australian National University gravitational wave research group plans to breach the standard quantum limit. However, the standard quantum limit has not yet been reached for a macroscopic mass, though several attempts have been made [13, 14]. The very fact that a standard quantum limited measurement has not yet been made is testament to the difficulty of the task.

1.3 Opto-Mechanical Noise Cancellation

To demonstrate the way quantum correlations could be used to breach the SQL, this project aimed to imprint a laser with classical correlations in the same quadratures as radiation pressure noise and shot noise. This classical analogue is called 'classical opto-mechanical noise cancellation' and is directly analogous to breaching the SQL. The quantum fluctuations are replaced by phase and amplitude modulation. The size of the modulations may then be far larger than their quantum equivalents, making the cancellation far more visible. The modulations can be arbitrarily correlated without the need of squeezing. These differences make classical, opto-mechanical noise cancellation possible to observe in a bench-top environment.

The most challenging part of this project is to design a system influenced by, and sensitive to radiation pressure. There are very few experimental papers which study the effects of these radiation pressure interactions. A common reference, Dorsel *et. al.* 1983, observes radiation pressure effects in a passive system [15]. Acquisition of expertise with radiation pressure sensitive systems is an important step in reaching the SQL. In addition the study of radiation pressure interactions is fundamentally interesting in its own right.

1.4 Overview of Thesis Structure

Chapter 1 provides a cursory examination of gravitational wave detection and the standard quantum limit. The quest to reach the SQL is given as the motivation for the experiment.

Chapter 2 outlines the theory of Hermite-Gaussian beams and presents a simple analysis of the Michelson interferometer and the Fabry-Perot cavity.

Chapter 3 presents two common means to control interferometers, offset locking and RF locking. The analysis of these schemes is limited to their application in the experimental work.

Chapter 4 describes the design of the moveable mirror used to sense radiation pressure. Intrinsic in the design process is the development of a model to predict the size and nature of radiation pressure interactions.

Chapter 5 presents the experimental design in detail. The three main sections being environmental noise suppression; optical design; and control schemes.

Chapter 6 displays the key data recorded during the experiment, including the measurement of opto-mechanical noise cancellation.

Chapter 7 discusses what further work is necessary to remove any ambiguity from the results and to increase understanding of the system. A brief summary of the project is also given.

The appendices provide information about the pre-project design of the flexure, the radiation pressure model, and the electronic schematic of the PZT servos.

Chapter 2

Interferometry

The field of interferometry began more than a century ago with such revolutionary experiments as the Michelson-Morely experiment in 1887 [16]. Since then, interferometry has become a widely used, high sensitive measurement technique. This chapter provides a simple introduction to two important and simple interferometers, the Michelson interferometer, and the Fabry-Perot cavity. A brief introduction to Gaussian beams is presented before these interferometers are analysed.

2.1 Hermite-Gaussian Beams

Under ideal conditions, a laser will produce a Gaussian beam at its output. To understand the interaction of lasers with optics it is useful to understand Gaussian beams. This section introduces the Gaussian beam and the basis set the Hermite-Gaussian beams, which contains the Gaussian beam is its first member [17, 18]. Hermite-Gaussian beams are not only important because they provide a basis set for optical fields, but also because their modes come very close to describing the transverse modes of optical resonators.

Gaussian beams have the smallest angular divergence of the wavefront normals necessary to satisfy the wave equation [19]. Thus, they are very ray like in behaviour, making them easy to work with experimentally and allowing many simplifying approximations to be made about their behaviour. Gaussian beams reproduce themselves under Fourier transform. The far- and near-field intensity distributions have the same form, and transformation through a lens produces another Gaussian beam.

There are several important parameters used to describe a Gaussian beam. For description of these parameters, and for further work, the axis of propagation is assumed to be the z-axis, with the x- and y-axes describing the transverse profile of the beam. The line $x = y = 0$ defines the peak transverse intensity of the beam.

The waist is the point along the axis of propagation where the transverse profile is spatially smallest. The Rayleigh length (or Rayleigh range), z_0, is the distance along the axis of propagation from the waist to the point where the beam is twice as large in cross sectional area as it is at the waist. The beam size (or spot size) is a measure of the transverse size of the beam at a given point on its axis of propagation.

The spot size is defined by the points where the electric field amplitude has dropped to be a factor of $1/e$ times its amplitude on the line $x = y = 0$. The spot

size at the waist (or waist size) is labelled w_0 and the general spot size is $w(z)$. The radius of curvature describes the curvature of the beam on the axis of propagation. Though the beam curvature is actually parabolic, the distinction between parabolic and spherical curvature is unimportant around the centre of the intensity profile (i.e. near $x = y = 0$). The spherical curvature assumption is known as the paraxial approximation. The radius of curvature is always expressed as a function of z, $R(z)$, so as not to confuse it with the intensity reflectivity of a mirror, which also has the symbol R. The following equations show how these parameters relate to each other:

$$z_0 \equiv \frac{\pi {w_0}^2}{\lambda} \tag{2.1}$$

$$w(z) = w_0 \left(1 + \frac{z^2}{{z_0}^2}\right)^{\frac{1}{2}} \tag{2.2}$$

$$R(z) = z\left(1 + \frac{{z_0}^2}{z^2}\right) \tag{2.3}$$

$$\eta(z) = \arctan\left(\frac{z}{z_0}\right) \tag{2.4}$$

The final equation defines $\eta(z)$ which is a phase term known as the Gouy phase. It describes the phase shift between the Gaussian beam and a spherical wave which propagates from the centre of the waist. This phase shift is caused by the diffraction of the beam.

With these parameters defined it is possible to describe a Gaussian beam in Cartesian co-ordinates. However, it is simpler to first define the co-efficients needed to create the Hermite-Gauss modes, and to define all the Hermite-Gaussian beams simultaneously. The Hermite-Gauss modes are 'variants' of the Gaussian beam whose transverse profile is described by the Hermite polynomials. The first few Hermite polynomials are:

$$\begin{aligned} H_0(x) &= 1 \\ H_1(x) &= 2x \\ H_2(x) &= 4x^2 - 2 \\ H_3(x) &= 8x^3 - 12x \end{aligned} \tag{2.5}$$

Further terms can be defined from these using the relation:

$$H_{n+1}(x) = 2xH_n(x) - 2nH_{n-1}(x) \tag{2.6}$$

Using all the above equations it is now possible to completely describe the electric field of a Hermite-Gaussian beam (real and imaginary components) at any point in

space using the equation [18]:

$$E_{lm}(x,y,z) = E_0 \frac{w_0}{w(z)} H_l\left(\sqrt{2}\frac{x}{w(z)}\right) H_m\left(\sqrt{2}\frac{y}{w(z)}\right) \times \exp\left[-\frac{x^2+y^2}{w^2(z)} - \frac{ik(x^2+y^2)}{2R(z)} - ikz + i(l+m+1)\eta(z)\right] \tag{2.7}$$

Where k=$2\pi/\lambda$ is the wavenumber of the laser light. Notice that the argument of the Hermite polynomials is no longer a single variable. The l and m terms refer to the mode numbers. Each mode has a different intensity profile in the transverse directions and experiences a different Gouy phase shift. These modes are usually referred to as the TEM_{lm} modes. Figure 2.1 shows the transverse profile of the TEM_{00} and TEM_{10} modes at the $z = 0$ plane. Notice the nodal line at $x = 0$ for the TEM_{10} mode.

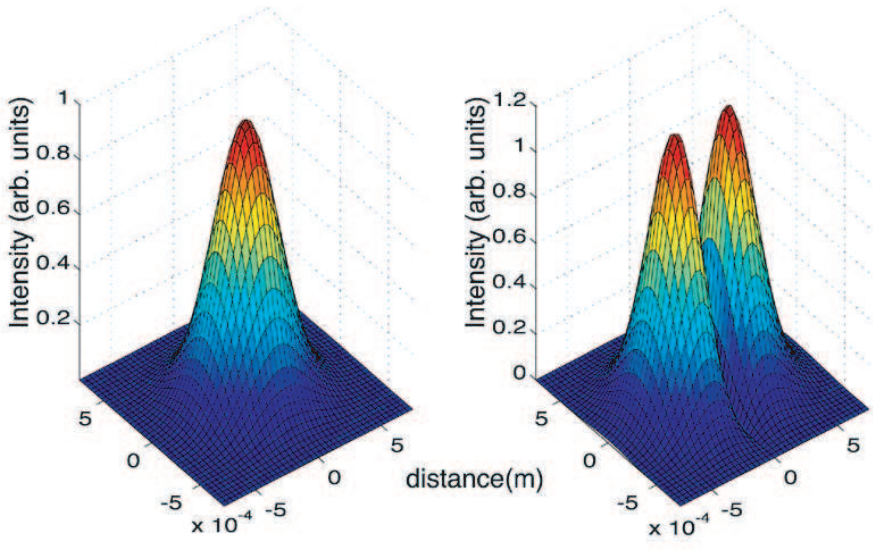

Figure 2.1: The transverse intensity distribution of the TEM_{00} and TEM_{10} modes. Both spatial axes are in units of metres (x10^{-4}).

The diffraction profile of the Hermite-Gaussian beams is illustrated in figure 2.2. For the TEM_{00} and TEM_{10} modes a cross section on the $y = 0$ plane is taken. The z-axis points into the page, and the figure shows a view 'down' the laser, showing the beams contracting to a waist (of 300μm with λ =1064 nm), and diverging downstream.

Equation 2.7 defines a beam azimuthally symmetric around the z-axis. As an extension to this formula it is possible to define an elliptic beam. In this case, the x and y axes have independent Rayleigh lengths, waist sizes, and waist positions. It is then possible to calculate the independent spot sizes, radii of curvature, and Gouy phases for each axis using equations 2.2 to 2.4. Elliptic beams may be converted to circular beams, and vice versa, using cylindrical lenses or tilted spherical lenses.

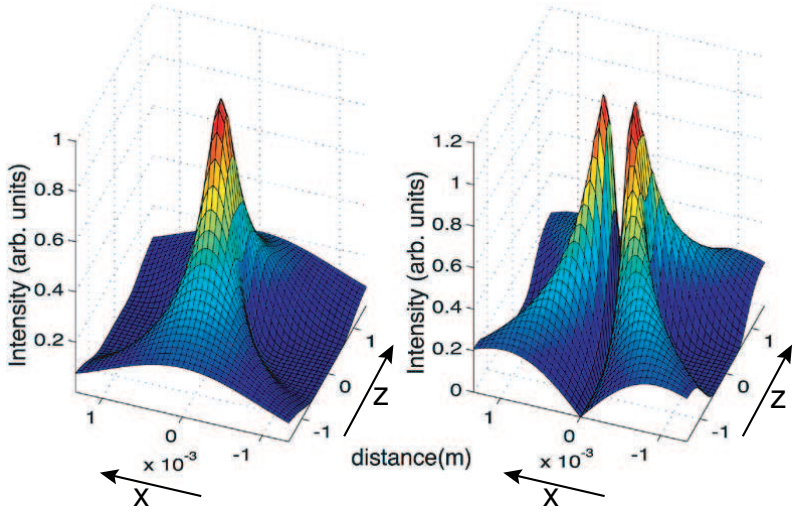

Figure 2.2: The intensity profile of the TEM_{00} and TEM_{10} modes along the axis of propagation. The z-axis is the axis of propagation with unit of metres, and the x-axis is a transverse axis with units of metres ($x10^{-3}$).

Experimentally the full description of Gaussian beams is most important for optical resonators, such as the Fabry-Perot cavity discussed below. For most other experimental purposes, laser beams are well enough collimated to be treated as rays. In deriving equations in the ensuing sections it will often be assumed that there are no adverse effects from mismatched spot sizes, or imperfect wavefronts. These assumptions are quite robust, with any error usually resulting in the addition of loss terms with each imperfect interaction. If sufficient care is taken during optical design, these losses are minimised.

2.2 The Michelson Interferometer

The Michelson interferometer has a long history from its first use in the famous Michelson-Morley experiment to current gravitational wave detectors. The Michelson interferometer in its simplest modern form consists of a laser, a beamsplitter (usually with equal reflectivity and transmissivity), and two highly reflective end mirrors. A detector at the output diagnoses the condition of the interferometer by measuring the optical power.

Interferometry is remarkably useful because it converts phase information into intensity information. An interferometer measures length differences with sensitivity on the same scale as the wavelength of light (which is both small and difficult to measure), and converts this into optical power differences, which are easy to measure. In exchange for this high sensitivity, good alignment, and a stable operating environment are required.

Successful interferometry requires good fringe visibility. Fringe visibility is a

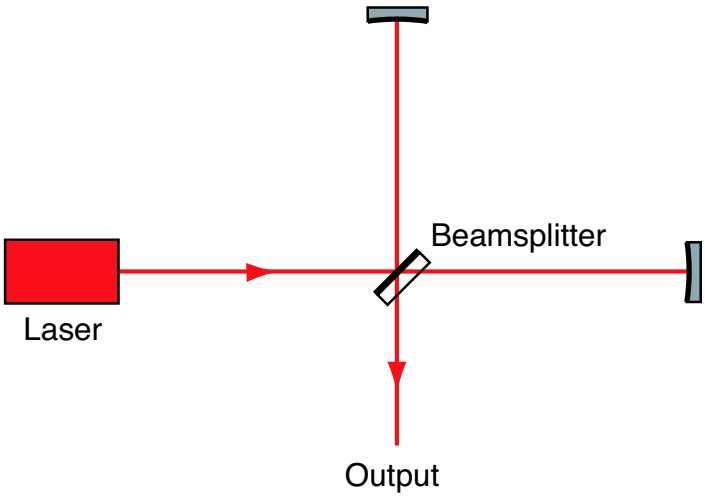

Figure 2.3: The simplest form of Michelson interferometer consists only of an input light source, a beamsplitter, and two arm mirrors.

measure of how efficiently the beams in each arm interfere when they return to the beamsplitter. The fringe visibility is given by:

$$F.V. = \frac{P_{max} - P_{min}}{P_{max} + P_{min}} \tag{2.8}$$

Where P_{max} is the output power when the light in the arms constructively interferes at the beamsplitter, and P_{min} is the output power when it destructively interferes. A fringe visibility of 1 implies perfect interference.

There are several factors which affect fringe visibility. The first is coherence. If the travel time in the arms is different, then the light emitted by the source at one time will interfere with light emitted at a different time. If the source is coherent over time scales longer than the difference in travel times, the two beams can still interfere efficiently. Experimentally there are two extreme limits for coherence: a white light source and a laser. For the white light source, which is incoherent, the arm lengths (and hence travel time) must be exactly matched for interference to occur. For a laser, the coherence time (or length) can be as high as a millisecond (or 300 km in length). Thus it is possible to have vastly different arm lengths, and still have coherent interference when a laser is used.

The next factor affecting fringe visibility is overlap. Two factors affect overlap, the alignment, and the spot size. Firstly, the two end mirrors must reflect the light in the arms back to the same spot on the beamsplitter, otherwise the beams will simply not overlap. Secondly, since laser light diffracts, if the arms are different lengths, the beams returning from each arm will be different sizes. To allow for this it is necessary to either; match arm lengths, place correcting lenses in the arms (or use curved end mirrors), or choose an input beam such that the spot size difference

is small.

The final factor (for the purpose of this thesis) in fringe visibility is beam distortion. If the beams are distorted differently in each arm, by a reflection from an imperfect mirror, for example, they will interfere imperfectly. These three factors combine to reduce the interference efficiency, and thus decrease the fringe visibility.

2.2.1 Interference Condition

In this section the output intensity from a simple Michelson as a function of arm lengths is given. The simplifying assumptions used are:

•Perfect coherence
•Matched spot sizes
•No beam distortions and
•Lossless beamsplitter and end mirrors

Deviations from these conditions result in reduced fringe visibility, as discussed in the previous section. For this and other calculations, the convention used is where reflected beams obtain no phase shift and transmitted beams obtain a $\pi/2$ phase shift.

For a Michelson interferometer consisting of a beamsplitter with electric field amplitude reflectivity r, and transmissivity $t = \sqrt{1-r^2}$, with input electric field amplitude E_{inc}, and arms of length L_1 and L_2, the output electric field is given by:

$$E_{out} = irtE_{inc}\left(\mathrm{e}^{\frac{i2\omega L_1}{c}} + \mathrm{e}^{\frac{i2\omega L_2}{c}}\right) \tag{2.9}$$

where ω is the angular frequency of the optical field. Here, and further on in this thesis, the electric field will be defined to have units of $\sqrt{\text{Watts}}$, so that the optical power is given by:

$$P = E^*E \tag{2.10}$$

Defining the electric field in this way is useful because photodetectors produce an output voltage proportional to incident optical power.

Equation 2.9 can be simplified by using the average arm length and arm length difference instead of the individual arm lengths. The values for average length and arm length difference are defined in the obvious way:

$$\begin{aligned} L &= \frac{L_1 + L_2}{2} \\ \Delta L &= L_2 - L_1 \end{aligned} \tag{2.11}$$

Solving for L_1 and L_2 and substituting into equation 2.9 gives:

$$E_{out} = 2irtE_{inc}\mathrm{e}^{\frac{i2\omega L}{c}}\left(\cos\frac{\omega\Delta L}{c}\right) \tag{2.12}$$

The interference condition, and output power, of the Michelson interferometer is only affected by the ΔL parameter. Thus to change the output power, it is necessary to change the arm length difference. Changing the average arm length has the effect

of changing the phase of the output light, since this alters the total travel time from the input to the output of the Michelson.

To better illustrate the effects of changing length values, figure 2.4 shows the output power and electric field phase from a Michelson interferometer. In the first column the average arm length is increased, in the second the arm length difference is increased, and in the third the length of one arm is increased. For ease of plotting the length axes are centered around zero, however, this represents a microscopic deviation around a macroscopic average value of $L_1 = L_2 = 0.50$ m.

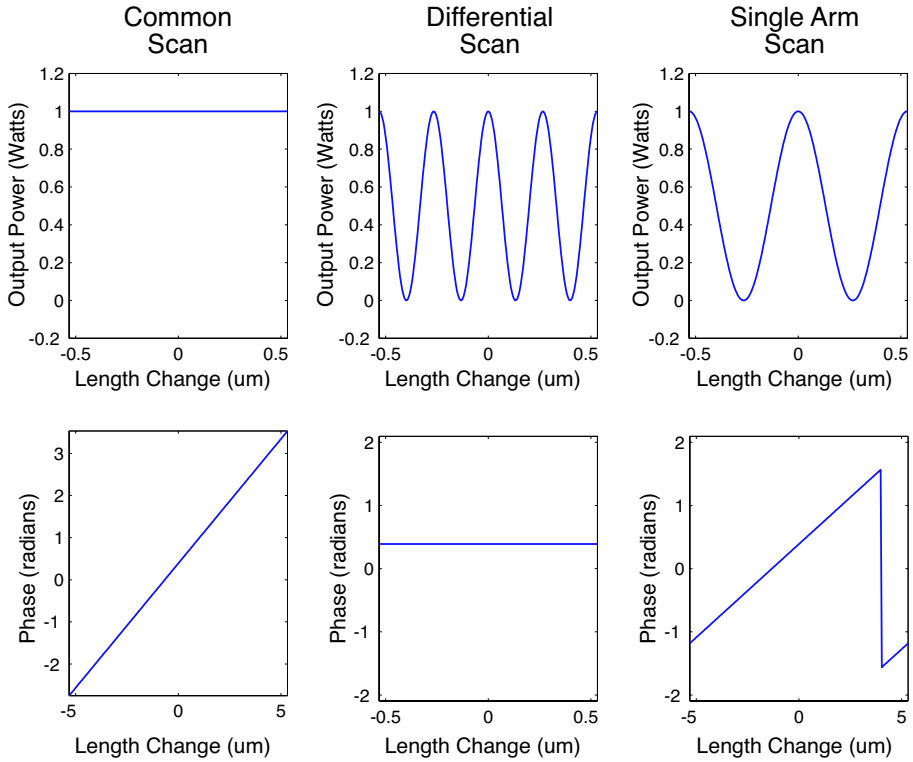

Figure 2.4: The output power and phase of a Michelson interferometer as the average arm length, the arm length difference, or the length of one arm are increased. Input power is unity.

Despite the antiquity of the Michelson interferometer it continues to be of relevance to modern optics. All the current generation of laser interferometric gravitational wave detectors are advanced Michelson interferometers. It would be ironic if at the start of this century, the quest to directly detect gravitational waves should fail, and the Michelson interferometer would again force a change to the way space is understood by returning a null result.

