

Gordon, Neil Andrew (2015) *Characterisation and control of coupled optical springs for future gravitational wave detectors.* PhD thesis.

http://theses.gla.ac.uk/6532/

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Glasgow Theses Service http://theses.gla.ac.uk/ theses@gla.ac.uk

Characterisation and Control of Coupled Optical Springs for Future Gravitational Wave Detectors

Neil Andrew Gordon, MSci.

School of Physics and Astronomy College of Science and Engineering University of Glasgow

A thesis submitted to the University of Glasgow in accordance with the requirements for the degree of Doctor of Philosophy.

July 4, 2015

ABSTRACT

The detection of gravitational waves in the years to come will be the result of decades of work by institutions around the world to design and construct instruments to probe the characteristic strain of space-time on the smallest scales. Once detected, an entirely new field of study will be open to astronomy, allowing probing of the interior structure of massive celestial objects and insight into the very beginning of the universe. However, this will by no means signal the end of this work. Efforts will continue to push the sensitivity limits of detectors ever lower, at the same time widening our gravitational wave horizon to encompass sources at greater distances.

Soon, gravitational wave detectors are expected to be operating at the Standard Quantum Limit throughout much of their detection bandwidth. Novel techniques will need to be employed to probe beneath this level. One such method involves the use of optomechanical rigidity, or "optical springs". This technique couples the suspended optics of a Fabry-Pérot cavity together using only the radiation pressure force transferred between them, transforming the system into the harmonic oscillator regime and thus increasing sensitivity in a narrow band about the associated resonance. If this technique is to be applied to a large-scale gravitational wave detector, the effect must be well-characterised and robust control methods investigated. Importantly, the optical spring effect is observed to arise in any gravitational wave detector operating with high power, providing further motivation for the study of control systems to cope with them. It has further been suggested that multiple optical springs may be employed in a single system, either for improved optical stability such that electronic feedback may be reduced, or else to favourably re-shape the noise spectrum of the instrument for wider-band sensitivity improvement.

We present the design and commissioning of an experiment at the Glasgow 10 m Prototype Interferometer Laboratory to investigate characterisation and control methods for coupled optical spring systems. The experimental system consists of two 10 m cavities coupled mechanically by a shared end test mass. Each cavity can be detuned to facilitate opto-mechanical coupling, the combined effects of which are studied.

The design of this experiment is covered from initial simulations through mechanical de-

sign, testing and installation. Robust digital control loops are built, providing techniques for reliable lock-acquisition of high-finesse optical cavities and maintenance of stability in the presence of strong opto-mechanical effects.

A number of experiments are performed to examine the interactions between optical springs in both cavities and the control loops which maintain them. We observe marked power stability over significant periods of time to enable precision measurement of optical spring resonant responses in both cavities, and confirm the power-dependence of the optical spring effect. Techniques are described for the maintenance of stable cavity lock in the presence of strong optical springs and anti-springs, and complex responses resulting from couplings of multiple springs to loop gains are characterised.

This work will inform the design of future gravitational wave detectors, which are expected to employ optical spring technologies in order to push further beneath the Standard Quantum Limit.

CONTENTS

	Abs	tract	i		
	Con	tents	v		
	List	of Tables	vi		
	List	of Figures	х		
	Ack	nowledgements	xi		
	Ove	rview	civ		
1	Theoretical Foundations and Sources of Gravitational Radiation				
	1.1	Introduction	1		
	1.2	General Relativity and Gravitational Waves	2		
	1.3	Sources	4		
		1.3.1 Gravitational Collapse	4		
		1.3.2 Pulsars	5		
		1.3.3 Compact Binary System Inspiral	5		
		1.3.4 Stochastic Backgrounds	6		
2	Det	ection of Gravitational Waves	7		
	2.1	A Brief History of Gravitational Wave Detectors	7		
	2.2	Indirect Observation	8		
	2.3	Laser Interferometry	8		
	2.4	Sources of Noise	11		
	2.5	Infrastructure, Suspensions and Seismic Isolation	11		
	2.6	Optics and Thermal Noise	13		
	2.7	Laser Noise	14		
	2.8	Quantum Noise	15		
	2.9	Current Status and Future Upgrades	16		
3	The	eory and Experimental Concepts of Opto-Mechanical Rigidity	19		
	3.1	Operating at the Standard Quantum Limit	19		
		3.1.1 Simple Quantum Meters	19		