2.3 The Fabry-Perot Cavity

The Fabry-Perot cavity is another form of interferometer. Unlike its cousin the Michelson, which is a two beam interferometer, the optical cavity is a multiple

beam interferometer. In its standard form, a Fabry-Perot cavity consists of two mirrors, front and back, with laser light incident on the front mirror. Once the system reaches steady state, typically on sub micro-second timescales, there are three fields of interest, the reflected field, the circulating field, and the transmitted field. This section is dedicated to explaining the behaviour of these fields.

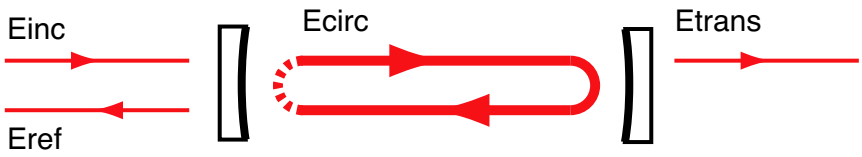

Figure 2.5: A schematic diagram of a Fabry-Perot cavity showing the three electric fields of interest.

Conceptually it is simplest to analyse a cavity as a multiple 'bounce' system [20]. In this picture, a small 'packet' of light is injected into the cavity. It circulates around the cavity, reflecting off (and transmitting through) both mirrors before interfering with the 'second' packet of light, though the system actually consists of a continuous chain of such packets. If this process is iterated infinitely, and the phasors of these packets are vector summed, a steady state solution is reached.

It is through the multiple bounce method that the effective bounce number can be determined. Since there is a fixed round-trip loss for each individual packet, there will be a characteristic number of bounces before the packet is considered insignificant. The bounce number also defines a time constant for each packet, since each round trip of the cavity takes a known time.

If the phasors for all the packets inside the cavity add up exactly in phase, the circulating field will become large compared with the input field. The circulating field increases until the power lost in each round trip is equal to the input power. This condition defines resonance - when the circulating field is at a maximum for a given input field.

2.3.1 Cavity Resonance Conditions

There are two requirements for a cavity to be on resonance. The first is that the beam is self-reproducing spatially (the transverse mode condition), and the second is that the circulating field components interfere constructively (the longitudinal mode condition).

The transverse condition is satisfied when the beam is replicated identically after completing one round trip of the cavity. Equivalently, if the radius of curvature and spot size are the same after one round trip, then the transverse condition is satisfied. For example, if a cavity consists of two flat mirrors, then after one round trip the beam will have diffracted and it will not be identical to the incident beam at the same position. However, if the back mirror is curved with the same radius of curvature as

the circulating field at that point, then the beam is 'folded' exactly back on itself, and the beam is replicated at all points in the cavity. Thus, most cavities are made with spherically curved mirrors.

The longitudinal condition requires that each phasor in the circulating field have the same phase. An equal statement requires the cavity round trip be an integer number of wavelengths in length. Thus, when a beam travels a full round trip of the cavity it is in the same part of its cycle as the new input beam. There are two ways to achieve this condition, either by changing the cavity length, or changing the input wavelength.

The transverse condition gives rise to the cavity stability criterion. There are some cavities for which there are no stable resonant modes. A cavity with one flat mirror and one convex mirror is unstable. Similarly, if the mirrors are too concave, the cavity is again unstable. The stability can be determined from the radii of curvature and the separation of the mirrors. A cavity is stable if the following inequality is satisfied [19]:

$$0 \leqslant \left(1 + \frac{L}{ROC_1}\right)\left(1 + \frac{L}{ROC_2}\right) \leqslant 1 \tag{2.13}$$

where ROC_1 and ROC_2 are the radii of curvature of the two mirrors and L is the length of the cavity (or the separation of the mirrors). Note that concave mirrors have negative radii of curvature.

The modes of a stable cavity are very similar the Hermite-Gauss modes. Since the behaviour of the Hermite-Gauss modes is known it is possible to calculate the resonant mode structure of a cavity with given length and mirror radii of curvature using the Hermite-Gauss modes as a model. The Rayleigh length of the cavity is given by:

$$z_0{}^2 = \frac{-L(ROC_1 + L)(ROC_2 + L)(ROC_2 + ROC_1 + L)}{(ROC_2 + ROC_1 + 2L)^2} \tag{2.14}$$

which gives the cavity waist through equation 2.2;

$$w_0 \equiv \sqrt{\frac{\lambda z_0}{\pi}}$$

Also, the waist position inside the cavity is found using:

$$\begin{aligned} z_1 &= \frac{-L(ROC_2 + L)}{(ROC_2 + ROC_1 + 2L)^2} \\ z_2 &= z_1 + L \end{aligned} \tag{2.15}$$

where z_1 and z_2 are the positions of the mirrors relative to the waist at z=0.

The equations 2.14 and 2.15 define the mode matching for the cavity. Incident light is coupled efficiently into the cavity mode if it has the same waist size and position as the cavity mode. If the incident light is not in this mode, the 'mismatched' components are reflected off the cavity. To couple light in efficiently, lenses are placed in the beam path to give the desired waist. Mode matching is an essential

part of experimental cavity physics.

Note that even though all the Hermite-Gauss modes have the same radius of curvature, they will not in general experience the same round trip phase delay in a cavity. The Gouy phase shift (equation 2.4) is weighted by the mode number (l+m+1) in the equation for the Hermite-Gaussian beam (equation 2.7). Thus, the different spatial modes receive a different longitudinal phase shift in the cavity, causing them to be, in general, non-degenerate. It is possible for a higher order spatial mode to be degenerate with the fundamental TEM_{00} mode if its round trip Gouy phase shift *difference* is exactly 2π. Cavities are generally designed so that this does not occur for any mode with significant amplitude.

2.3.2 The Cavity Equations

Although it is conceptually easier to understand the summation approach to cavity fields, it is mathematically simpler to analyse a set of self-consistent steady state equations [17]. These equations are derived in an intuitive way from the fields shown in figure 2.5. Note, the circulating field is defined at the point infinitesimally inside the front mirror, headed in the same direction as the incident field. If the amplitude reflectivities of the mirrors are r_1 and r_2 for the front and back mirrors respectively and assuming the system is lossless the cavity relations are:

$$E_{ref} = r_1 E_{inc} + i r_2 t_1 E_{circ} e^{i\phi} \quad (2.16)$$

$$E_{circ} = i t_1 E_{in} + r_2 r_1 E_{circ} e^{i\phi} \quad (2.17)$$

$$E_{trans} = i t_2 E_{circ} e^{i\phi/2} \quad (2.18)$$

where ϕ is the round trip phase shift of the cavity defined by $\phi = 2\omega L/c$ and, as before, the convention of a $\pi/2$ phase shift on transmission is used. The cavity relations are then solved in terms of the input field, E_{inc} to give the cavity equations:

$$E_{ref} = E_{inc} \frac{r_1 - r_2 e^{i\phi}}{1 - r_1 r_2 e^{i\phi}} \quad (2.19)$$

$$E_{circ} = E_{inc} \frac{i t_1}{1 - r_1 r_2 e^{i\phi}} \quad (2.20)$$

$$E_{trans} = E_{inc} \frac{-t_1 t_2 e^{i\phi/2}}{1 - r_1 r_2 e^{i\phi}} \quad (2.21)$$

These three equations define the three fields of importance when studying a cavity. The power and phase of these three fields is shown in figure 2.6, for three different pairs of mirror reflectivites. Mirror reflectivites are usually quoted in terms of power, usually as percentages. These are converted to amplitude reflectivities by taking the square root of the fractional power reflectivity, i.e. $r = \sqrt{R}$. The phase of the circulating field does not provide any useful information and is only included for completeness. As an addendum to the above equations, the maximum circulating power per input power, or the power recycling factor, is given at $\phi = 0$. For the specific case where there is no loss in either the mirrors or the cavity medium, the

power recycling factor is given by:

$$\frac{P_{circ,max}}{P_{inc}} = \frac{it_1}{1 - r_1 r_2} \tag{2.22}$$

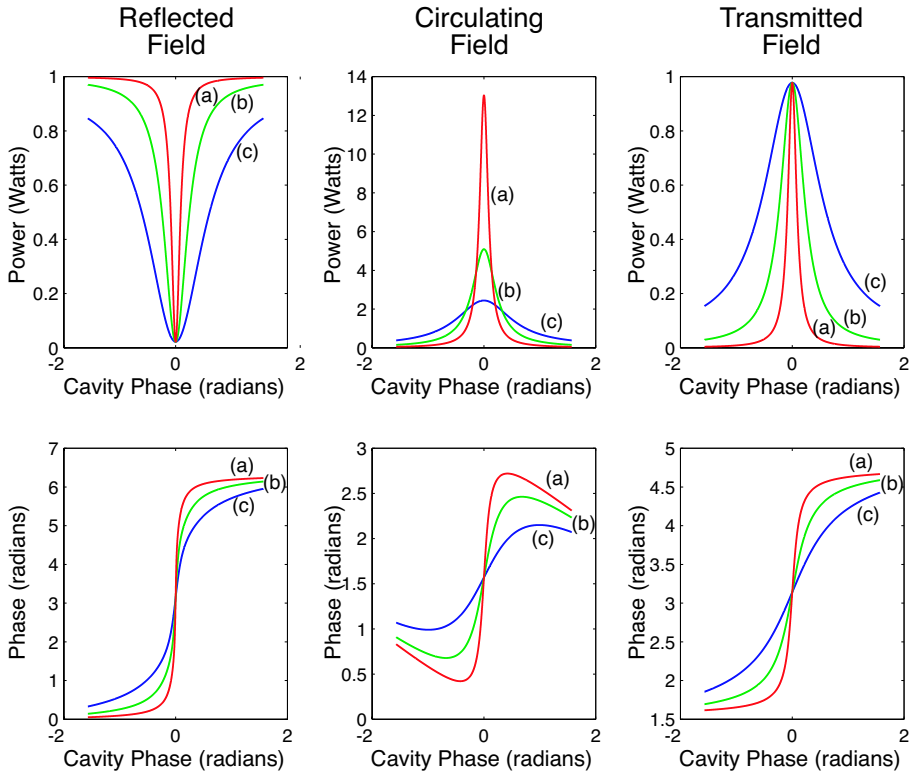

Figure 2.6: Plots of the Reflected, Circulating, and Transmitted powers, with their respective phases, for the parameters (a) R1=50% R2=60% (b) R1=75% R2=80.1% (c) R1=90% R2=92.5%. Input power is unity

The height of the circulating peak, and the 'sharpness' of the resonance is measured by a value called finesse ($\mathcal{F}$). The finesse is a common term, and is defined in several different ways. However, the different values converge for $\mathcal{F} \gtrsim 10$. The first definition, based on mirror amplitude reflectivities, is:

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} \tag{2.23}$$

The second definition is based on cavity loss, both in the mirrors and in the intra-cavity medium. It is not stated here. The third definition is based on the relative width of the transmitted peak compared with the inter-peak spacing. The inter-peak spacing is called the free spectral range (FSR), and the peak width is defined as the full width at half maximum height (FWHM). The FWHM is also known as the linewidth (ν) or bandwidth of the cavity. The finesse is related to these values

by:

$$\mathcal{F} = \frac{\text{FSR}}{\text{FWHM}} \tag{2.24}$$

where the FSR is;

$$\text{FSR} = \frac{c}{2L} \tag{2.25}$$

The FSR is the distance between adjacent longitudinal modes in laser frequency. However, it is equally valid to express the FSR in terms of the cavity length. Since each longitudinal mode is an integer number of wavelengths in round trip length, an FSR is one wavelength in length. Similarly the FWHM is $\lambda/\mathcal{F}$ metres 'long' in cavity round trip length. The calibration from wavelength to finesse to FWHM will be used later in this thesis. Further discussion of cavities is contained in chapter 3, in particular, control of cavities, and their interaction with phase and frequency modulation.

Chapter 3

Control

This chapter outlines the theory behind two common control techniques in optics. The first is DC locking, which is a more general form of offset locking. The second is RF cavity locking, which includes the commonly used PDH (Pound Drever Hall) locking [21]. Before these topics are covered, a brief introduction to the essentials of control is given.

3.1 Control of Optics

Control systems for optics consist of four main parts: the plant, sensor, servo, and actuator. The plant is the chief component and the core of the experiment. The sensor diagnoses the condition of the plant. In optics the sensor usually includes a photodetector. The servo stage filters the output of the sensor and amplifies or attenuates it as necessary before feeding it to the actuator. The actuator controls the plant.

For the purposes of this thesis, the plant is considered to be either an interferometer or a cavity. In general, a cavity will be locked to resonance, and an interferometer locked to an arbitrary position on a fringe. To lock either of these systems, the optical path length must be precisely controlled.

The sensor stage of the control system may have several components. Optically, the most important is the photodetector. However, for the case of RF locking, the sensor also includes a modulator and a mixer.

The servo stage consists of two (or more) components. The two core components are the servo and the high voltage amplifier. The servo filters and amplifies the input signal from the sensor. Generally this involves the output signal being suitable for the actuation stage, which will have limited actuation bandwidth. The servo must have low noise. The second stage is a high voltage amplifier. It provides flat gain and its output may vary from -200 to +200 volts. Both servos and high voltage amplifiers are designed and built in-house. An additional stage could be a pre-amplifier, to boost the size of the signal from the sensor before it reaches the servo.

The actuator adjusts the length of a cavity, or the arm length difference of a Michelson interferometer, by moving a mirror attached to a piezo-electric transducer (PZT). PZT's are ceramic devices that expand or contract under input voltage. The PZT's used in this experiment accept input voltages from -100 to 200 volts, which allows them to move a little more than one micron. Typically, PZT-mirror systems

are able to respond to inputs up to 20 kHz, though in practice the mounting system often has some lower frequency resonances prohibiting, actuation at frequencies much higher than 1 kHz.

For a control system to be successful the sensor must provide an error signal. The error signal must provide both the polarity and magnitude of the deviation of the plant from its desired operating point. The error signal must be zero crossing, so that at the desired condition the plant is unaffected by the control system. The derivation of an appropriate error signal is a major component in the design of an optical control system. The following sections discuss two techniques for deriving an error signal for particular plant and sensor combinations. The two techniques described are DC locking and RF locking, for the specific cases of a Michelson interferometer and a Fabry-Perot cavity respectively.

3.2 DC Locking a Michelson Interferometer

DC locking is a remarkably simple and effective means of controlling the resonance condition of a Michelson interferometer. The Michelson interferometer is the plant and a photodetector on the output, which measures a small fraction of the transmitted beam, is the sensor. The servo stage is a combination of a filtering gain stage, a PZT-PID servo, with a flat gain stage on the output, a high voltage amplifier. The actuation comes from PZTs on one or both of the arm mirrors of the interferometer. Critical to DC locking is a DC voltage offset source. The offset is subtracted from the output of the sensor to create an error signal.

For a high fringe visibility interferometer, the transmitted power of the Michelson varies sinusoidally and is close to zero at the minimum. After photo detection, this corresponds to a sinusoidally varying output voltage which, like the output power, does not cross zero. If a constant voltage of the appropriate size is subtracted from the photo voltage, the total will cross zero at some points. The zero crossing points are the lock points of what is now an error signal. Varying the size of the offset changes the lock points. By switching the polarity of the feedback to the actuator, it is possible to lock to either the positive or negative voltage slope, and thus to either side of an interference fringe.

DC locking uses the output of the Michelson as an error signal. Since the output has zero derivative at the maximum and minimum, it is impossible to lock these points using DC locking. Also, since the slope of the output power is less near the turning points, the gain of the servo stage must be higher to keep the interferometer locked. Locking close to the end-points is a special case of DC locking called offset locking. Offset locking is currently planned for use in the second generation of LIGO detectors [9].

DC locking is useful for a great many applications apart from Michelson interferometers. It is simple, cheap, robust, and versatile. However, the inability to use it to lock to turning points in power make its application unsuitable at times.

3.3 RF Cavity Locking

Radio frequency cavity locking schemes are varied in concept and application. Universally they rely on imposing an RF modulation on the laser beam. After interacting with the cavity, the laser light is detected on a high speed photodetector, with a detection bandwidth greater than the RF modulation frequency. For brevity, only RF locking schemes using frontal phase modulation (PM) are described here.

In frontal PM schemes, the incident light is phase modulated at some radio frequency (typically 10-500 MHz). The modulated light interacts with the cavity and is detected either on transmission or reflection. The output signal from the detector is then demodulated (or mixed down) to DC and low pass filtered to provide an error signal. The elements here are discussed in order; phase modulation, interaction with a cavity, detection, and demodulation and filtering.

3.3.1 Phase Modulation

Phase modulation changes the phase of the electric field. This is usually achieved by changing the optical path length using an electro-optic modulator (EOM), which changes the refractive index of a crystal by applying a voltage across it. The output light is phase modulated in a manner proportional to the input voltage on the modulator.

In the phasor picture, by rotating the frame of reference at the optical frequency, the unmodulated beam has length E_0 and phase ϕ. The phase modulation leaves the length unchanged, but the angle becomes time dependent. The zero to peak modulation, δ, can be expressed in terms of the maximum phase change from the average position, usually in radians. For small modulation depths, the 'wobbling' phasor can be decomposed into an unmodulated phasor (the carrier), with two time varying phasors (the sidebands). This decomposition is shown graphically in figure 3.1. The two sidebands vector sum with the carrier to change its phase without changing its amplitude. Mathematically, phase modulation is represented by:

$$E = E_0 \mathrm{e}^{i(\omega t + \delta \sin(\omega_m t))} \tag{3.1}$$

where the unmodulated beam is $E_o \mathrm{e}^{i\omega t}$ and ω_m is the angular frequency of the modulation. If $\delta \ll 1$, this can be decomposed to make:

$$E_0 \mathrm{e}^{i\omega t}(1 + i\delta \sin(\omega_m t)) = E_0 \mathrm{e}^{i\omega t}\left(1 + \frac{\delta}{2}\mathrm{e}^{i\omega_m t} - \frac{\delta}{2}\mathrm{e}^{-i\omega_m t}\right) \tag{3.2}$$

In figure 3.1 the sidebands cancel out, but when time is advanced one quarter of the modulation period, the sidebands will add constructively in the imaginary plane (since they counter-rotate relative to the carrier). When the imaginary plane sidebands are then added to the carrier, the vector sum is of length E_0 (for small δ), and is tilted with respect to the vertical. It is possible to express amplitude modulation in the same picture by simply flipping the sign of one of the sidebands. Thus, AM sidebands cancel in the imaginary plane and add in the real plane, changing

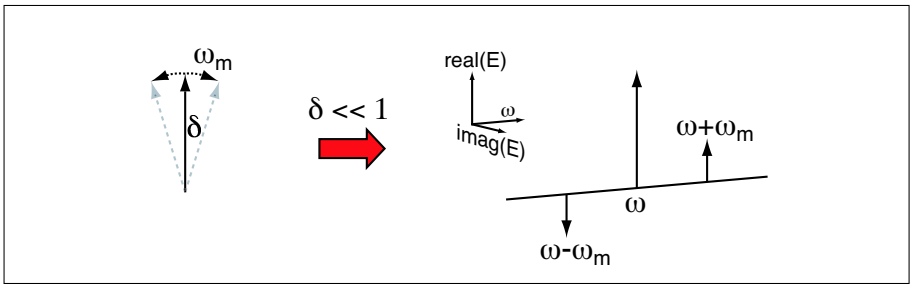

Figure 3.1: In the first picture the phasor is length E_0 oscillates in phase at frequency ω_m. For small modulation depth, this is decomposed into a stationary carrier with two superimposed sidebands, which vector sum to give the same phase modulation at frequency ω_m.

the length but not the angle of the carrier. Amplitude modulation is represented mathematically by:

$$E_0 e^{i\omega t}(1 + \delta \sin(\omega_m t)) = E_0 e^{i\omega t}\left(1 - i\frac{\delta}{2}e^{i\omega_m t} + i\frac{\delta}{2}e^{-i\omega_m t}\right) \tag{3.3}$$

The sideband picture of modulation will be used in the following sections to understand the interaction, detection, and derivation of an error signal after interaction with a cavity.

3.3.2 Interaction with a Cavity

The resonance condition of a cavity is dependent on the input optical frequency. The sidebands are at a frequency different from the cavity, and as such will experience a different resonance condition. If the modulation frequency is high compared with the bandwidth of the cavity, then when the carrier is on resonance the sidebands will be far from resonance. If the modulation frequency is low compared with the cavity bandwidth, then the sidebands will experience approximately the same resonance condition as the carrier. These two sideband interaction regimes are analysed separately.