		3.1.2	Discrete Linear Position Measurement
		3.1.3	Continuous Measurement
		3.1.4	Standard Quantum Limit
		3.1.5	Free Mass Standard Quantum Limit
		3.1.6	Harmonic Oscillator SQL
	3.2	Radia	tion Pressure and Optical Cavity Dynamics
		3.2.1	Field and Power Equations for a Fabry-Pérot Cavity
		3.2.2	Optical Springs
		3.2.3	Frequency-Dependent Optical Spring
		3.2.4	Multiple Optical Springs
		3.2.5	Motivation for this Thesis
		3.2.6	Previous Applications
4	Cor	nceptu	al Experimental Design and Implementation 45
	4.1	Overv	iew \ldots \ldots \ldots \ldots \ldots \ldots 46
	4.2	The G	$Hasgow 10 m Prototype \dots 47$
	4.3	Model	lling and Design Study
	4.4	Laser	Amplifier Implementation
		4.4.1	Laser Amplification
		4.4.2	Amplifier Design
		4.4.3	Gain Measurements
	4.5	Seismi	ic Isolation and Suspension Design
		4.5.1	Dynamics of Suspended Optics
		4.5.2	Design of Suspensions and Pre-Isolators
		4.5.3	Suspension Design and Characterisation
		4.5.4	Actuation and Damping
		4.5.5	Local Control
		4.5.6	Global Control
	4.6	Install	lation
		4.6.1	Determination of CTM Wedge Angle
		4.6.2	Mode Matching
		4.6.3	Finesse Measurement
5	Dig	ital Co	ontrol Systems Design 80
	5.1	Digita	l Feedback and the CDS
		5.1.1	A Brief Introduction to Control Theory
		5.1.2	Digital Control
		5.1.3	The CDS

	5.2	The Po	ound-Drever-Hall Technique	. 88
	5.3	Contro	ol Systems Design and Implementation	92
		5.3.1	Frequency Stabilisation	92
		5.3.2	CDS Timing Delay and Servo Bandwidth	. 94
		5.3.3	Lock Acquisition by Finesse Transition	. 97
		5.3.4	Lock Acquisition and Maintenance by Widening of the Error Signal	98
		5.3.5	Servo Design	103
		5.3.6	Global Actuation Calibration	107
	5.4	Auto-A	Alignment of Cavity Optics	108
		5.4.1	The Sidles-Sigg Instability	108
		5.4.2	Auto-Alignment	111
6	Mul	tiple (Optical Springs in Coupled Cavities: Results	115
	6.1	Measu	rement Technique	116
	6.2	Precisi	on Optical Spring Measurements	119
	6.3	Single	Frequency Injections and Local Readout	123
	6.4	Power	Dependence	128
	6.5	Anti-S	prings and Loop Instability	129
	6.6	Dynan	nic Response of the Modified Error Signal	131
	6.7	Double	e Springs: Modification of Dynamics	132
	6.8	Double	e Springs: Simultaneous Measurement	136
		6.8.1	Coupling of Spring Effects into On-Resonance Cavity $\ldots \ldots \ldots$	137
		6.8.2	Suppression of CTM Motion	139
		6.8.3	Coupling of Loop Gain Features	140
		6.8.4	Coupling of Spring Strengths	141
	6.9	Summe	ed Response	141
	6.10	Conclu	isions	143
7	Con	clusior	ns and Future Work	144
A	Circ	uit Di	agrams	148
	A.1	Signal	Photodetector	148
	A.2	Local	Control Coil Driver	149
	A.3	Global	Control Coil Driver	150
Bi	bliog	raphy		151