In the case of high modulation frequency, when the carrier is on resonance the sidebands are non-resonant. Since the sidebands are non-resonant, they will be reflected off the cavity. It is therefore assumed that the reflected field is to be monitored. The description for the transmitted field is similar, but the sidebands are highly attenuated. If the carrier then drifts a small way from resonance, it will obtain a phase shift, $\Delta\phi$, from the dispersion of the cavity. However, the sidebands are far from resonance, and so experience effectively zero dispersion. Thus, the carrier experiences a phase shift relative to the sidebands. The reference frame is then rotated to keep the carrier phase constant, and it now appears that the sidebands experienced the phase shift. The rotated sidebands are then decomposed into the original PM, and some AM. The interaction with the cavity converted some

PM to AM. This decomposition is shown in figure 3.2.

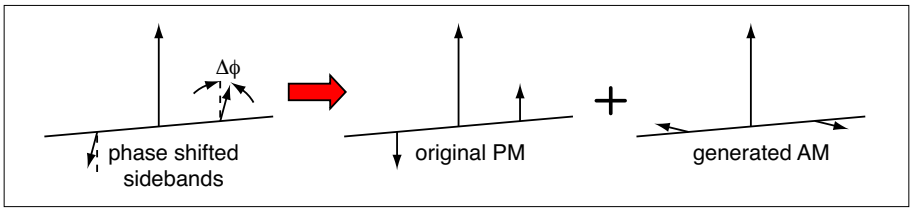

Figure 3.2: The effect of a phase shift of the carrier with respect to the sidebands. Some PM is converted to AM. Figure reproduced from reference [22]

The electric field sidebands following this interaction are described by [22]:

$$E = E_0 e^{i\omega t}\left(1 + i\delta\sin(\omega_m t)\cos(\Delta\phi) - \delta\sin(\omega_m t)\cos(\Delta\phi)\right) \tag{3.4}$$

Here the first term is the carrier, the second is the incident phase modulation attenuated by the $\sin(\Delta\phi)$ term, and the third term is the generated amplitude modulation. Equation 3.4 describes the interaction of high frequency phase modulation after interaction with a Fabry-Perot cavity.

When the modulation frequency is low compared with the cavity linewidth, the effect on the sidebands is quite different. In this case, when the cavity is on resonance, the sidebands are very near to resonance, and so are transmitted along with the carrier. In this case, the transmitted field will be analysed, though the arguments are valid for both transmitted and reflected fields. The sidebands receive a different phase shift to the carrier on both transmission and reflection, but this does not serve to generate AM, since one sideband is phase advanced and the other equally delayed. It has the same effect as advancing the sideband picture in figure 3.1 in time slightly. As such, this phase effect will be ignored.

When the carrier drifts from resonance, it and the sidebands receive approximately the same phase shift, since the phase dispersion of the cavity is linear well inside the linewidth. As the carrier drifts, one sideband will approach resonance, and the other will be driven further from resonance. This changes the relative heights of the sidebands on either transmission or reflection. The differential height change of the sidebands converts the incident PM into AM, as shown in figure 3.3.

Note that the AM sidebands in figure 3.3 have different phase relative to the PM sidebands than those in figure 3.2. This change in relative phase is important, and is briefly mentioned in section 3.3.3. Mathematically, the conversion of PM to AM through differential magnitude changes is expressed as:

$$E = E_0 e^{i\omega t}\left(1 + i\alpha\sin(\omega_m t) - \alpha\cos(\omega_m t)\right) \tag{3.5}$$

Where α is the change in height of the sidebands, and is dependent on the transmissivity of the cavity at the sideband frequencies.

In the regime between high and low frequency modulation, a linear combination of these two effects will occur. Since the phase of the AM is different in the

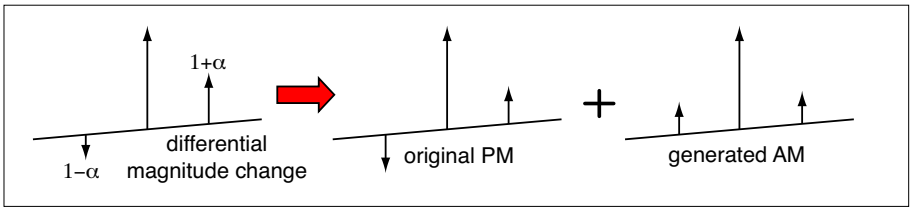

Figure 3.3: PM is converted to AM when the sidebands experience differential magnitude changes. Figure reproduced from reference [22]

two cases, the phase of the output AM depends on the relative strength of the two effects. In all cases, the amount of AM generated is determined by the difference from resonance. The next section explains how the conversion from phase to amplitude modulation allows a standard detector to sense the modulation, and hence the resonance condition of the cavity.

3.3.3 Detection, Mixing, and Filtering

Photodiodes convert light into current. They operate at a given quantum efficiency, which specifies what fraction of incident photons are converted to electrons in the current. The quantum efficiency is typically around 90% for detectors operating at $\lambda = 1064$ nm. The photo-current passes through a trans-impedance amplifier, which produces an output voltage proportional to the input current. Although the gain and efficiency of different detectors is unique, the output voltage is still proportional to the input optical power. Equations which specify photo-voltage in terms of incident electric field will thus have an implicit constant of proportionality.

Here the optical power is defined by the relation $P = E^*E$, where E* is the conjugate of E. Since it is assumed that the voltage from the detector is directly proportional to the optical power, the photo-voltage is given by:

$$V_{pd} = E^*E \tag{3.6}$$

where the constant of proportionality (with units of volts per Watt) is left out. This equation shows that phase modulated light, defined in equation 3.1, has constant magnitude upon detection.

When a phase modulated beam interacts with a cavity near resonance, amplitude modulation is created. This amplitude modulation is detectable. Converting equations 3.2 and 3.3 into voltages, the detected voltage after interaction with a cavity in the high frequency regime is:

$$V_{pd} \approx P_0(1 - 2\delta \sin(\omega_m t) \sin(\Delta\phi)) \tag{3.7}$$

and in the low frequency regime;

$$V_{pd} \approx P_0(1 + 2\alpha\delta\cos(\omega_m t)) \tag{3.8}$$

In both cases the δ^2 terms are neglected. P_0 is the average power of the detected light. Both $\Delta\phi$ and α give information about the resonance condition of the cavity. To derive an error signal, the information encoded at the modulation frequency is reduced to DC by demodulation. Demodulation, or mixing, multiplies the signal by a sinusoidal reference at the modulation frequency.

The phase of the demodulation source relative to the generated amplitude modulation of the beam is important. The best way to get the information from the modulation frequency is to demodulate in phase with this signal. In equation 3.7 this means multiplying by a $\sin(\omega_m t)$ local oscillator, and for equation 3.8 multiplying by a $\cos(\omega_m t)$ local oscillator. The phase of the local oscillator is measured with respect to the incident phase modulation. In practice, the local oscillator both generates the PM and provides the demodulation source.

When the photo-voltage is demodulated there are terms at DC, ω_m, and $2\omega_m$. These terms are low pass filtered to strip away the high frequency terms. The low pass filter leaves a window around DC large enough so the resulting term can provide information about the changes in the resonance condition up to a MHz or so. The result of demodulating and low pass filtering equations 3.7 and 3.8 with the appropriate local oscillator phase is:

$$V_{demodulated} = -P_0\delta\sin(\Delta\phi) \tag{3.9}$$

for high frequency modulation. For low frequency modulation:

$$V_{demodulated} = P_0\delta\alpha \tag{3.10}$$

Around resonance neither the transmitted nor reflected power varies to first order, so both power and modulation depth are constants. Thus the error signal is proportional only to $\sin(\Delta\phi)$ in the high frequency case and α for low frequencies. Both these terms are zero on resonance and have opposite sign on either side of resonance. Thus an error signal which provides information on the distance and direction from resonance is generated by phase modulation, detection, and demodulation with the correct phase. The particular case of locking a cavity using high frequency modulation, and measuring the reflected power, is known as Pound-Drever-Hall (PDH) locking. It is a commonly used technique in interferometer control, and is an important sub-class of wider RF locking techniques.

The error signal generated in this way is proportional to the imaginary component of the transmitted (or reflected) field of the cavity. Figure 3.4 shows the imaginary component of the cavity plotted about resonance. Notice the linearity of the trace close to resonance. The slope of the error signal in this central section is twice the peak-to-peak height divided by the peak-to-peak separation [23].

In the intermediate regime the error signal is still linear and proportional to

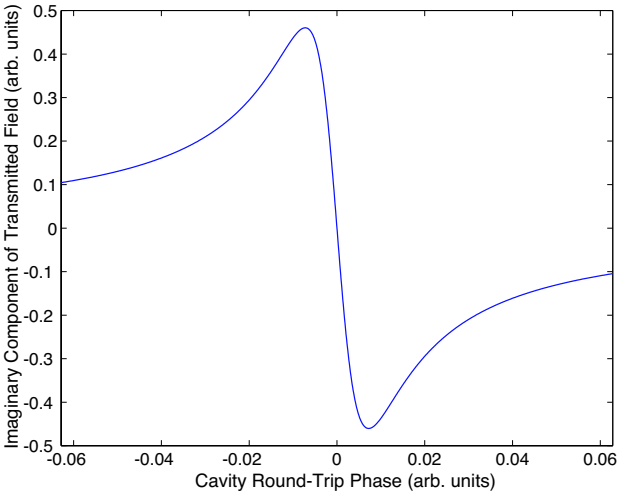

Figure 3.4: The imaginary field of the cavity on transmission has the same shape as the error signal around resonance. It is linear for small motions, and crosses zero at resonance.

length (or frequency) changes near resonance. The required demodulation phase is non-trivial, though this is no practical hindrance since the phase between the local oscillator and the amplitude modulation signal can be altered by delaying one signal.

Chapter 4

Flexure Design

The momentum of light is notoriously small. To measure radiation pressure, with the modest laser power available in a lab environment, it is necessary to have a very light mirror. This chapter addresses the design of the moveable mirror. Also, the interactions between radiation pressure, an optical cavity, and a moveable mirror are modelled. The model is used as an aid to understanding the system and as an independent verification that the effects observed are indeed caused by radiation pressure.

4.1 Simple Harmonic Oscillators

The aim of this section is to present all the formulae used in modelling the interaction of the moveable mirror with light pressure. To do so, it is necessary to describe the basic design of the moveable mirror.

The term 'moveable mirror' is perhaps misleading. More accurate is 'mirror on a spring'. The mirror is mounted on a thin cantilever of niobium, which is in turn attached to a solid base. The mirror is glued on, but the rest of the construct is monolithic.

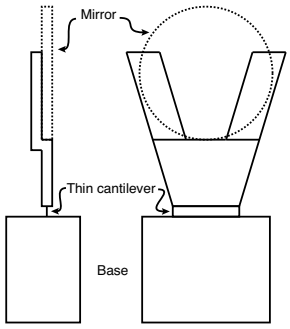

Figure 4.1: The thin cantilever shown is of order one hundred microns thick. The mirror is 1mm thick, as is the plate on which it is mounted. Barring the mirror, the system is monolithic.

A force exerted on the mirror stresses the cantilever, or flexure membrane, causing it to bend. The simple cantilever design is advantageous because it exhibits

single moded behaviour with a spring constant easily modified by changing the cantilever geometry. The bending of the flexure membrane is an angular action, and as such the flexure system has an angular spring constant. The equations governing the motion of simple harmonic oscillators are expressed in angular terms to match the flexure motion. The angular equations are equivalent to the linear equations and the two can be interchanged by substituting all the appropriate rotational terms with their linear analogues. A list of the relevant terms and their analogues is contained in reference [24].

The key parameters for simple harmonic oscillators are the resonance frequency (ω_0), the quality factor (Q), the angular spring constant (k_a), and the moment of inertia (I). The resonance frequency is related to the moment of inertia and spring constant by [25]:

$$\omega_0 = \sqrt{\frac{k_a}{I}} \tag{4.1}$$

The quality factor is used to calculate the angular damping co-efficient (b_a):

$$b_a = \frac{I\omega_0}{Q} \tag{4.2}$$

with these terms defined, the steady state angular response, in radians, of the oscillator is given by:

$$\theta = \frac{T_0}{\sqrt{I^2({\omega_0}^2 - \omega^2) + {b_a}^2\omega^2}} \tag{4.3}$$

Where T_0 is the amplitude of the torque incident on the cavity at frequency ω, such that the time dependent torque is $T = T_0 \sin(\omega t)$. The phase of the steady state motion relative to the input is given by:

$$\tan\phi = \frac{b_a\omega}{I({\omega_0}^2 - \omega^2)} \tag{4.4}$$

With equations 4.3 and 4.4 the transfer function of a mirror-spring system can be plotted once the angular spring constant, the quality factor, and the moment of inertia are known. These parameters are dependent on both materials and geometry. The flexure material is niobium. Niobium was chosen because it has a high intrinsic quality factor, is relatively cheap, and is easy to work with. Also, the Australian gravitational wave research community has significant expertise in designing high-Q niobium oscillators. See, among others, references [26, 27].

The relevant material properties are the Young's modulus and density of niobium [28]:

$$\begin{aligned} E_y &= 105 \times 10^9 \quad \text{Pa} \\ \rho &= 8570 \quad \text{kgm}^{-3} \end{aligned} \tag{4.5}$$

The density was used in calculations of the moment of inertia of the flexure-mirror system. The Young's modulus (or modulus of elasticity) was used to determine the

angular spring constant [29, 30]:

$$k_a = \frac{E_y I_a}{l} \tag{4.6}$$

where l is the length (or height) of the flexure membrane and I_a is the moment of area of the flexure membrane. The moment of area is defined by the thickness and width of the membrane:

$$I_a = \frac{W t^3}{12} \tag{4.7}$$

It is important that W $\gtrsim$ 30t. If this condition is satisfied, then the flexure motion is largely single moded. Since the membranes are typically of order 100 microns thick, the flexure must be at least a few millimetres wide. It is implicitly assumed that the base is heavy compared with the 'moving' components. The simple equations presented so far were used to model radiation pressure effects in the pre-project stage. The work completed in the pre-project design phase, including calculations of radiation pressure effects under realistic conditions, is summarised in appendix A.

4.2 Radiation Pressure Response

To determine if the effects seen in this experiment were radiation pressure it was important to model the strength of the radiation pressure effects independent of fitted parameters. The model was created using a mix of theoretical and experimental parameters. This section describes the key aspects of the model. Some of the results derived from the flexure response model are presented in chapter 6. Although the model discussed below could have been used for any of the three flexures, it was only used for the flexure with the highest resonance frequency and the 240 μm thick membrane.

The important values for calculating the cavity length changes caused by radiation pressure are: moment of inertia (I), Angular spring constant (k_a), quality factor (Q), incident laser power (P_{inc}), amplitude modulation depth (δ), circulating power (P_{circ}), and the height of the mirror above the rotation axis (h_m).

The moment of inertia was determined using the software package Pro-Engineer. Prior to construction, a 3-D model of the flexure was created, as shown in figure 4.2. After adding the mirror to the 3-D model and supplying Pro-Engineer with appropriate values for density (of Niobium, and of BK7 for the mirror), the moment of inertia was determined. There was some substantial ($\sim$10%) error in the moment of inertia because it was not obvious where to place the rotation axis, since the bend of the flexure membrane is continuous.

The angular spring constant was calculated using equation 4.6. The bulk geometric values of the flexure membrane, the length and the width, were assumed to be the design values. The thickness was adjusted from the design parameter until the resonance frequency given by the model was equal to that observed in the experiment.

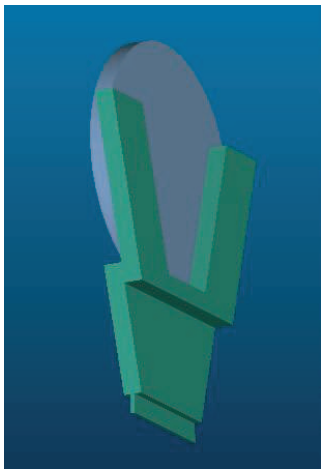

Figure 4.2: The flexure (or flexure-mirror system). The base is not shown in this image, but it attaches to the bottom of the flexure membrane.

Incident laser power was measured immediately in front of the cavity using a power meter. The modulation depth was defined as the zero to peak fractional amplitude of the modulation. It was a variable, and could be set to a particular experimental value to check experimental results. The circulating power was determined using both experimental and design values. First, the circulating power was calculated using equation 2.20 and the design values for the mirror reflectivities. This was later refined by reducing the reflectivity of the mirrors to match the experimentally measured cavity finesse. The circulating power converts to optical force through photon momentum, giving:

$$F_{opt}(t) = F_0 \sin(\omega t)$$

with

$$F_0 = \frac{2\delta P_{circ}}{c} \tag{4.8}$$

The factor of two is present in the radiation pressure force because the back mirror of the cavity reflects most of the circulating power, thus doubling the momentum kick from each photon.

To convert an optical force to a torque, and to convert an angular displacement to a cavity length change, the height of the mirror above the rotation axis had to be known specifically, the height of the laser spot above the rotation axis. The spot was assumed to be in the middle of the mirror, and the position of the mirror taken from the design parameters of the flexure system. The quality factor was taken directly from the experiment, as discussed in chapter 6.

The complete matlab code used to model the interaction of the flexure-mirror system with radiation pressure is included as appendix B.

Figure 4.3 shows a sample plot from the refined model. The vertical scale shows

the position fluctuations of the moveable mirror as a function of driving frequency. For this plot, P_{inc} = 250mW, δ = 1, and Q = 1000. Note that there were not enough data points in the plot to accurately resolve the resonance peak.

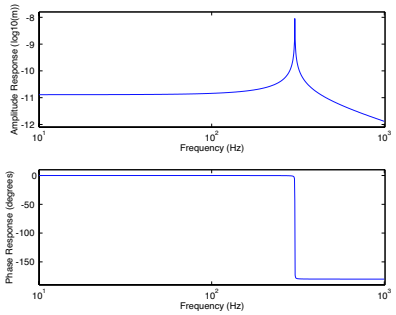

Figure 4.3: A sample of the calibrated transfer functions generated by the refined model for radiation pressure effects.

Calibrated transfer function plots were crucial in the geometric design of the flexure system, as presented in appendix A. The amplitude response provide the most important information and the phase response is only included here for completeness. They also provide key support for the argument that the amplitude modulation effects discussed in chapter 6 are from radiation pressure.

4.3 The Optical Spring

Prior to the commencement of the experimental part of this project, the optical spring effect was not considered. The interaction of the mechanical spring with the radiation pressure force was examined more closely in the context of experimental questions.

When free from other effects, the torque exerted on the mirror is linearly proportional to its angular displacement. For small displacements, this is linearly proportional to the horizontal displacement of the mirror. The slope of this torque versus angle relation is the angular spring constant. However, when the mirror is part of an optical cavity, as the horizontal position of the mirror is changed, so is the resonance condition of the cavity. This creates a changing force on the mirror which is the same shape as the circulating power plots shown in figure 2.6. After converting this force into a torque, the sum of the torques acting on the mirror is no longer linear with distance.

The changing slope of the torque versus distance line means the effective spring constant of the 'mirror' is changing with displacement, and as such the resonance frequency also changes. Figure 4.4 shows the predictions of the model for input power of 250 mW. The resonance frequency changes are predicted to be large, approximately 10% of the mechanical resonance.

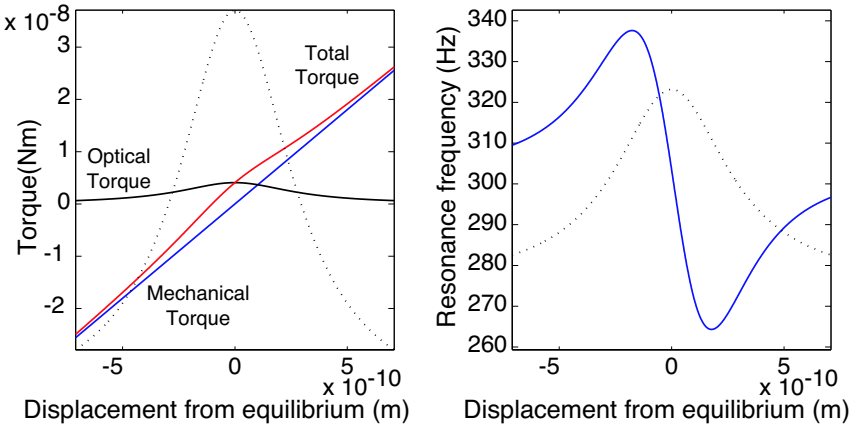

Figure 4.4: The left plot shows the mechanical and optical torques, and the sum of the two. The right plot shows the resonance frequency changing from its mechanical value of 303 Hz. On both plots the dotted line shows the circulating power (arb units) as a length scale reference. Cavity resonance is assumed to be at the equilibrium position of the flexure.