LIST OF TABLES

3.1	Comparison of experimental parameters for a proposed "double optical	
	spring" Advanced LIGO configuration and the MOSES apparatus $\ . \ . \ .$	40
4.1	Updated system parameters after further modelling and long-term opera-	
	tion of the system	51
4.2	Distribution of rubber-stacks and approximate vertical resonant frequencies	
	in the two cavity-housing vacuum tanks	64
4.3	Mode frequencies in five degrees of freedom for each of the five cavity-	
	forming suspensions	65
4.4	Wire and suspended mass properties required to calculate suspension modes	66
5.1	Timing delay of the various input/output configurations of CDS	95

LIST OF FIGURES

1	Schematic of the Glasgow 10 m Prototype Laboratory vacuum system	XV
1.1	Effect of a gravitational wave of $+$ and \times polarisation on a ring of particles perpendicular to the direction of propagation	4
2.1	Diagram of a simple interferometer	9
2.2	Representative design for an advanced gravitational wave detector with	
	Fabry-Pérot arm cavities, set up for power and signal recycling	10
2.3	Advanced LIGO noise curves	17
3.1	Sketch of a linear measurement scheme	20
3.2	Total quantum noise power spectral densities for the free mass and har-	
	monic oscillator regimes	25
3.3	Schematic diagram of a simple two-mirror cavity	27
3.4	Simulated Airy function showing how intra-cavity power varies with tuning	28
3.5	Schematic diagram of a suspended two-mirror cavity acted upon by radia-	
	tion pressure	29
3.6	Simulated intra-cavity power and optical spring constant while sweeping	
	detuning	32
3.7	Variation of opto-mechanical restoring force and damping with detuning	
	and observation frequency	37
3.8	Graphical demonstration of the stable optical trap method	38
3.9	Sketch of the three-mirror coupled-cavity system proposed for this work	40
3.10	Simulation of quantum noise levels for a coupled-cavity double-optical-	
	spring system	41
4.1	Schematic of the Glasgow 10 m Prototype Laboratory vacuum system	47
4.2	Examples of contour plots used to pin down experimental parameters in	
	the early stages of the design study	49
4.3	Simple schematic diagram of the experimental setup	50
4.4	Sketch of a three-level lasing atomic system	52

4.5	Photo of amplifier crystal set up for double-pass gain	54
4.6	Optical setup for double-passing of the signal beam through the laser amplifier	55
4.7	Measured beam radii of pump and signal beams through the amplifier crystal	56
4.8	Gain measurements of the laser amplifier	57
4.9	Schematic diagram of a mirror suspended from four wires	59
4.10	Measurement of seismic noise in the laboratory, and simulated filtering	
	thereof by a two-stage pendulum	61
4.11	Sketch of the seismic isolation apparatus in each vacuum tank \ldots .	62
4.12	Tests of rubber stack resonant frequencies	63
4.13	Mechanical transfer functions of three degrees of freedom of the CTM sus-	
	pension	68
4.14	Pictures of suspension design and installation	69
4.15	Simulated force per unit current for a single coil-magnet actuator	70
4.16	Sketch of the optic positions in the Central Test Mass	74
4.17	Mode-matching solution for both cavities	76
4.18	Beam radius measurements for mode-matching purposes	77
4.19	Measurements of left- and right-hand cavity finesse from the reflected cavity	
	signals	79
51	Block diagram of a generic control system	81
5.2	Example of a CDS filter bank	01
		Xh
5.3	Portion of Simulink model describing signal routing in CDS	80 87
5.3 5.4	Portion of Simulink model describing signal routing in CDS	86 87
5.2 5.3 5.4	Portion of Simulink model describing signal routing in CDS	80 87 88
5.2 5.3 5.4	Portion of Simulink model describing signal routing in CDS	80 87 88 89
 5.2 5.3 5.4 5.5 5.6 	Portion of Simulink model describing signal routing in CDS	86 87 88 89 91
5.2 5.3 5.4 5.5 5.6 5.7	Portion of Simulink model describing signal routing in CDS	80 87 88 89 91
5.2 5.3 5.4 5.5 5.6 5.7	Portion of Simulink model describing signal routing in CDS	80 87 88 89 91 93
5.2 5.3 5.4 5.5 5.6 5.7 5.8	Portion of Simulink model describing signal routing in CDS	80 87 88 89 91 93
5.2 5.3 5.4 5.5 5.6 5.7 5.8	Portion of Simulink model describing signal routing in CDS	 80 87 88 89 91 93 94
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9	Portion of Simulink model describing signal routing in CDS	 86 87 88 89 91 93 94 96
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 5.10	Portion of Simulink model describing signal routing in CDS	80 87 88 89 91 93 94 96
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 5.10	Portion of Simulink model describing signal routing in CDS	 80 87 88 89 91 93 94 96 97
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 5.10 5.11	Portion of Simulink model describing signal routing in CDS	 86 87 88 89 91 93 94 96 97
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 5.10 5.11	Portion of Simulink model describing signal routing in CDS	 86 87 88 89 91 93 94 96 97 99
5.2 5.3 5.4 5.5 5.6 5.7 5.8 5.9 5.10 5.11 5.12	Portion of Simulink model describing signal routing in CDS	80 87 88 89 91 93 94 96 97 99

5.14	Block diagram of the main elements of the control loops	104
5.15	Block diagram of the RF setup	105
5.16	Simulation of open-loop gains external and internal to the CDS, and a	
	comparison of gain of the combined simulated and measured OLTF \ldots	106
5.17	Strength of the global control actuation, calibrated in metres per volt of	
	feedback signal	107
5.18	Illustration of the two coupled torsional modes of the Sigles-Sigg instability I	109
5.19	Observation of radiation-pressure-driven instabilities in the reflected DC	
	signals from each cavity	111
6.1	Example of frequency shaping applied to injected noise signal in CDS for	
	in-loop spring measurements	117
6.2	Coupled spring Optickle simulation presented as an open-loop transfer	
	function of the control loop	118
6.3	Open-loop transfer functions showing opto-mechanical resonances for var-	
	ious detunings of each cavity	121
6.4	Comparison between observed and theoretical DC spring constants	122
6.5	Injection signals and readout cavity displacement sensitivities for the "local	
	readout" method	124
6.6	Estimates of intra-cavity power obtained using the local readout setup	126
6.7	Estimates of intra-cavity power obtained by monitoring CTM position as	
	one cavity loses lock	127
6.8	Open-loop transfer functions showing opto-mechanical resonances for two	
	different input powers.	129
6.9	Effect of anti-springs on the open- and closed-loop gain of the system	130
6.10	Comparison of OLTF measurements and modified PDH signal response	100
0.11	showing increased phase margin at higher detunings	133
6.11	In-loop measurements showing modification of an optical spring in the	105
C 10	right-hand cavity by various detunings of the left-hand cavity	135
0.12	Observation of an optical spring in one cavity coupling into the measured	190
C 19	OLIF of the adjacent on-resonance cavity	138
0.13	Simultaneously-performed open-loop transfer functions showing apparent	190
611	Suppression of CTM motion as observed by both cavities	138
0.14	the closed loop response of one cavity into the open loop response of the	
	other	140
		140

to left-hand cavity
vity system 143
gnals \ldots \ldots \ldots 148
150

ACKNOWLEDGEMENTS

A lot of my friends have obtained what they call "real" jobs. Most of them even pay taxes. A great many people have helped me continue to do neither of these things over the last three and a half years, and for that I will be eternally grateful.

A great deal of the blame for all this ought to be borne by Ken Strain, Stefan Hild and Bryan Barr. I doubt I would have learned a fraction of the things I have without Ken's tireless supervision and seemingly boundless knowledge and interest, nor Stefan's enthusiasm and penchant for asking exactly the right awkward question. Bryan deserves special mention for the sheer force of will required to explain to me time and again every single concept in this thesis (and many more besides) without resorting to violence. I may regret putting this in writing, but he has certainly earned himself a pint.