The study of these optical spring effects, and other possible interactions between radiation pressure and the flexure, presented exciting new experimental possibilities. Although there was not time enough to explore these new interactions to much depth, some data were recorded. This data is analysed and compared with the radiation pressure model in chapter 6.

4.4 Chapter Summary

The simple equations governing the motion of damped simple harmonic oscillators presented have been presented. They were used to plot the amplitude response of a flexure-mirror system for a given optical force input. These plots allowed the development of a suitable oscillator geometry. The model was refined after the commencement of experimentation. The refinements corrected numerical values and allowed for the modelling of unexpected new effects, principally the optical spring effect. The results chapter presents the comparisons between the refined model and experimental results, whereas this chapter has provided the necessary tools for such a comparison.

Chapter 5

Experimental design

The design of this experiment can be grouped into several different categories. The movable mirror was discussed in the previous chapter. This chapter deals with environmental noise suppression, optical layout and components, and the control scheme.

5.1 Environmental Noise Suppression

Any mechanical system capable of sensing radiation pressure will also be susceptible to a great deal of other noise. This section deals with the noise requirements of the experiment in general and, in particular, with the techniques needed to ensure that non-radiation pressure motion is minimised.

There were four main sources of noise likely to limit the sensitivity of the experiment: acoustic, electronic, thermal, and vibrational. The gravitational wave research group has had significant experience with very high sensitivity interferometry, and has developed low noise photodetectors and good control electronics (servo's and high voltage amplifiers). For this reason electronic noise was assumed to not be a limiting noise source, and this assumption was shown to be justified. Also, in terms of absolute sensitivity the experiment will not be approaching thermal noise, so we can discount this from our noise concerns.

It should be noted that the laboratory environment was unusually quiet to begin with. The concrete slab for the laboratory is twice as thick as a normal foundation, and it is separated by several centimetres from both the walls and the foundations of the rest of the building.

Vibrational and acoustic noise were highlighted as the two main noise sources. To combat these two sources several measures were taken. The first stage of isolation came from the optical bench. The table is a TMC optical bench floated on active gas legs. The legs were filled with nitrogen at approximately 900 kPa. As such, there was no solid contact between the table top and the ground.

Secondly, the test cavity in the experiment (and as such the flexure) is mounted on a custom made 12 kg base. This base was then placed on 4 sorbothane hemispheres. Sorbothane is a visco-elastic damping material, which has a critically damped resonance [31]. It is ideal for isolating the flexure from table vibrations, so long as it is properly loaded.

Thirdly, for acoustic isolation a box to enclose the test cavity was built. It was

constructed from one inch thick medium density fibreboard, which was subsequently lined with inch thick closed cell foam. The fibreboard, being relatively dense, was selected to provide the bulk of the acoustic absorption. The foam was added to increase damping of internal modes. To assist the internal damping, the foam was added in such a way that no two sides were parallel, the corners were padded, and other asymmetric features were created.

Picture 5.1 shows all the essential features. The large stainless steel base holding the test cavity is not in contact with the table. There is a thick layer of sorbothane on the table which marks the perimeter of the box. This serves several practical purposes, firstly cables are clamped in the sorbothane, reducing the vibrational 'short-circuiting' effect of the cables. Secondly, there is no air gap between the bottom of the box and the table, increasing acoustic insolation. Thirdly, the box is always in the same place and orientation. The final thing to note is the acoustic suppression chamber, which clearly shows the thickness of the wood and foam, some of the asymmetric features, and the input window.

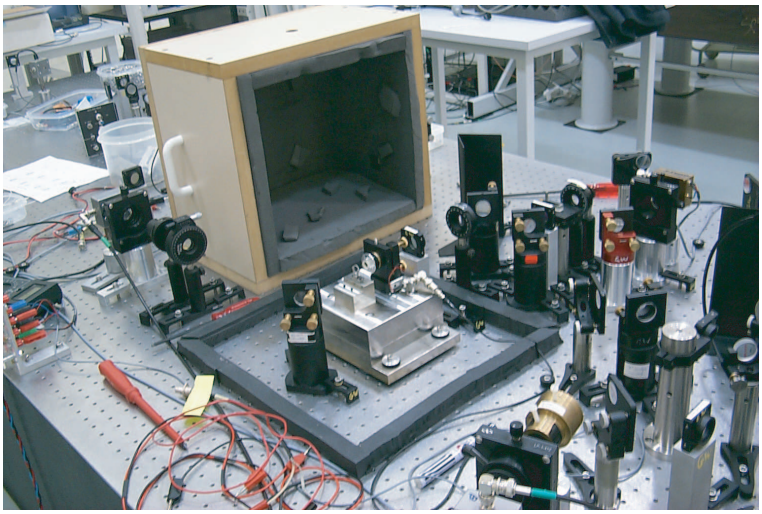

Figure 5.1: The experiment, showing the underside of the noise suppression chamber, and the section of the experiment it covers. Note the 12 Kg slab holding the optical cavity is held off the table by a damping and isolation material.

With all these measures in place the environmental noise was greatly attenuated. Quantitative measurements of the reduction were not made since it was impossible to operate the experiment without them in place. It should also be noted that almost all data were recorded at night, since ambient daytime noise made it virtually impossible to operate the experiment during the day.

5.2 Optical Layout

The optical layout describes the optical and opto-electronic components of the experiment. The two most important features here are the amplitude modulator and

the test cavity. The design parameters of both components, and the experimental performance of the amplitude modulator are discussed. The performance of the test cavity is analysed in chapter 6. Other important features mentioned are the laser (used for frequency modulation), the photodetectors, and the power management and calibration.

5.2.1 The Amplitude Modulator

The generation of amplitude modulation was a crucial part of the experiment. There were several requirements:

- •Full modulation depth, from zero to full power
- •Large acoustic frequency range, from 1 Hz to 1 kHz
- •No phase modulation effect
- •Highly stable output

To achieve these goals a purpose built modulator was used: a common mirror Michelson interferometer.

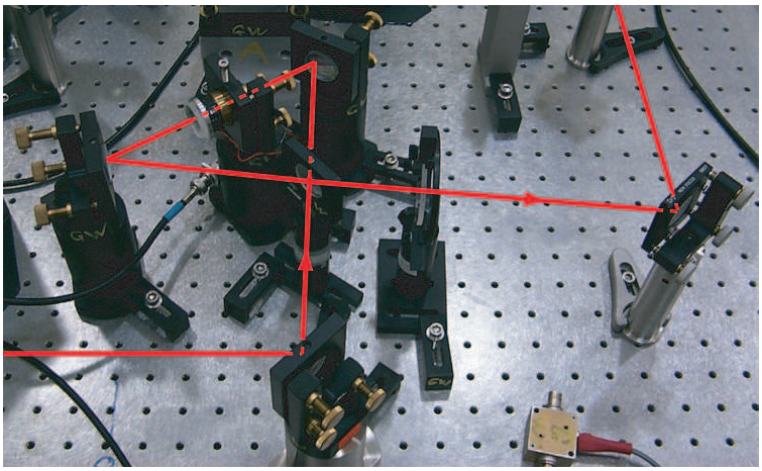

Figure 5.2: The common arm Michelson used as our amplitude modulator. Note the common mirror, mounted on a PZT.

A standard Michelson interferometer allows full control of output intensity from 0-100%. To modulate the output it is necessary to adjust the length of the arms relative to each other, which can be achieved by modulating the length of one arm. However, this introduces equal parts phase modulation and amplitude modulation, as shown in figure 2.4. We could modulate the length of each arm exactly out of phase and with equal depth, but this is experimentally challenging since it requires that the PZT's in each arm have matched responses, an undesirable complication.

The common mirror Michelson 'folds' the beams in each arm so they reflect off the same mirror, at the same angle, as shown in picture 5.2. Modulation of the position of this middle mirror changes only the relative length, ΔL. The average length, L, is unchanged, and as such the amplitude, but not the phase is modulated.

5.2.2 The Laser and Frequency Modulation

The laser used in this experiment is a lightwave LW 126. It produces 700 mW of output power in a slightly elliptical spatial mode with an elliptical polarisation. It is a non-planar ring oscillator (NPRO), and has good long term frequency and amplitude stability, after about half an hour of operation. It has a ND:YAG crystal, and has an output wavelength of 1064 nm. ND:YAG lasers are the standard lasers used in gravitational wave research.

A frequency modulator is necessary to complement the amplitude modulator. The laser controller has two front panel co-axial cable connections which control the laser frequency, labelled fast and slow. The slow input alters the temperature of the lasing medium, giving a large dynamic range of 10 GHz. The response time, however, limits the usable bandwidth to ~ 100 mHz. The fast input controls a PZT, which stresses the laser crystal, altering the output over a range of 150 MHz. The bandwidth of the PZT response is ~ 150 kHz.

The response of the laser frequency to voltage at the fast input is linear over a large range of input amplitudes and frequencies. Over a broad frequency span, and a variety of modulation depths, the output modulation depth remains constant for a given input voltage. For this reason the laser was chosen as the frequency modulator.

It is important to note the difference between phase and frequency modulation at low frequencies. At frequencies well inside the cavity bandwidth phase modulation has little effect. In the time domain, this is because the phase of the circulating field is 'cycled' out of the cavity before the input phase has changed significantly. Thus, low frequency phase modulation does not change the resonance condition of a cavity significantly. Frequency modulation, on the other hand, has exactly the same effect as length modulation. At low frequencies, frequency modulation simply drives the cavity off resonance. Phase and frequency modulation have directly comparable effects, it is only the magnitude of these effects that is affected by the modulation frequency.

5.2.3 The Test Cavity

The test cavity formed the core of this experiment. The cavity has two major effects. Firstly, it increases the sensitivity to the position of its constituent mirrors by a factor proportional to finesse and secondly, it has a circulating optical field which is larger than the input, again by a factor proportional to finesse. However, it also adds complications. To conduct the experiment, the cavity has to be on resonance, which requires some form of control. There are many other issues, some of which may be dealt with using appropriate design parameters.

Having a higher finesse gives both higher sensitivity and larger circulating power, both desirable. The trade-off again comes in control. It was desirable to be in the regime where measurements could be made outside the locking bandwidth of the cavity control system. This, in turn, required that any noise outside that same bandwidth be insufficient to drive the cavity far from resonance. Thus it was neces-

sary to choose a value of finesse which gave both reasonable circulating power and noise requirements.

Since the flexure is an angular spring, it tilts the reflected beam as it changes the length of the cavity. Since, for a given angle tilt, the vertical offset scales linearly with length, we chose a relatively short cavity. The cavity was designed to be linear because it is possible to have the input and reflected beam use the same port, suitable for situating the cavity inside a box.

With the geometry and finesse constrained, the radii of curvature of the mirrors had to be decided. The movable mirror was very thin, and as such difficult to polish with precision. A flat was selected to keep the manufacture simple. The front mirror was designed with a radius of curvature appropriate to a cavity of the given length and finesse. The design parameters are shown in table 5.1.

Parameter	Design value
Reflectivity of fixed mirror	99 %
Reflectivity of light mirror	99.7 %
ROC of fixed mirror	0.10 m
ROC of light mirror	inf
Cavity length	15 mm

Table 5.1: Design parameters for the test cavity

These design values were used to infer the cavities other properties, shown in table 5.2. The values were calculated using the cavity formulae, equations 2.19 to 2.21.

Parameter	Calculated value
Finesse	481
Free Spectral Range (FSR)	10 GHz
FWHM (or linewidth)	21 MHz
Power recycling factor	236
Spot size at waist	110 μm

Table 5.2: Design parameters for the test cavity

The linewidth of the cavity described above is important because it justifies the assumption that this cavity is always in steady-state. Changes in the cavity circulating field are effectively instantaneous when compared with the acoustic frequencies used in the amplitude and frequency modulators. Whilst large bandwidth is good for this reason, it is important that it is of the same order as the laser PZT modulation depth, so that large deviations from resonance can be generated using laser frequency modulation.

The experimentally measured values for finesse and length are presented in chapter 6. The FSR, linewidth, and power recycling factor were inferred from these values. The spot size in the cavity was not measured, but the calculated value was used for mode matching into the cavity. Following the cavity design came its incorporation into the optical layout of the whole experiment.

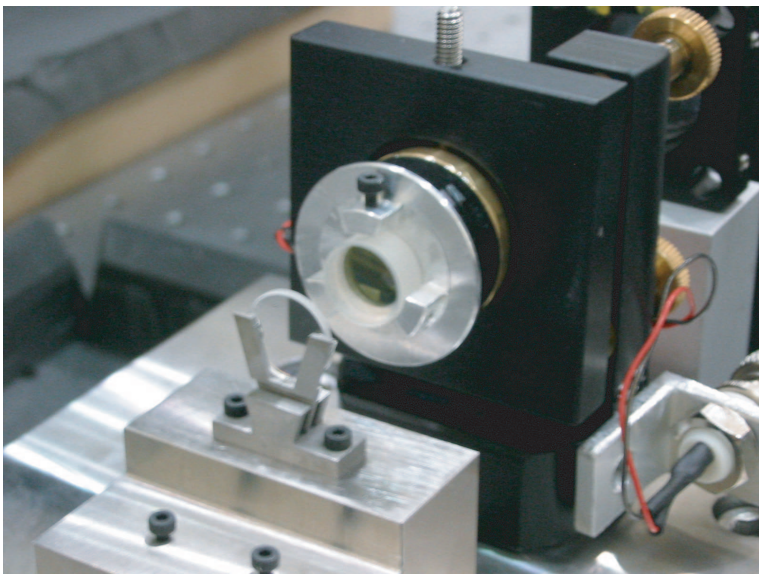

Figure 5.3: The test cavity and moveable mirror.

5.2.4 Experimental Schematic

This section is a description of the optical layout represented schematically as figure 5.4. The discussion will follow the beam path as closely as possible, from the laser to the final detector.

The input beam preparation consists of all the optics in the first row. The isolator, necessary to stop optical feedback into the laser, only transmits input light of the correct polarisation. Since high power levels were needed, it was important to maximise this transmission. As such, the elliptical polarisation coming out of the laser had first to be linearised by the quarter wave plate ($\lambda/4$), and then rotated to the correct axis by the half wave plate ($\lambda/2$). In general, the optics are optimised for vertical polarisation, so the polarising beamsplitter (PBS) ensured that only vertical polarisation was reflected into the rest of the experiment. In combination with the $\lambda/2$ plate after the isolator the downstream power can be adjusted to an arbitrary level. This was most useful since it is often unwieldy to work with 700 mW of power. The beam was then collimated by lens 1 (L1), which has the long focal length f=750 mm.

The next two optics are steering mirrors, which were designed to reflect the beam with low loss at a broad range of incidence angles. They align the beam into the beamsplitter of the amplitude modulator. It should be noted that to align the Michelson, the middle mirror was left out, and the system was aligned as a Sagnac interferometer [32]. This simplified alignment vastly, since the interferometer was especially sensitive to misalignment of the middle mirror.

On the output of the amplitude modulator, lens 2 (f=1000 mm) re-collimates the beam, which was then steered into the broadband phase modulator, Gsänger PM 25 intra-cavity phase modulator. Lenses 3 and 4 both have f=200 mm. They focus the

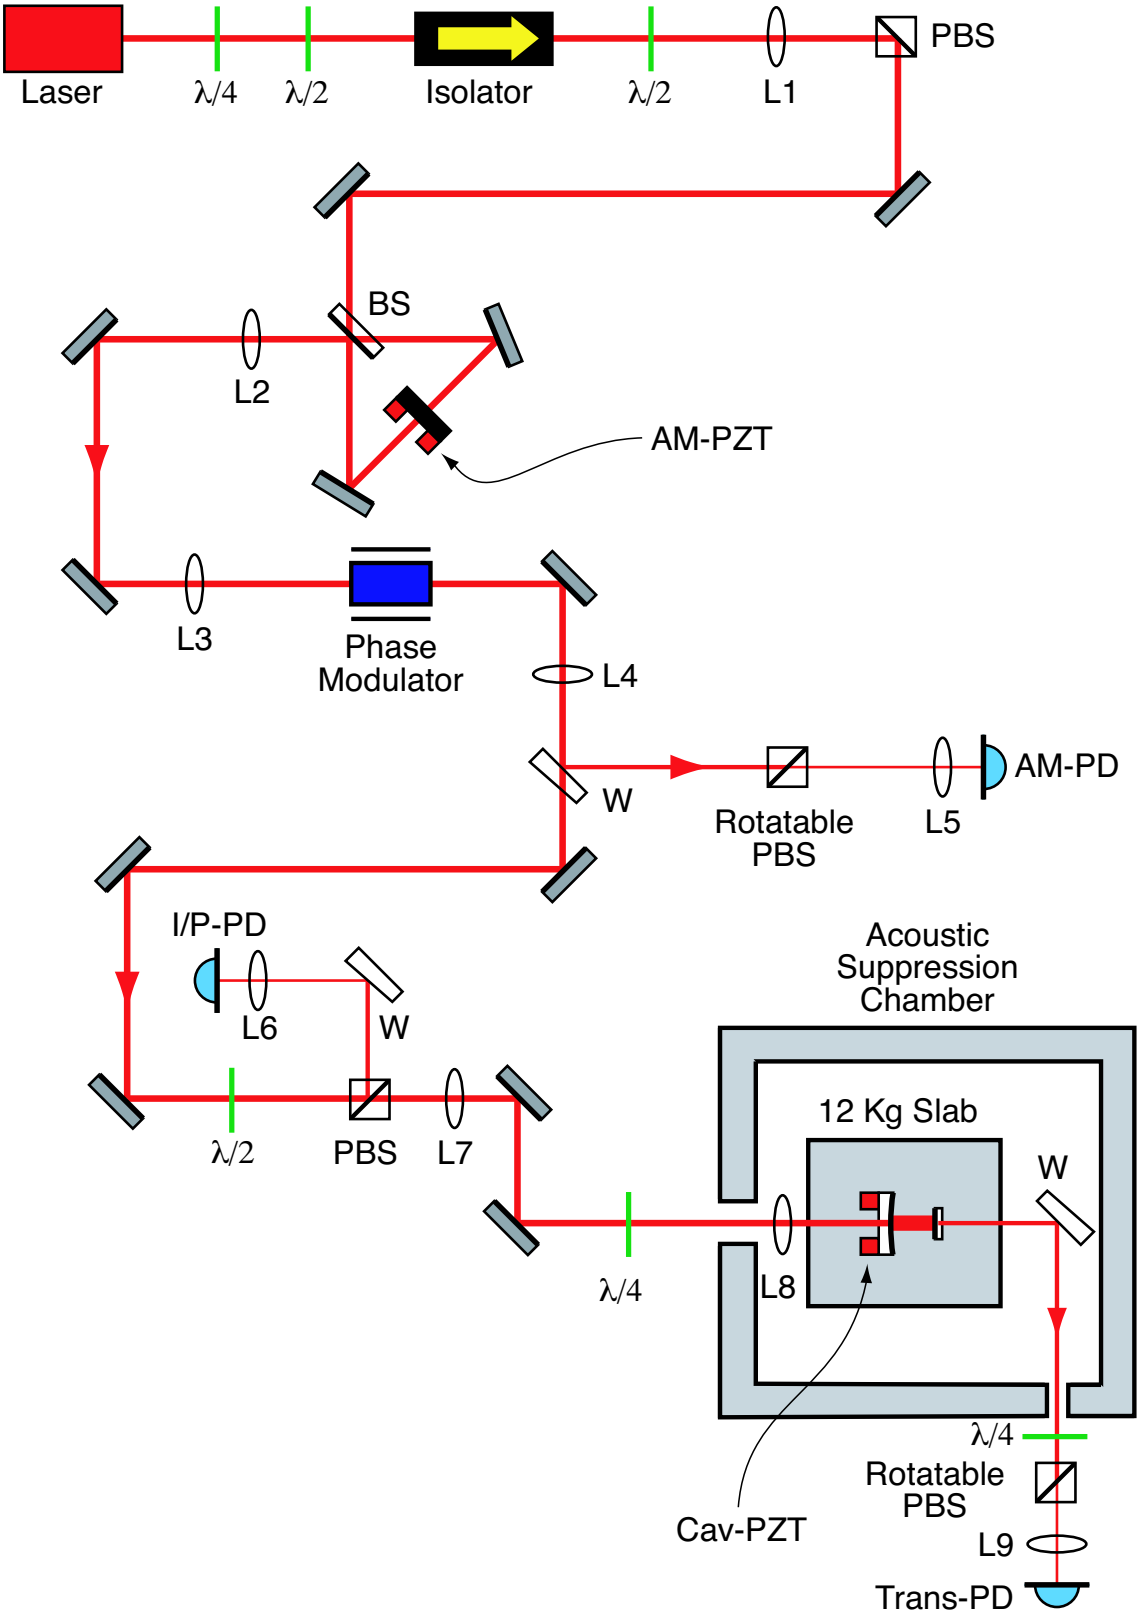

Figure 5.4: A full diagram of the optical components of the experiment.

beam to a small size to pass through the phase modulator, and then re-collimate it on the other side. The lenses were positioned 2×f apart, so that the waist size of the beam incident on the first lens is the same as the waist size of the beam after the second lens. The waist position was offset, but this was corrected for by adding optical delay. It was important not to change the waist size so that mode matching and collimation properties were constant. The wedge which follows (W) is an uncoated substrate. As such, it simply reflects a small percentage (typically around 5%) of the incident light off each surface. The front and back surfaces are angled at 3 degrees relative to each other, to eliminate any parasitic etalon effects, and provide two spatially separated beams at the output.

One of the outputs from the wedge was attenuated by a rotatable PBS. By rotating the PBS the transmitted beam goes from being vertical (with respect to the table) to horizontal and back. Since the input beam is purely vertical, it can be attenuated to an arbitrary level (limited only by the extinction ratio of the PBS, typically 1000:1). Lens 5 (f=25 mm) then focused the beam onto a low noise, DC-coupled transimpedance photodetector [33]. This photodetector was used to determine how much light is transmitted through the amplitude modulator. It was a crucial part of the AM locking scheme, as discussed below in section 5.3.1. From here on this detector will be referred to as AM-PD (for amplitude modulation photodetector). There is a piezo-electric transducer (PZT) attached to the common mirror of the Michelson. This was used for control of the amplitude modulator.

The beam was then steered around to another $\lambda/2$ and PBS arrangement. These two alter the optical power incident on the cavity, again limited only by the extinction ratio of the PBS. The reflected port of the PBS was directed onto a wedge (for constant attenuation), through lens 6 (f=25 mm), and onto the detector labelled I/P-PD (for input photodetector), This detector records how much of the power incident on the PBS is transmitted onto the cavity, as a percentage.

On the transmitted port of this PBS the beam passes through two more lenses, 7 (f=1000 mm) and 8 (f=200 mm), both used for mode matching. The quarter wave plate before the cavity converts the polarisation to circular. Once light reflects off the cavity, the direction of the circular polarisation is reversed, from left-handed to right-handed or vice-versa. When it then passes through the quarter wave plate again it becomes vertically polarised (orthogonal to the incident horizontal polarisation). Thus it is possible to analyse the reflection from the cavity at the final unused port of previous PBS. However, for this experiment the reflected port was not used, but the quarter wave plate was never removed since it provided extra isolation against optical feedback into the laser.

The two steering mirrors before the cavity were crucial for aligning the input beam into the cavity mode. The input beam had to be re-aligned often, since the base holding the test cavity would move as the sorbothane sagged. The test cavity was discussed above, though now the PZT on the front mirror of the cavity is explicitly drawn and labelled. This PZT provided the actuation on the cavity length. The wedge on the transmitted port of the cavity served two purposes: firstly it attenuated the beam to hit the detector, and secondly it steered the beam out of a second hole in the acoustic suppression chamber. The quarter wave plate

and rotatable PBS outside the box again provided variable attenuation, before the beam was focussed (lens 9, f=25 mm) onto the transmitted port detector, Trans-PD.

There were two reasons for using the transmitted port. Firstly, during initial alignment of the cavity it was far easier to analyse the transmitted port, since the cavity acts as a mode-cleaner. Secondly, high levels of power were incident on the cavity. Even when it was on resonance only a small fraction of the incident light was transmitted onto the Trans-PD, making the power much easier to manage. The Trans-PD was used for cavity control and readout. It provided the bulk of the important data for this experiment.

5.3 Control and readout scheme

Control forms a crucial part of most interferometry experiments. The aim is to keep the system operating at the ideal point by suppressing noise, but at the same time to keep the control scheme from interfering with the desired measurement bandwidth. There were no measurements to be taken with the Michelson interferometer, so the ideal situation was for very large locking bandwidth with large noise suppression. For the cavity, on the other hand, the error signal (or control signal) was used as the readout of the cavity resonance condition - the critical data. Consequently it was desirable to have a very low locking bandwidth, so the error signal was not disturbed by the feedback signal at the frequencies of interest. The locking for the two schemes is discussed at greater length below.

5.3.1 Michelson Interferometer Control Scheme

To use a Michelson as an amplitude modulator it is crucial to control its interference condition. Since the Michelson was used in the larger experiment, it was also important that the control scheme be simple, robust, and independent. To this end DC locking was used as outlined in section 3.2.

DC locking involves taking a small tap-off of the power transmitted through the Michelson. This tap-off created a voltage on the amplitude modulation photodetector (AM-PD). This voltage has an approximate range from -20 mV to -5 V. To lock the Michelson, an electronic offset between +20mV and +5V was added to the above signal, and the resulting signal was put through a servo and fed back to the PZT shown on the common mirror in the Michelson (see picture 5.2).

DC locking is perhaps the simplest scheme to lock a Michelson with, and although it is difficult to lock very close to either a dark fringe or a bright fringe (since the error signal is now symmetric), it provided a locking range more than sufficient for this experiment.

The question of how to inject the modulation signal was also answered neatly. Figure 5.5 clearly shows the extra ports that can used for signal injection, it also illustrates the remarkable simplicity of the scheme, both conceptually and electronically. By adding a sinusoidal signal to the electronic offset used for locking, the lock point changes to track the combined input. Thus, at any frequency where the

closed loop locking gain was high, any injected modulation signal was tracked very closely by the output.

The DC locking scheme locks the output power level, rather than the position on the interference fringe. This has two effects. First, any input power fluctuations were suppressed, and second any input modulation (or signal) was tracked by output power, and not the PZR position. Thus even at high modulation depths the output power is unaffected by the sinusoidal nature of the Michelson curve.

To test the response of the modulator, its transfer function was measured from input into the servo (changing the lock point) to output at an independent, downstream photodetector. The transfer function, see figure 5.6, has several features. There are two resonances, one at 257 Hz, the other at approximately 1 kHz. These are resonances in the mechanical mount holding the common mirror. Neither presented a particular problem, as they are both quite narrow in frequency and small in amplitude. There was also a flat phase roll-off of approximately 39 degrees from 1 Hz to 1.2 kHz. This corresponded to a time delay of approximately 90 μs.

The flatness of the response of the amplitude modulator was more than satisfactory. Excluding the resonances, the modulation depth change by approximately 0.5 dB from 1 Hz to 800 Hz. The long term stability of both the locking and modulation were excellent. It was rare to have the modulator 'drop lock' during a night of experimentation, and lock acquisition (or re-acquisition) took very little effort.

5.3.2 Cavity Control and Readout Scheme

The control of the test cavity was both the most complex and important part of the control scheme. The locking scheme used was transmission locking, a subset of RF cavity locking. Phase modulation sidebands are added at a frequency generally at or below the linewidth of the cavity. In this case 19 MHz. The sidebands were then partially transmitted through the cavity, as discussed in section 3.3. The error signal then came from two effects: differential attenuation of the sidebands and nonlinearity of the cavity phase response. The optimal demodulation phase was then

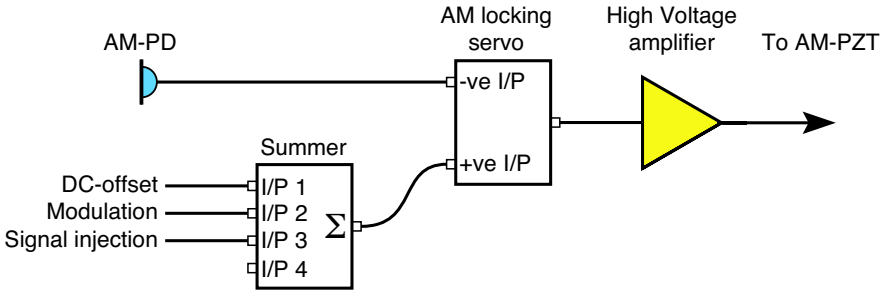

Figure 5.5: A schematic of the locking scheme used to control the amplitude modulator. It has enough ports for a DC offset (for locking), a modulation signal, and any other signal used to analyse the response.

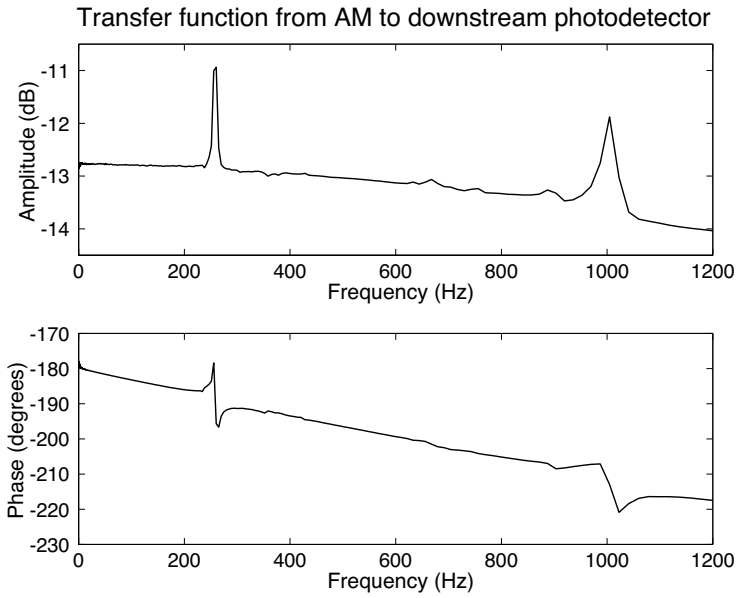

Figure 5.6: The amplitude and phase components of the transfer function from input at the amplitude modulator, to output at an independent, downstream photodetector. Note the resonances at 257 Hz and 1 kHz, and the phase roll-off corresponding to a delay of 90 μs.

determined empirically.

The modulation frequency was also chosen empirically. Several different frequencies were used, but this was the one that seemed to have the best error signal. Once chosen the modulation frequency was not altered throughout the experiment.

For the electronic gain stages there were several requirements. There had to be a broad range of gain values, with fine control. Also, it was necessary to have several points at which to either inject or monitor signals. To achieve this both a Stanford SR 560 pre-amplifier and the cavity servo were used for variable gain effects, and then a standard high voltage amplifier to drive the PZT. The pre-amp had a spare input and output port, and it was largely used for both result taking and locking loop analysis. The servo had one spare input and output - both of which were largely unused. The HV amp had one spare input, which was not used. For a schematic illustration of the locking scheme see figure 5.7. The servo itself was significantly modified before the principal data taking. These modifications are discussed briefly in chapter 6 and the open loop transfer function of the modified servo is included in appendix C.

There are two limits suitable for recording data in a control loop. The first is in the high gain limit. In this case, all noise (and signal) is suppressed by the locking system, so the error signal is quiet. However, the correction point signal (at the output of the servo) must be tracking the other mirror precisely, so it contains the necessary cavity length information.

However, the test cavity demonstrated poor locking performance in the high gain

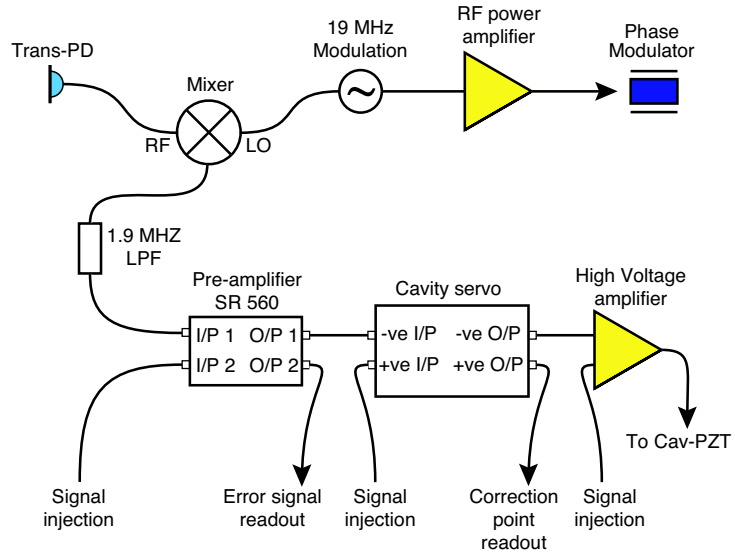

Figure 5.7: The schematic for the locking scheme used to control and analyse the cavity length

limit, due to instabilities when feeding back above the mechanical resonance. Thus the other good recording regime was used - very low gain. In this case, only noise at low frequencies was suppressed, and the signal frequencies did not interact with the servo. This in turn required that the noise at the signal frequencies was insufficient to cause the cavity to drop lock.

Thus, when the test cavity was locked with sufficiently low bandwidth, the error signal diagnosed the resonance condition of the cavity. In the case where the flexure mirror motion was the dominating noise source, the spectrum of the error signal provided the spectrum of this motion. Since the front mirror is essentially rigid, and the laser has stable frequency, the low noise assumption is justified in the absence of signal injection. Thus the error signal was used to diagnose the position of the flexure mirror only.

5.4 Chapter Summary

This chapter has presented in detail the design features of this experiment. The environmental noise suppression techniques necessary for a measurement this sensitive have been discussed. The optical layout necessary to fully prepare, analyse, and inject the laser into the test cavity has been presented and explained. The control scheme necessary for the Michelson interferometer and the test cavity has been discussed. It was shown that the error signal from the test cavity would provide the best measurement of the cavity resonance condition - and as such the cavity length.

Chapter 6

Results

The most important data recorded during this project are presented in the following sections. The results fall into three categories, cancellation data, both single frequency and broadband; the observation of non-linear interactions; and the verification that the cancellation and non-linearity were caused by radiation pressure.

6.1 Non-linear Effects

The first non-linear effects were recorded while scanning the length of the test cavity across the cavity resonance. The power transmitted through the cavity was monitored using the trans-PD connected to a HP 54602B 150 MHz CRO. When the cavity was scanned across resonance, by continuously increasing or decreasing the length using the cavity PZT, a peak was seen on the CRO corresponding to the resonance of the cavity. If the cavity was a linear system free from power dependent effects (such as radiation pressure), then the resonance peak would have been symmetric and independent of the direction of the length change, for any input power. However, a clear power dependence was observed in the shape of the resonance peak.

The cavity was first scanned with low incident power of approximately 50 mW. Symmetric, identical peaks were seen on the increasing- and decreasing-length scans. The cavity was then scanned with high power, an increase by a factor of approximately 8. The peaks are asymmetric and dissimilar. The power incident on the cavity was changed by varying the output power from the Michelson. The traces of both these scans are shown in figure 6.1.

The traces in figure 6.1 require explanation case by case. The two cases are; approaching the resonance by decreasing the cavity length, and approaching resonance by increasing the cavity length.

As the cavity length decreases towards resonance, the circulating power increases, which exerts a radiation pressure force on the moveable back mirror of the cavity. The cavity then increases in length, moving further from resonance. Through this process the circulating power rises more slowly than if there was no moveable mirror. When the cavity reaches resonance, the radiation pressure force is at a maximum and the mirror is at its furthest extreme from equilibrium. When the length of the cavity then decreases a little beyond resonance, the force drops and the cavity length decreases further from resonance. This runaway effect causes the transmitted power to drop very quickly. Thus, when the cavity decreases in length towards resonance

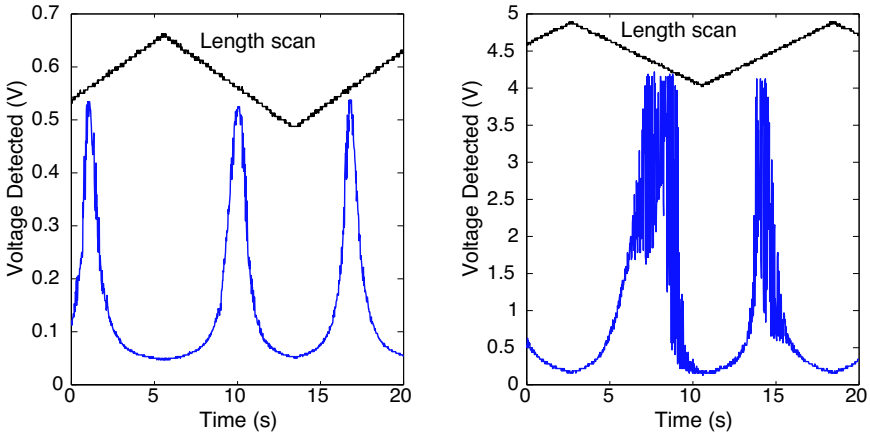

Figure 6.1: Scans of the cavity resonance at low and high power. The high power trace shows significantly non-linear behaviour. The top trace in each plot is proportional to the cavity length.

the transmitted power will increase slowly, reach its full height, and drop suddenly. This is seen in the first peak of the high power trace in figure 6.1.

When the cavity length increases towards resonance, the circulating power increases, which increases the radiation pressure force on the back mirror. The cavity then increases in length, and comes closer to resonance, increasing circulating power further. This is a runaway effect and the transmitted power has a very short rise time. The cavity is just shorter than resonance length prior the sharp rise in circulating power begins. The higher circulating power pushes the back mirror rapidly through resonance and the cavity continues to lengthen. As the cavity length increases further, the force on the back mirror lessens and the cavity shortens, keeping it closer to resonance than it otherwise would have been. Thus, when the cavity increases in length over resonance, a narrow peak with reduced height, sharp rise time, and longer decay time is expected. These features are seen on the second peak of the high power scan in figure 6.1.

The fuzz seen on both peaks in the high power trace was due to ringing of the flexure at its mechanical resonance frequency. The ringing was a dynamic effect, and is not covered by the theoretical treatment in chapter 4. It occurred when the circulating power in the cavity changed rapidly, exciting the mechanical resonance at 303 Hz. The ringing decays eventually to give a steady state response. Modelling and studying these dynamic effects is beyond the scope of this project.

Figure 6.2 resolves the ringing from figure 6.1. The ringing frequency is 287 Hz, not the expected frequency of 303 Hz. This was the first experimental evidence for the optical spring effect discussed in section 4.3 and this measurement inspired further investigation of the optical spring effect.

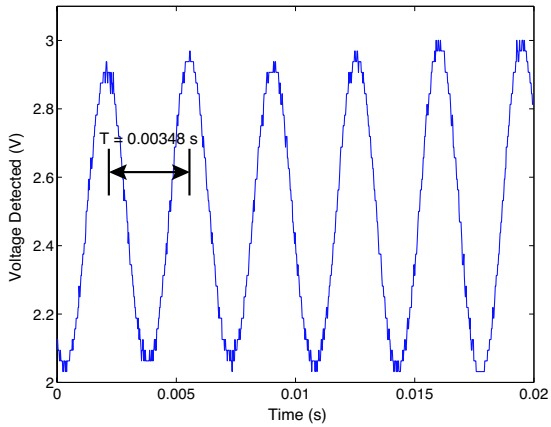

Figure 6.2: A close-up of the 'fuzz' on the high power cavity scan. The time between adjacent peaks is approximately T=0.00348s, corresponding to a frequency of 287 Hz.

6.2 Cavity Locking and Frequency Shifting

It was a major experimental challenge to lock the test cavity in a stable manner, using the locking scheme shown in figure 5.7. The first difficulty was in choosing an appropriate locking bandwidth. The cavity front mirror was mounted on a PZT and the mount holding the PZT had a number of acoustic frequency resonances; in particular one at 900 Hz. This resonance made locking with high bandwidth all but impossible. Locking with the unity gain point inside the measurement bandwidth, between 10 Hz and 1 kHz, was possible but highly undesirable. The error signal was monitored to measure the cavity length, if the gain had changed significantly in the signal bandwidth then the data would be difficult to interpret.

Thus, it was decided to lock with low bandwidth. The cavity servo was modified to make locking with low bandwidth simpler. The open loop gain of the cavity was reduced, the corner frequency of the output low pass filters was reduced, and the pre-integrator offset was made less sensitive. The circuit diagram of the servo (pre-modification) is shown in appendix C, along with the post modification open loop transfer function. After these changes the cavity was successfully locked with low bandwidth. However, having low bandwidth meant that noise at low frequency, and in particular at the resonance frequency, was unsuppressed by the locking loop. Despite the measures taken to suppress environmental noise, as discussed in section 5.1, it was not possible to lock the cavity in a stable manner during the day due to this low frequency noise. Most data were recorded at night to allow stable locking, and to improve signal to noise.

With the modified servo, the cavity was consistently locked with a bandwidth of less than 1 Hz. Figure 6.3 shows a typical closed loop transfer function with a locking bandwidth of approximately 900 mHz. This form of closed loop transfer function is also known as the sensitivity function, since it measures how much closed loop suppression there is at each frequency. To obtain this closed loop response, a swept sine signal was injected into the spare port of the SR 560 pre-amplifier, and

the output was monitored at the spare output of the same pre-amplifier. The signal was amplified by 20 dB before it was analysed, so when the closed loop gain was low, the output was amplified by 20 dB. The magnitude of the output decreased from 20 dB as the closed loop gain increased. When the transfer function is 3 dB below the unsuppressed value, the closed loop gain is unity. Thus, it was possible to measure the unity gain frequency (which is equivalent to the locking bandwidth) from the closed loop sensitivity function.

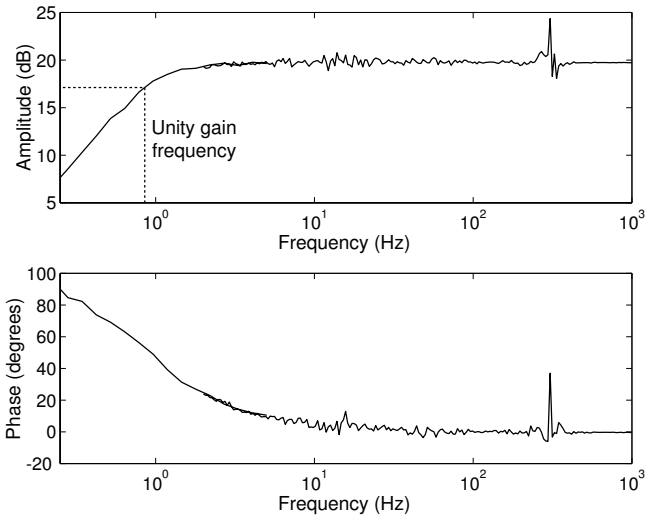

Figure 6.3: The sensitivity function of the cavity locking loop. The unsupressed value was at 20 dB because the injected signal passes through a flat gain stage before measurement. Unity gain frequency was approximately 900 mHz.

When the cavity was locked with this low bandwidth the error signal was monitored on the output of the SR 560 pre-amplifier. A SR 785 dual channel digital dynamic signal analyser was used to monitor the error signal. The SR 785 has a maximum bandwidth of 100 kHz and a minimum linewidth (or resolution bandwidth) of approximately 0.5 mHz. It has a source output capable of generating swept sine signals, for recording transfer functions, and bandwidth limited white noise, used in the broadband cancellation section. It was used to record the transfer function in figure 6.3. A sample of the spectrum of the error signal (excited by background noise) is shown in figure 6.4. The ordinate usually has units dBVrms, which are defined by $dBVrms = 20\log_{10}(V_{rms})$. Since the excitation caused by ambient noise was modified by the transfer functions of the noise suppression stages, it was not considered to be 'white', and as such the background noise spectrum is not the same as the mechanical transfer function of the flexure.

There were unavoidable electronic offsets in the cavity locking loop. These offsets caused the cavity to be locked slightly off resonance, by shifting error signal so the zero crossing point was no longer exactly at resonance. Locking off resonance has the effect of changing the effective spring constant, and hence resonance frequency, as shown in figure 4.4.

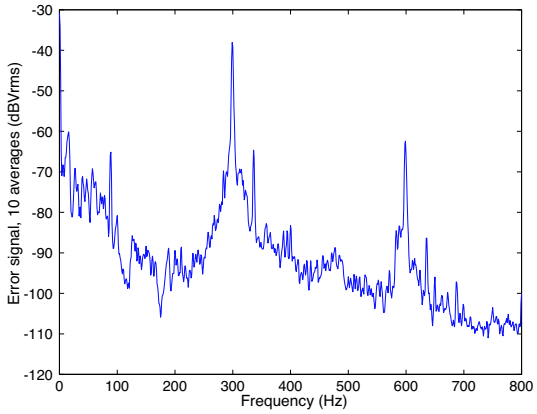

Figure 6.4: The spectrum of the cavity locking error signal excited by ambient noise. The flexure resonance is at approximately 303 Hz. Linewidth= 1Hz.

During observations of frequency shifting it was noted that to lock the cavity with good stability and low bandwidth, the effective frequency had to be lower than the unmodified mechanical resonance frequency. The reason for this optical stability effect was not understood completely, and is a topic for further work. To consistently lock the cavity with the same bandwidth, the resonance frequency was modified by injecting a DC offset into the spare input of the cavity servo. This DC offset was adjusted until the resonance frequency was approximately 301 Hz.

It was important to measure the mechanical resonance frequency of the flexure, free from optical effects. This was done by reducing the input power to its minimum level. It was noted that at this low power, the resonance frequency did not change as a function of cavity servo offset. From the low power background noise trace, the mechanical resonance frequency was then measured to be $f_{res} = 303$ Hz.

To measure the strength of the frequency shift, a trace was recorded with the standard locking offset, giving a resonance frequency of 302 Hz. The DC offset was then reduced until the cavity locking was at the limit of its stability. At this point the resonance frequency was 293.5 Hz. When the DC offset was increased, the locking became unstable. The gain of the pre-amplifier was increased by a factor of 10 to regain stability. The maximum resonance frequency obtained by increasing the offset was 310 Hz. Increasing the gain of the pre-amplifier by 10 meant that to display the error signal on the same scale as the other two, it had to be divided by 10 (a subtraction of 20 dB). The voltage offset used to reach the higher resonance frequency also had to be divided by 10 to keep its size relative to the error constant. Figure 6.5 shows the traces of the three different frequency shifts, each labelled with the appropriate voltage offset and resonance frequency.

The frequency shifts were an exciting and unexpected experimental discovery which also provided the first means to test quantitatively the radiation pressure model.

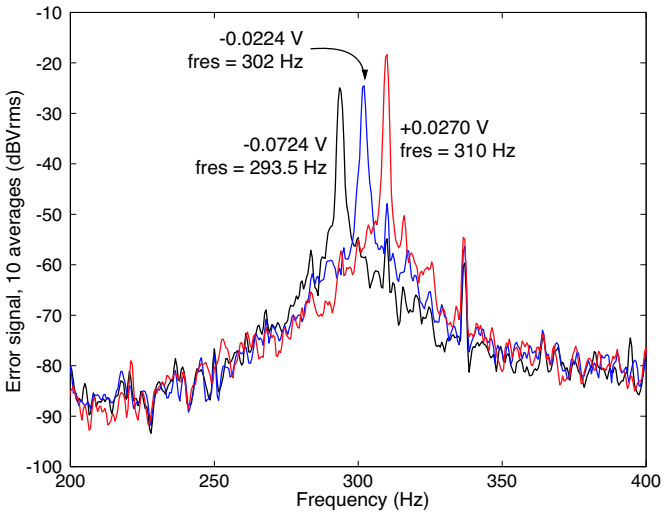

Figure 6.5: At different lock points the resonance frequency of the flexure was different. The voltage offsets, injected into the cavity locking servo, and the resulting resonance frequencies are shown for each peak.

6.3 Radiation Pressure

To verify that radiation pressure caused the observed non-linearities, it was critical that the radiation pressure model qualitatively and quantitatively agree with the experimental results. The non-linear scans and frequency shifting were the first observations of non-linear effects. Both these effects are qualitatively consistent with the radiation pressure model presented in chapter 4.

The frequency shifting measurements were the first quantitative test of the radiation pressure model. The error signal offsets which caused the frequency shifts were converted into cavity length changes. The length change versus resonance frequency was then superimposed on a plot like that in figure 4.4. To convert the voltage offsets into position offsets the slope of the error signal (measured in volts per metre) was determined. The error signal peak-to-peak separation is the same as one FWHM and the FWHM is the FSR divided by the finesse. The FSR is one wavelength, in round trip length, and half a wavelength in mirror displacement. The slope in the centre of the error signal is twice the average peak to peak slope. Combining all these premises the slope of the error signal for mirror displacement is given by:

$$\epsilon_{slope}(V/m) = \frac{4\mathcal{F}\epsilon_{sig,pp}}{\lambda} \tag{6.1}$$

where $\epsilon_{sig,pp}$ is the peak to peak height of the error signal. Note that the error signal can have positive or negative slope depending on the phase of the local oscillator, so it was sometimes necessary to flip the sign of this conversion factor. The experimental values for the required parameters were $\mathcal{F} = 434$, $\epsilon_{sig,pp} = 1.0$ volts, and $\lambda =$

1064 nm. The resulting error signal slope is:

$$\epsilon_{slope} = 1.63 \times 10^9 \quad \text{V/m} \tag{6.2}$$

The offset values shown in figure 6.5 were added to offsets already present in the locking loop. To allow for this, the offsets were converted to displacements from resonance, and then all were shifted until the 302 Hz point intersected the plot of resonance frequency versus displacement, as shown in figure 6.6. The model then predicts values for the frequencies for the other two offsets. The comparison between the model frequency and the actual frequency at these points tests quantitatively the supposition that the frequency shift was caused by radiation pressure.

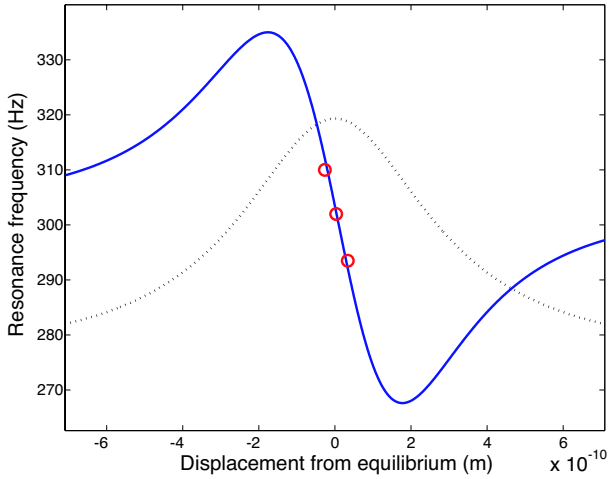

Figure 6.6: The radiation pressure prediction of the resonance frequency as a function of mirror position. The experimental values are the superimposed circles.

Exp. Freq. (Hz)	Offset (mV)	Corrected Disp. (m)	Model Freq. (Hz)
293.5	-72.4	3.41×10^{-11}	290.8
302	-22.4	3.49×10^{-12}	n/a
310	27.0	-2.68×10^{-11}	312.7

Table 6.1: Quantitative analysis of the frequency shift.

In the calculation of the frequency shift from the model there was significant error in both the moment of inertia (approximately 10%) and in the circulating power estimate. The circulating power was inferred from the predicted value shown in chapter 5, scaled by the ratio of the experimental finesse to the predicted finesse. Using this method, the power recycling factor was estimated to be 200. This implied that the average circulating power on resonance was approximately 45 Watts. However, to accurately determine the power recycling factor, the total cavity loss and the mirror reflectivities must be known. Also, the input power measurement (of 230 mW) had approximately 5% error due to power meter error. Given the ex-

perimental noise, there is excellent agreement between model and experiment. This is very strong evidence that radiation pressure causes the frequency shifts.

The next measurements used to quantitatively test the radiation pressure model were the measurements of the AM induced peak in the error signal. The amplitude modulator was driven at a given frequency and the modulation depth adjusted until a good signal to noise peak was seen. The modulation depth, relative to the average power, was recorded. Frequency modulation was also imposed on the input beam, at a slightly different frequency to the AM frequency (usually 10 Hz away), as shown in figure 6.7. The AM and FM drives were then held constant, and the average input power decreased. The power on the transmitted detector was held constant by increasing the fraction of transmitted power detected. The heights of the AM and FM peaks were recorded as a function of power at four different frequencies, 40 Hz, 180 Hz, 270 Hz, and 480 Hz. The 180 Hz measurements were recorded under different conditions to the measurements at the other three frequencies, therefore, the 180 Hz data are largely ignored during this analysis.

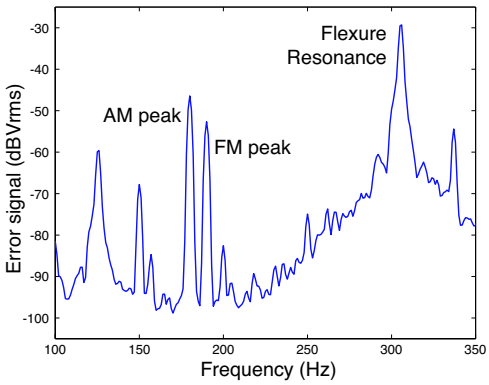

Figure 6.7: AM and FM induced peaks are at 180 and 190 Hz respectively.

The FM peaks were used as a comparative tool. The FM effect was independent of input power, confirming the validity of the measurement strategy. If radiation pressure was the only effect converting incident AM into output error signal, then the red lines fitted to the AM peak height measurements in figure 6.8 would be straight and linearly proportional to power. However, in three plots the AM line is not straight, and there is an AM effect in the error signal which is constant with input power.

A likely cause of the power independent AM effect is modulation of the error signal size. The equations defining the error signal for small deviations (equations 3.9 and 3.10) are dependent on the input power. If this fluctuates, the size and slope of the error signal also fluctuates. If the cavity error signal has some offset, then the size of the offset is modulated by the incident AM. This causes a peak at the AM frequency. It is not only offsets which mix with the AM at the error signal. Deviations from resonance at any frequency mix with the amplitude modulation, causing peaks at the sum and difference frequencies. Intermodulation products between the

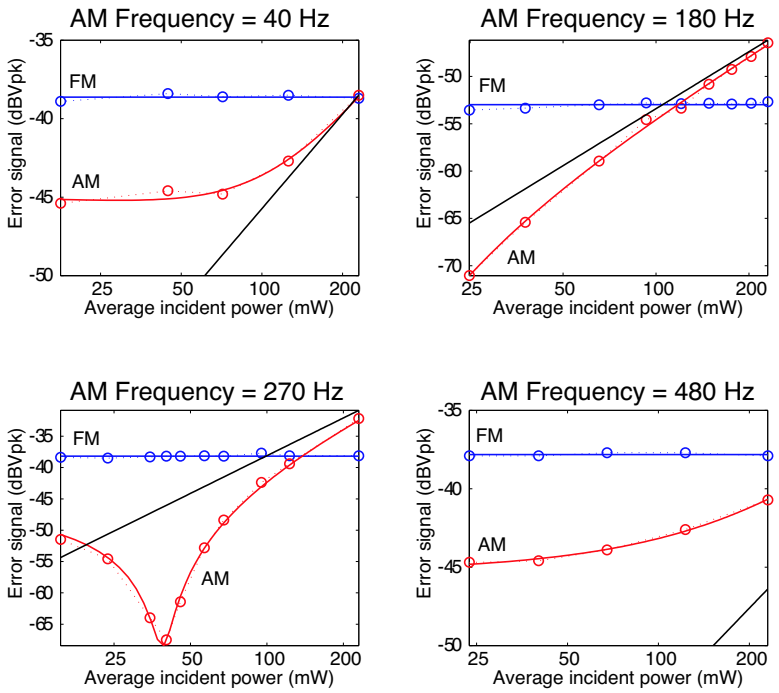

Figure 6.8: The error signal peak height caused by AM and FM drives as a function of power, at four different drive frequencies. The third line which is straight and diagonal is the power dependent component of the AM fit.

AM and the resonance frequency were observed, supporting this hypothesis.

The power dependent and power independent effects have a phase relation which is constant with power, but changes with frequency. The red lines in figure 6.8 are fitted to the AM data points using a term which scales linearly with power, a power independent term, and a phase relationship between the two. For the three consistent plots (the 40 Hz, 270 Hz, and 480 Hz plots), the constant offset term is approximately the same. In all cases the fitted curve matches the data closely.

The power dependent part of the AM fit is plotted separately on each graph as a straight, black, diagonal line. This linear component of the fit separates the motion induced through optical effects and the power independent component. The amplitude of motion implied by the power dependent component is compared with the amplitude of motion predicted by the radiation pressure model.

The data presented in table 6.2 show that the power dependent motion calcu-

Frequency (Hz)	Exp. fit (m)	RP model (m)	Ratio
40	1.02×10^{-11}	1.21×10^{-11}	1.19
270	2.47×10^{-11}	5.74×10^{-11}	2.32
480	4.14×10^{-12}	7.89×10^{-12}	1.91

Table 6.2: The power dependent component of the experimentally measured motion is compared with model predications for RP induced motion.

lated from experimental data is substantially different from the radiation pressure model predictions. The discrepancies are explained in two stages. Firstly, all the experimental values are lower than the model predictions. A possible cause of this is peak broadening in the error signal spectrum. The incident AM has very narrow linewidth, but this may have been broadened either by mixing with DC noise (a non-linear process in the cavity), or by the windowing function of the spectrum analyser. If the peak is broadened, the peak height represents only a fraction of the motion. The area under the peak must be integrated to determine the actual motion. Integrating the peak area would increase the experimentally observed motion by approximately a factor of two. The increase would be similar at each frequency.

Secondly, the 40 Hz motion is too great with respect to the 270 Hz motion. The most likely candidate for this discrepancy is radiometry. Radiometric effects typically have timescales of minutes, but the dielectric coating on the movable mirror is very thin, and can change temperature rapidly. Air molecules colliding with the 'hot' coating gain energy, and the recoil momentum is higher than the incident momentum. Thus, the coated surface receives a small unbalanced momentum kick.

There are three important things to note about this sort of radiometry. Firstly it is an optical, power dependent effect almost indistinguishable from radiation pressure over small bandwidths. Secondly, the effect relies on temperature changes in the mirror surface and is likely to decrease in strength as frequency increases. Thirdly, it is almost impossible to remove the radiometric effect without putting the system in vacuum. There was probably a significant radiometric effect at low frequencies (below 200 Hz). This is consistent with the low frequency roll-off of the error signal response shown in figure 6.10. Radiometry could also have caused the higher than expected response to 40 Hz AM shown in table 6.2.

The inconsistency between the 480 Hz and 270 Hz experimental peaks, compared with the model, is within experimental error. The error occurs because there are so few data points to fit the power dependent part of the AM curve of the 480 Hz plot in figure 6.8. To confirm the hypotheses here regarding peak width and radiometry, considerable further work is required. The plans for further testing are discussed at some length in chapter 7.

The radiation pressure model was largely in agreement with the experimental values, both qualitatively and quantitatively. Where there were differences, a consistent explanation was given for a possible source of the errors. Despite the power independent effect and the possible pollution of radiometry, radiation pressure effects are present and observable throughout the signal bandwidth used for cancellation. In addition, radiation pressure dominates the non-linear response around resonance ,in the frequency range 200-350 Hz.

6.4 Single Frequency Cancellation

Single frequency cancellation was the major goal of this project, and successful measurement of such cancellation was sufficient for completion of the project. To generate correlations between frequency and amplitude, two phase locked Stanford

Research Systems 0-30 MHz signal generators were used to drive the respective modulators. The relative phase of the two signals was adjustable over a full cycle, with increments of one degree. The absolute amplitude of each signal could be varied over the full range of each signal generator output.

A method was developed to consistently achieve cancellation. The amplitude of the AM drive amplitude was adjusted until the peak height was approximately 35 dB above background noise. The drive amplitude was then held constant throughout the single frequency cancellation experiment. The AM drive frequency was set to the desired cancellation frequency. The FM signal was then added at a slightly different frequency, about 10 Hz separated from the AM peak. Figure 6.7 in section 6.3 shows AM and FM induced peaks side by side. The amplitude of the FM drive was altered until the AM and FM peaks were the same height. The FM frequency was then changed to be the same as the AM frequency. At this point the AM and FM became phase locked. The relative phase of the two drives was adjusted until the combined peak was a minimum. The amplitude of the FM drive was then tweaked to reduce the combined peak further. The phase of the FM drive was then reversed to see the height of the peak in the constructive interference limit.

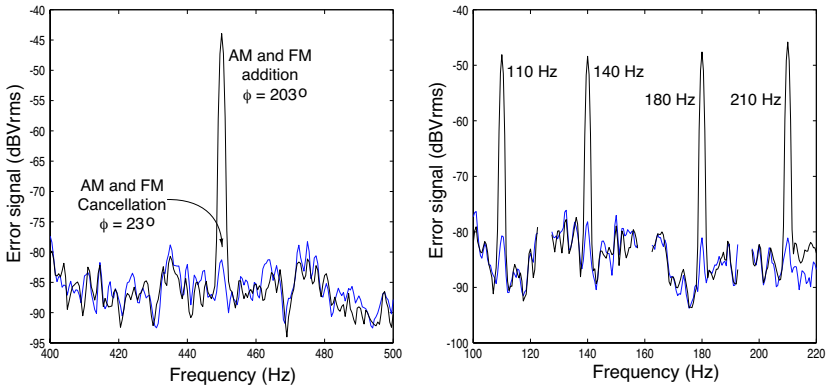

Figure 6.9: The left hand plot shows cancellation at 45 Hz with the relative phase of the drive signals shown. The right hand plot shows cancellation at four other frequencies, with the data windowed around the relevant frequency.

For all frequencies where cancellation was observed, the amplitude of the FM drive and the relative phase of the drives were recorded. The relative phase between the drives at cancellation and addition was measured using a lock-in amplifier. At each cancellation frequency, the two peaks were cancelled to the noise floor, approximately 35 dB of cancellation. Figure 6.9 shows two plots, the first is single frequency cancellation at 450 Hz with the relative phase of the two drives explicitly marked. The two traces show constructive and destructive interference between the AM and FM signals. Note that the background noise features are unaffected by the state of the interference. The second plot shows cancellation at four other frequencies, with only the data immediately around the cancellation frequency shown. Table 6.3 lists all the frequencies at which signal cancellation was recorded, with the respective

FM drive and relative drive phase at cancellation.

Frequency (Hz)	FM Drive (mV)	Phase at Minimum
20	70	-60°
40	50	-68°
65	32	-76°
80	28	-81°
110	23	-104°
140	22	-119°
180	24	-144°
210	30	-159°
270	101	-190°
330	145	-377°
390	55	-381°
450	37	-383°
510	30	-385°
800	20	-389°

Table 6.3: The frequencies at which signal cancellation was recorded. The AM drive was constant at 75 mV. The phase at minimum is the phase difference between the AM and the FM drives at best cancellation.

The phase response of the flexure in steady state should be flat for frequencies greater than 10 Hz away from the flexure resonance, as shown in figure 4.3. It was already known that the response of the frequency modulation to input drives was flat in both phase and amplitude. Thus, the phase of cancellation in 6.3 was surprising, since it shows significant change with frequency. To better understand this, the transfer function from the AM drive to the error signal was recorded. The swept sine output from the SR 785 signal analyser was injected into the AM signal injection port (shown in figure 5.5). The output was the cavity error signal, measured at the unused output of the SR 560 pre-amplifier (shown in figure 5.7). The amplitude and phase response of the cavity error signal was recorded. The response at resonance was not recorded because the motion excited at resonance was too large for the cavity to stay locked.

Superimposed on the transfer function in figure 6.10 are the data points from table 6.3. To make the data from the table compatible with the transfer function it was necessary to convert the FM drive amplitude to dB units using a global scaling parameter. The single frequency cancellation data clearly fits the AM to error signal transfer function very well.

With the success of the single frequency signal cancellation came the completion of the major goal of this experiment. Mechanical motion was generated using amplitude modulation, and the effect of that motion on the error signal was cancelled using frequency modulation. It is important to note that the signals were cancelled in a coherent process rather than being suppressed by something like a feedback system. The cancellation achieved always reduced the input peaks to the noise floor, usually a suppression of 35 dB, or a factor of approximately 50 in amplitude.

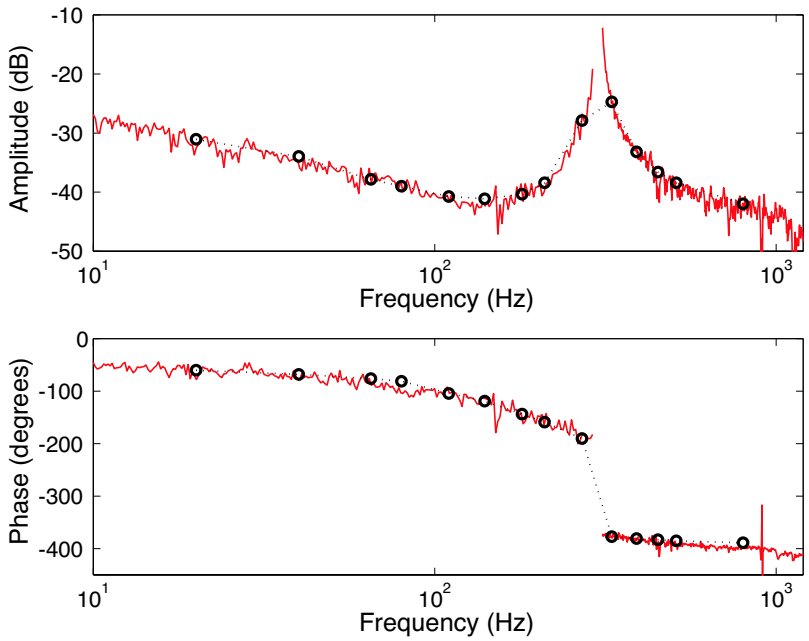

Figure 6.10: The AM drive to cavity error signal transfer function. Also plotted are the data from table 6.3

6.5 Broadband Cancellation

A more ambitious goal of this project was to cancel 'noise' over some finite bandwidth. This was a much more challenging problem experimentally. For single frequency cancellation, two phase locked sources were used. It was simple to adjust their relative amplitude and phase. For noise cancellation, however, a single source was used, to make sure the AM and FM drives were highly correlated. The noise is split and one path was given a phase and amplitude shift.

To achieve cancellation over the full signal bandwidth required matching the relative amplitudes and phases of the noise drives to the AM to error signal transfer function in figure 6.10. This was not possible with the equipment avaliable, and in general requires either a very complex analogue filter, or a programmable digital filter. The decision was then made to attempt the cancellation over a small, flat part of the transfer function response. The region chosen was from 450 to 550 Hz.

Shifting the relative amplitudes over this bandwidth was straightforward, and was done using a variable attenuator placed on the FM drive. Shifting the phase by an amount constant over the measurement bandwidth, in a precise and controllable way, was far from easy. Several attempts were made before one was successful.

For the successful attempt, the AM noise drive was connected to a SR 560 pre-amplifier. The input was split before the pre-amplifier, with one part going directly to the summer box on the input of the amplitude modulator. The other part was processed by the pre-amplifier, which high pass filtered the signal at 10 kHz, with a 6 dB per octave roll-off. From 450-550 Hz, this gave a relatively flat phase shift of

90°. The high pass attenuation was also relatively flat over this bandwidth. A flat gain of approximately 200 was added to correct for the attenuation. The output was connected to the AM input summer box. The two noise inputs at the AM summer box now had a 90° phase difference, so by adjusting their relative amplitudes (using the fine scale adjustment of the pre-amplifier gain), the phase of their sum changed.

Through this rather convoluted mechanism the phase and amplitude of the two noise drives could be adjusted quite freely. The response of the phase shift system was also sufficiently flat over the 100 Hz bandwidth to allow moderate cancellation. A more precise method of controlling the phase of the drives is planned in the further work section.

To test how the broadband cancellation interacted with an independent signal in the cancellation bandwidth, a signal was injected into the error signal through the cavity locking servo. The data were recorded by measuring the error signal spectrum with five different levels of signal input: with only the background noise; the background noise with a single frequency signal injected into the cavity servo at 517 Hz; with bandwidth limited white noise connected to the amplitude modulator through the summer box (in the arrangement described above) with noise output adjusted until the signal was masked; with the same output noise used for the AM connected to the FM drive, and the variable attenuation adjusted until the level was the same as for the AM noise; and finally the noise source was used to drive the AM and FM simultaneously. The results are shown in figure 6.11.

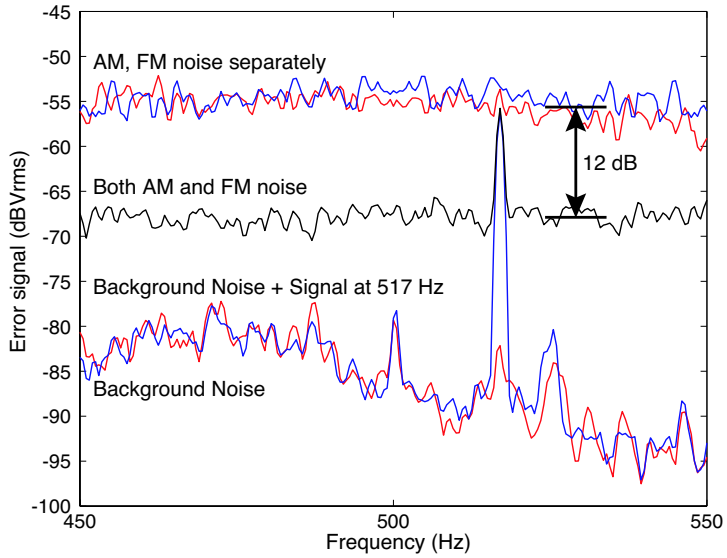

Figure 6.11: The signal to noise ratio is demonstrably improved when both the amplitude and frequency modulators are driven with noise, the effect of each was large enough to mask the signal, but when both were present, they partially cancelled.

The signal at 517 Hz was completely masked by either AM or FM noise. However, when both AM and FM were noisy, the noise level dropped by approximately 12 dB. This improved the signal to noise ratio of the 517 Hz signal from 1 to 4. The signal shape and height is unchanged after the cancellation, further proof that the AM and FM are interacting in a coherent manner. If the noise reduction was achieved using some sort of suppression, the signal would also have been suppressed, and the signal to noise ratio would have been constant. The successful measurement of broadband cancellation completed the most optimistic of the pre-project goals.

6.6 Results Summary

It has been shown that the experiment was susceptible to optical non-linear effects. Of these, the most interesting was the observation that the mechanical resonance frequency of the flexure was altered by interaction with the circulating power in the cavity. These effects were compared with a radiation pressure model of the system and it was determined that radiation pressure was the cause of the non-linearites.

Most importantly, correlations were imprinted on a laser to mimic the quantum correlations caused by squeezing. These correlations were created through modulating the beam in the same quadratures as those which cause shot noise and radiation pressure noise. It was shown that when the phase of the correlation was correct, the noise caused by the modulations was reduced. This noise cancellation is the classical analogue of the quantum process which allows breaching of the standard quantum limit.

Chapter 7

Summary and Further Work

At the end of an experimental project there is usually further work that could be done. In this case, further work will be done and the material presented in this chapter will be used as a guide for that work. For this reason, the suggestions made here are experimentally viable and will be implemented in the coming months. A summary of the project outcomes is also presented.

7.1 Further work

The further work planned is divided into two categories: refinements and revisions. Refinements are small changes which do not involve significant alterations of the experiment. Revisions require major changes, but may yield the greatest results. It may not be possible to implement all the suggestions in the revisions section.

7.1.1 Refinements

The detector used to measure how much power is incident on the cavity (I/P-PD) will be changed. The current detector is a high power detector and has approximately 25 mW of optical power incident on it. The temperature of the detector changed significantly with this much incident power, creating some long time constant drift in the output voltage. The planned update involves heavily attenuating the measured power using a non-absorptive method, such as a wedge, and a more standard low power detector. As an extension of this, all the multi-order waveplates currently in the main beam path will be replaced with zero order waveplates, to reduce temperature effects with high power.

Further modification of the servo is planned. The servo interacts with the cavity in a rather unpredictable manner, requiring that the locking bandwidth be very low. To improve this performance, the post unity gain roll off will be increased substantially from the current 1/f roll off. To do this, the third order elliptic filter currently at 900 Hz will be moved to approximately 20 Hz. Third order elliptic filters combine fast roll off with minimal phase disturbance. Further information regarding such filters can be found in reference [34].

The flexure will be immobilised. By immobilising the flexure, possibly with a wedge jamming the flexure between the movement limiters, radiation pressure and

radiometric effects will be (all but) eliminated. This will allow better diagnosis of AM effects not relating to motion of the flexure.

The average power of the input light will be increased. Previously the Michelson was locked half-way up a fringe, to give the greatest range of modulation depth. During the experiment, the modulation depth was almost always less than 20% of the average power. Thus, it should be possible to lock the Michelson with at least 75% of the incident power transmitted, and modulate around that point. Also, the tap-off which provides the Michelson locking signal is currently an uncoated wedge. This removes approximately 10% of the laser power. To reduce loss, it will be replaced by an anti-reflective coated surface, which should only reflect about 1% of the laser power.

Using the improved input detector arrangement and higher optical intensity, the AM to error signal transfer function will be re-recorded as a function of input power, from high to low. It is hoped that this will isolate the spectrum of the power dependent effects.

A superior means of controlling the low frequency phase will be developed to improve the broadband noise cancellation. The system will likely use a quadrature phase shifter, which has one input and four outputs which are 90 degrees apart. These outputs can be summed with variable gain to give arbitrary phase. Such a system was constructed during the experimental phase, but its performance was poor.

7.1.2 Revisions

The revisions planned here require substantially more work than the previous suggestions, but the pay-offs should by similarly higher.

Previously, there was a single beam incident on the cavity. It was high power, circularly polarised, and was modulated with the necessary AM, the cancelling FM, and the control scheme PM. However, AM interacts with the derivation of the cavity error signal in an unfortunate way. To correct for this, the high power beam will have linear polarisation, amplitude modulation, and frequency modulation. A second probe beam of the orthogonal polarisation will then be used to sense the resonance condition of the cavity. The probe beam will be generated by the same laser, and will most likely be taken from the currently unused output of the first PBS in the beam path in figure 5.4. The probe will carry the same FM as the high power beam and will also carry the phase modulation used to control and sense the cavity resonance condition.

Under this revised scheme, the high power beam now serves solely to affect the cavity through radiation pressure. The probe beam senses the condition of the cavity without disturbing it. As an additional benefit, the high power beam won't pass through the phase modulator, another cause of loss. This revision should remove the power independent error signal modulation effect described in section 6.3. Part of the experimental challenge will be to mode-match both beams to the cavity.

The construction of a small, portable vacuum chamber would greatly advance this experiment. It is not known yet whether such a system will be built, so only a

brief outline of the potential advantages will be given. Radiometry is an effect which is optical and power dependent, and as such is extremely difficult to differentiate from radiation pressure. However, radiometry is caused by the movable mirror interacting with the surrounding gas, so removal of the gas should eliminate radiometry. This is perhaps the greatest advantage of placing the system in vacuum.

Acoustic noise is all but eliminated by placing the system in vacuum. Also, if the vacuum chamber is small it may be possible to suspend it from some form of mechanical isolation on the table. This last suggestion is highly speculative, and as yet unplanned.

A higher power laser would generate larger radiation pressure effects. The acquisition of such a laser is planned, though it is unlikely to be used in this experiment.

Implicit in all the above suggestions is the assumption that appropriate data will be recorded after each improvement. It is hoped that most of the above suggestions will be implemented before February next year.

7.2 Conclusion

This thesis described the background and motivation for opto-mechanical noise cancellation in terms of gravitational wave detection. The concept of the standard quantum limit was introduced, along with a means of breaching that limit. The classical equivalent to breaching the SQL was described. Sufficient background material was presented to adequately describe the experiment and analyse the results. The experimental design, from pre-project work through to optical layout, was described. The justification for several important design features was given.

The important results were presented and explained. These were the observation of non-linear cavity scans; the measurement of an optical effect shifting a mechanical resonance frequency; comparisons of experimental data with a model of radiation pressure interaction; demonstration of single frequency cancellation; and demonstration of broadband opto-mechanical noise cancellation improving signal to noise ratio. The cancellation of motion induced by radiation pressure with a correlated frequency signal is directly analogous to using squeezing to breach the standard quantum limit.

Appendix A

Pre-project Flexure Design

The aim of the pre-project design was twofold: to determine whether the project was viable as an honours project, and if so to design the flexure system to be used in the project. Though this work was also completed by the author, it was not strictly performed during the standard honours year.

The two aims of the project became, in effect, one. A design was hypothesised, then tested using the first radiation pressure interaction model. The design was then modified in an effort to:

•Increase the DC (or below resonance) amplitude response

•Increase the resonance frequency, and

•Determine whether the induced motion was detectable.

To achieve these tasks, is was necessary to reduce the mass, and moment of inertia as much as possible. However, the moment of inertia was limited by several factors, including the dimensions and placement of the mirror. The dimensions of the mirror were chosen early in the design phase to be 12.7 mm in diameter and 1 mm thick and weighing approximately 300 mg. The placement of the mirror greatly affected the response. Having the mirror higher above the flexure increased moment of inertia, and hence decreased resonance frequency, however it increased both the optical torque and the linear displacement, by virtue of being a longer lever arm.

There were also the requirements that the flexure membrane be wide compared with its thickness and that it be possible to manufacture the flexure using electrical discharge machining.

Three designs were chosen to span a reasonable range of parameters. The first had a 120 micron thick flexure which was 1 mm long, and a lever arm 12.7 mm long. The second was the same as the first, but with a 240 micron thick flexure. The third had an 80 micron thick, 2 mm long flexure, with a 19.05 mm long lever arm. Each flexure system was labelled by the thickness of the membrane ('the 120 micron flexure').

The transfer functions of the three designs chosen were plotted as a function of frequency. The transfer functions are shown in figure A.1. The amplitude displacement is for a given optical force. To calculate the optical force an incident average optical power of 250 mW was used, with full modulation depth ($\delta = 1$), and a cavity finesse of 1000. A diagram of the flexures is shown in figure A.2. It constrains all the physical dimensions of the moving system. Note that the mirror was to sit on the little 'ledge' seen in the side view. A picture of the 240 micron design is shown

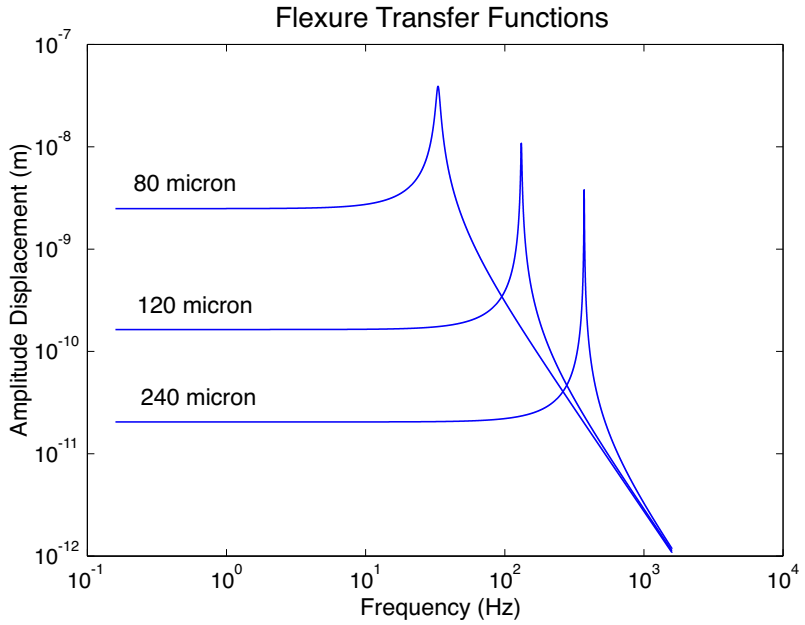

Figure A.1: The transfer functions of the three desisgns chosen for construction. Both the amplitude response and the resonance frequency change significantly between designs.

in 3-D in figure 4.2.

In summary, three designs were chosen for possible use a moveable mirror spring. The design spanned a wide range of parameters for spring constant and moment of inertia. They were designed using a radiation pressure model using realistic values for optical power, modulation depth, and cavity finesse. The lightweight mirror to be used in conjunction with the flexure designs was also chosen at this stage. The combined flexure-mirror systems were predicted to have resonance frequencies of 32 Hz, for the 80 micron flexure, 130 Hz, for the 120 micron flexure, and 370 Hz, for the 240 micron flexure. The predicted DC displacement was approximately 3 nanometers (80 micron flexure), 0.2 nanometers (120 micron flexure), and 20 picometers (240 micron flexure). The aim of choosing such a range of parameters was that one set would be suitable under experimental conditions.

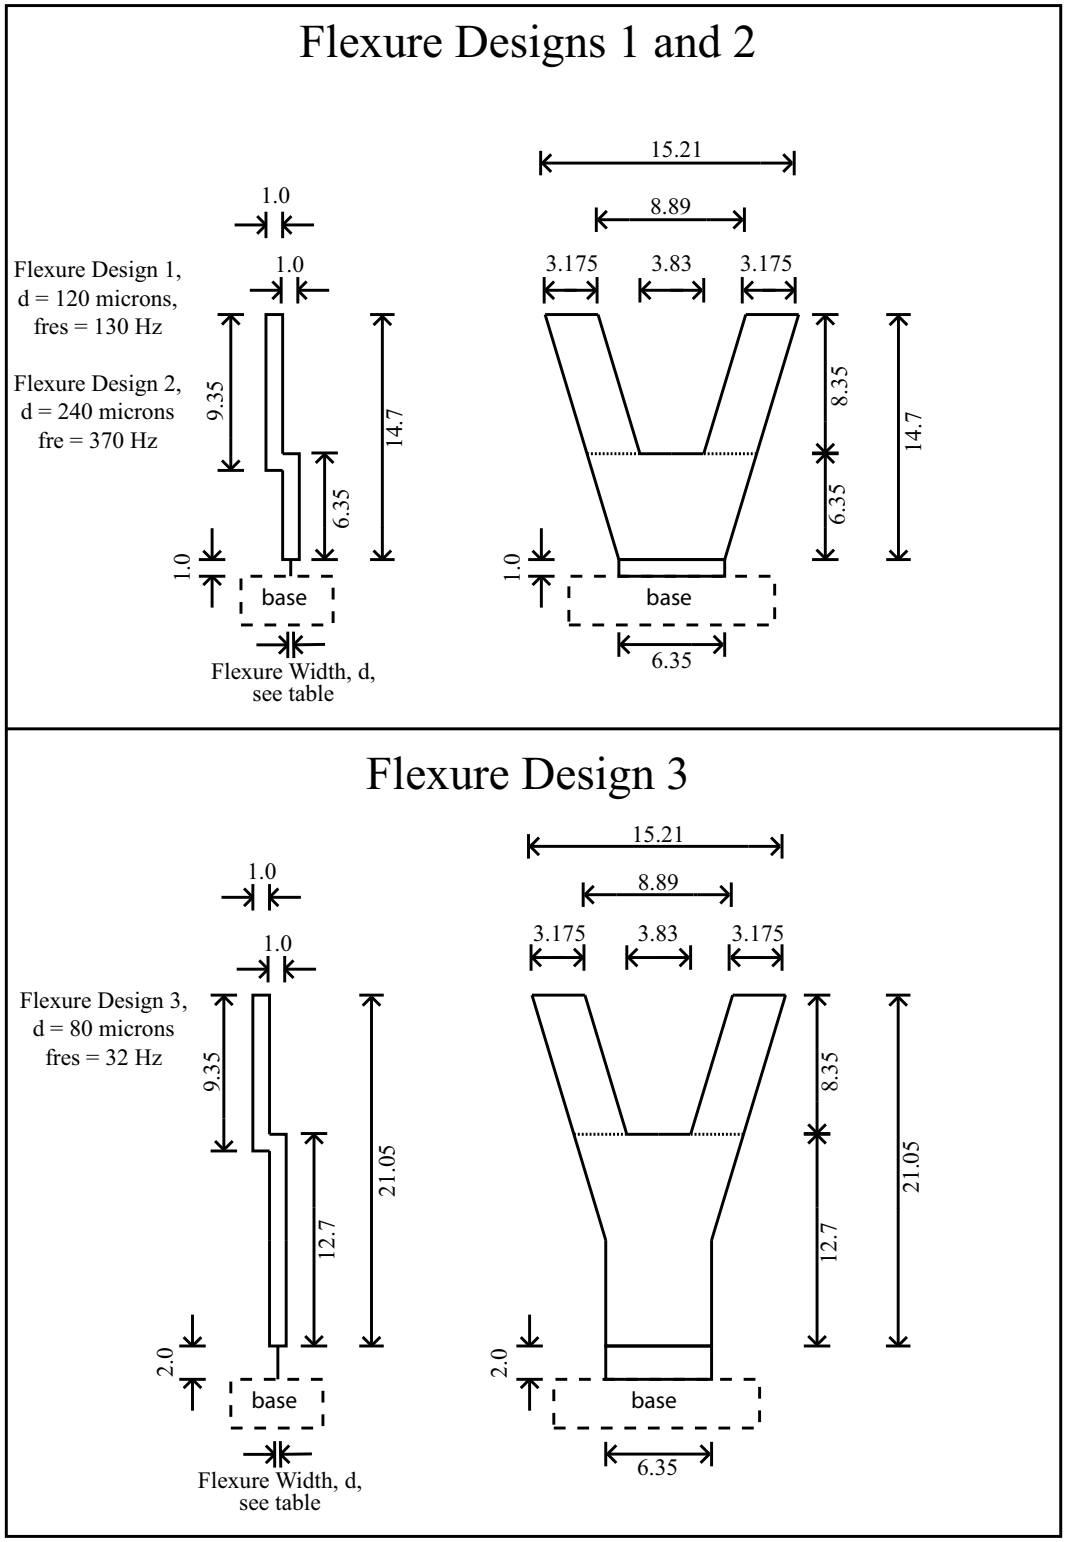

Figure A.2: The three designs selected at the end of the pre-project design phase.

Appendix B

Radiation Pressure Model

The Matlab code used to model the radiation pressure effects. The principal variables are the input power, modulation depth, and error signal size. The thickness was altered until a mechanical resonance of 303 Hz was calculated.

```
%% Program for plotting Flexure response using Pro-E %%
%% Values for mom-of-inertia, and measured finesse %%
clear all

Pin = 0.25; % Laser Power (Watts)
AMconst = 2.67; % Average incident power (arb units)
AMpp = 0.61; % Peak to peak modulation (arb units)
delta = AMpp/(2*AMconst); % Modulation depth
esigpk = 0.5; % zero-pk size of error signal (Volts)

r1 = sqrt(0.990); % Input coupler amplitude reflectivity
r2 = sqrt(0.9956); % Back mirror amplitude reflectivity
t1 = sqrt(1-r1^2); % Transmissivity of input coupler
c = 3e8; % Speed of light (m/s)
lambda = 1064e-9; % Wavelength (m)
wopt = c*2*pi/lambda; % Angular frequency of light (rad/s)
scan = lambda/1500; % Length scan (m)
nsteps = 1000;
x = linspace(-scan,scan,nsteps);
xrt = 2*x; % Round trip length change
phi = wopt*xrt/c; % Round trip phase of cavity

Ecirc = i*t1./(1-r1*r2*exp(i*2*phi));
Pcirc = Ecirc.*conj(Ecirc)*Pin; % Circulating power
Precyc = t1^2/(1-r1*r2)^2; % Power recycling factor
Fin1 = pi*(r1*r2)^(1/2)/(1-r1*r2); % Finesse
scr = get(0,'ScreenSize');
FHWM = lambda/Fin1;

F = 2*delta*max(Pcirc)/c; % Force on mirror (N)
```

```
%% Flexure Parameters %%
Ey = 105e9; % Youngs Modulus (N/m^2)
W = 6.35e-3; % Width of flexure (m)
l = 1e-3; % Length of flexure (m)
t = 202e-6; % Thickness of flexure membrane (m)
% Itot = 1.406e-7; % Moment of inertia of system (kgm^2)
Itot = 1.262e-7; % PRO-E mom of inertia (kgm^2)
Q = 1000; % Quality factor
hm = 1.27e-2; % Height of mirror centre above flexure (m)

Ia = W*t^3/12; % Moment of area (m^4)
ka = Ey*Ia/l; % Angular spring constant (Nm/rad)
T = F*hm; % Optical torque (Nm)
wres = sqrt(ka/Itot); % Undamped Resonance Frequency (Rad/s)
wdamp = wres*sqrt(1-1/(2*Q)); % Corrected resonance
fres = wdamp/(2*pi) % Resonance Frequency (Hz)
b = Itot*wres/Q; % Damping Co-efficient (kg m^2/s)

npoints = 1000;
fstart = 1;
ffin = 1.5e3;
w = logspace(log10(fstart*2*pi),log10(ffin*2*pi),npoints);
%% Driving frequency (Rad/s) %%

%% Amplitude response (in radians) of the mirror%%
theta = T./sqrt(Itot^2*((wres^2-w.^2).^2) + b^2*w.^2);
A = theta*hm; % Linear displacement of mirror
ampk = 20*log10(A*(esigpk*2)/(sqrt(2)*FHWM/2));

esig = 1; % Pk-pk size of error signal (V)
FWHM = lambda/Fin1; % Length scale of cavity linewidth
angx = x/hm; % Radian displacement of mirror in scan
Tx = ka*angx; % Mechanical torque (Nm)
Txopt = 2*Pcirc*hm/c; % Optical torque (Nm)
Txtot = Tx+Txopt; % Total torque (Nm)

for n=2:nsteps
katot(n-1) = (Txtot(n)-Txtot(n-1))*hm/(x(n)-x(n-1));
kaopt(n-1) = (Txopt(n)-Txopt(n-1))*hm/(x(n)-x(n-1));
% The total, and optical spring constants
end

%% The total resonance frequency %%
frestot = sqrt(katot/Itot)/(2*pi);
```

Appendix C

PZT Servo Design

The PZT servos used for both the Michelson locking and the test cavity locking are based around the same design. A schematic for a standard PZT servo is shown in figure C.2. The servos have a standard 1/f response at frequencies above the pole. The integrator, when active, cancels this pole with a zero at the same frequency to give infinite gain at DC with a 1/f roll off from DC. In addition, there is a third order elliptic filter which provides a very sharp roll off. This filter is generally at high frequencies, around 20 kHz, to cancel the mechanical resonance of the PZT. For the cavity locking servo, the elliptic filter was moved down to 900 HZ, to cancel a PZT mount resonance, and there are plans to lower it further in frequency.

The cavity servo was modified in other ways as well, these changes are briefly listed in chapter 6. The open loop transfer function of the servo, post modification, is included here as figure C.1.

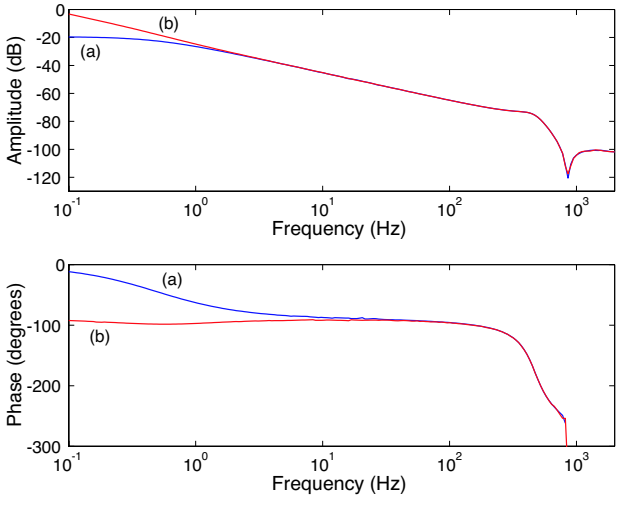

Figure C.1: The open loop transfer function of the cavity PZT servo with (a) the integrator off and (b) the integrator on.

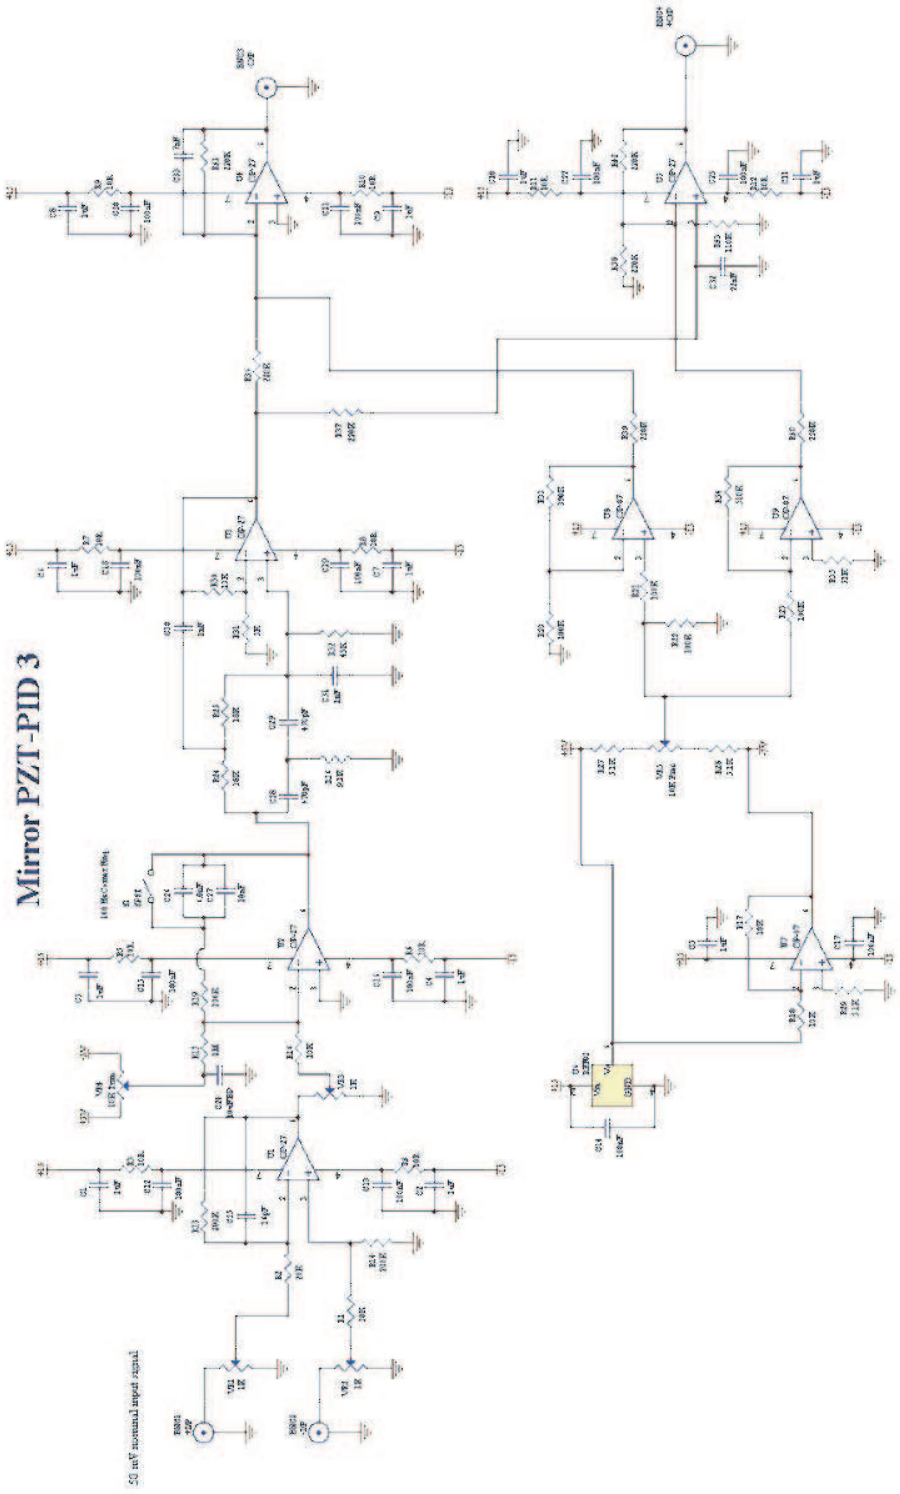

Figure C.2: A standard PZT servo circuit diagram

Bibliography

[1] A. Einstein, Ann. Phys. Lpz **49**, 769 (1916).

[2] A. Einstein, Preuss. Akad. Wiss. Berlin , 688 (1916).

[3] C. W. Misner, K. S. Thorne, and J. L. Wheeler, *Gravitation*, Freeman, San Francisco, 1973.

[4] P. R. Saulson, *Fundamentals of interferometric gravitational wave detectors*, World Scientific, Singapore, 1994.

[5] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gursel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, K. S. Thorne, R. E. Vogt, R. Weiss, S. E. Whitcomb, and M. E. Zucker, *LIGO: the Laser Interferometer Gravitational-Wave Observatory*, Science **256**, 325 (1992).

[6] A. Brillet et al., Virgo proposal to CNRS and INFN, unpublished, 1989.

[7] H. Lück, P. Aufmuth, O. S. Brozek, K. Danzmann, A. Freise, S. Goßler, A. Grado, H. Grote, K. Mossavi, V. Quetschke, B. Willke, K. Kawabe, A. Rüdiger, R. Schilling, W. Winkler, C. Zhao, K. A. Strain, C. Cagnoli, M. Casey, J. Hough, M. Husman, P. McNamara, G. P. Newton, M. V. Plissi, N. A. Robertson, S. Rowan, D. I. Robertson, K. D. Skeldon, C. I. Torrie, H. Ward, B. F. Schutz, I. Taylor, and B. S. Sathyprakash, *The status of GEO600*, Gravitational Waves: Third Edoardo Amaldi Conference, Pasadena. ed S. Meshkov , 119 (2000).

[8] M. Ando, K. Tsubono, and the TAMA collaboration, *TAMA Project: design and current status*, Gravitational Waves: Third Edoardo Amaldi Conference, Pasadena. ed S. Meshkov , 128 (2000).

[9] E. Gustafson, D. Shoemaker, K. Strain, and R. Weiss, *LSC White Paper on Detector Research and Development*, LIGO T990080-00-D (1999).

[10] C. M. Caves, K. S. Thorne, R. W. P. Drever, V. D. Sandberg, and M. Zimmermann, *On the measurement of a weak classical force coupled to a quantum-mechanical oscillator. I. Issues of principle*, Rev. Mod. Phys. **52**, 341 (1980).

[11] C. M. Caves, *Quantum-mechanical noise in an interferometer*, Phys. Rev. D. **23**, 1693 (1981).

[12] V. B. Braginsky, Y. I. Vorontsov, and K. Thorne, *Quantum nondemolition measurements*, Science **209**, 547 (1980).

[13] P. F. Cohadon, A. Heidmann, and M. Pinard, *Cooling of a mirror by radiation pressure*, Phys. Rev. Lett. **83**, 3174 (1999).

[14] I. Tittonen, G. Breitenbach, T. Kalkbrenner, and et. al., *Interferometric measurements of the position of a macroscopic body: Towards observations of quantum limits*, Phys. Rev. A **59**, 1038 (1999).

[15] A. Dorsel, J. D. McCullen, P. Meystre, E. Vignes, and H. Walther, *Optical Bistability and Mirror Confinement Induced by Radiation Pressure*, Phys. Rex. Lett. **51**, 1550 (1983).

[16] A. A. Michelson and E. W. Morley, Am. J. Sci. **34**, 333 (1887).

[17] A. E. Siegman, *Lasers*, University Science Books, Mill Valley, California, 1986.

[18] Amnon Yariv, *Optical Electronics in Modern Communications*, Oxford University Press, New York, 5th edition, 1997.

[19] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, John Wiley & Sons, Inc., New York, 1991.

[20] E. Hecht, *Optics*, Addison-Wesley, Reading, 2nd edition, 1987.

[21] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, *Laser phase and frequency stabilization using an optical resonator*, Appl. Phys. B **31**, 97 (1983).

[22] D. A. Shaddock, *Advanced Interferometry for Gravitational Wave Detection*, PhD thesis, Physics Department, Australian National University, Canberra, Australia, 2000.

[23] S. E. Whitcomb, 2002, Private communication.

[24] D. Halliday, R. Resnick, and J. Walker, *Fundamentals of Physics*, John Wiley & Sons, Inc., New York, 5th edition, 1997.

[25] P. A. Tipler, *Physics: for Scientists and Engineers*, WH Freeman & Co., 4th edition, 1999.

[26] L. Ju, D. G. Blair, M. Taniwaki, and R. Andrew, *The quality factor of niobium flexure pendulums*, Phys. Lett. A **254**, 239 (1999).

[27] L. Ju, M. Baker, and D. G. Blair, *High-Q niobium membrane flexure pendulum*, Phys. Lett. A **280**, 182 (2001).

[28] D. Roylance, *Mechanics of Materials*, John Wiley & Sons, Inc., New York, 1996.

[29] J. Winterflood, 2000, Private communication.

[30] D. G. Chetwynd and S. T. Smith, *Foundations of Ultraprecision Mechanism Design*, Gordon and Breach science publishers, Switzerland, 1992.

[31] Trelleborg Corporation, 2002, www.sorbothane.com.

[32] G. Sagnac, *L'éther lumineux démontré par l'effet du vent relatif d'éther dans un interféromètre en rotation uniforme*, C. R. Acad. Sci. **95**, 708 (1913).

[33] M. B. Gray, D. A. Shaddock, C. C. Harb, and H.-A. Bachor, *Photodetector designs for low-noise, broadband, and high power applcations*, Rev. Sci. Instrum. **69**, 3755 (1998).

[34] A. B. Williams and F. J. Taylor, *Electronic filter design handbook: LC, active and digital filters*, McGraw-Hill, New York, 2nd edition, 1988.