Fellow denizens of Room 350, Angus Bell and Mariëlle van Veggel, have kept me entertained and informed in equal measure. Honorary mention goes to Borja Sorazu, for being only a wall away should should information or argument be required, and Sheena Barclay, for the all-too-temporary role of Room 350 Christmas Decorator.

The rest of the Interferometry Group, past and present, also deserves mention here, as I can't think of a single one who hasn't helped me at multiple points throughout the last few years. In no particular order: John Macarthur, Sebastian Steinlechner, Christian Gräf, Jennie Wright, Sean Leavey, Ally Houston, Jan Hennig, Sabina Huttner, Daniela Pascucci and Stefan Danilishin. My gratitude also to everybody else in the department who's helped or otherwise existed for reasons academic- to caffeine-based.

Thanks to all my friends inside and outside the University, for reminding me what free time is and how best to use it. Most of all, thanks to my parents and family, without whom I would certainly not have managed this, and to Taylor, without whom I would certainly have gone mad. This, I promise, will be my last Ph.D.

If you're reading this to see if your name is here, consider yourself Acknowledged for effort. I'd stop here if I were you.

DECLARATION

I hereby declare that the work presented in this thesis is the result of my own work, except where explicit reference is made to the work of others. This work has not been presented in any previous application for a degree at this or any other institution.

NEIL GORDON

Here is a lesson in creative writing. First rule: Do not use semicolons. All they do is show you've been to college.

- Kurt Vonnegut (with apologies)

OVERVIEW

The current and future generations of gravitational wave detectors are expected to be limited throughout the bulk of their bandwidth by quantum noise. This arises from statistical fluctuations on the source laser light used as the meter, and its interaction with the optics of the interferometric detector. For the particular case where these optics are suspended as free masses to reduce the impact of seismic noise, the "free-mass Standard Quantum Limit" (SQL) applies. A number of methods have been proposed for increasing the quantum sensitivity of the detector beyond this point. One such method involves transforming the dynamics of the suspended optics from free masses to harmonic oscillators, with the result that a narrowband increase in sensitivity is observed about the associated resonant frequency. This can be achieved by detuning the arm cavities of the detector from the resonant point, creating a dependence of intra-cavity power upon the positions of the optics, with a strong linear component; this is termed the *optical spring*.

A single optical spring is inherently unstable, requiring robust control systems to maintain stability of the opto-mechanical system. Further, the associated increase in sensitivity is of very narrow bandwidth, typically tens of Hz in frequency. Proposals exist to employ multiple optical springs in a single system, relatively tuned such that their combination is stable without requiring further complexity in the control loops. In addition, they can be used to further re-shape the quantum sensitivity of the instrument. In order that such a system may be realised in a large-scale interferometric gravitational wave detector, it is important that the practical aspects of its operation are well-understood.

In this thesis, design is presented for apparatus at the Glasgow 10 m prototype interferometer laboratory for investigating the interactions between multiple optical springs and their bounding control loops. Figure 1 shows a schematic of this apparatus installed in the laboratory, taking the form of a three-mirror coupled-cavity system capable of sustaining several optical springs. This thesis begins by providing the motivation for this investigation, then covers the design and implementation of the system and a number of prominent measurements performed with the apparatus.

In Chapter 1 gravitational radiation is briefly derived from the basics of General Relativity, and expected sources are characterised. In Chapter 2 laser interferometry for



Figure 1: Schematic of the Glasgow 10 m Prototype Laboratory vacuum system showing the main components of relevance to the work of this thesis. The upper 10 m arm contains the multiple spring experiment, while the lower arm contains a 10 m long reference cavity that is employed to stabilise the laser frequency.

the detection of gravitational waves is introduced, and more advanced detector systems discussed. The most prominent noise sources contributing to the sensitivity limit of a gravitational wave detector are described. In Chapter 3 the SQL and optical spring effect are both derived, specific motivation is given for the operation of a detector with two or more optical springs, and previous investigations into optical springs are discussed to place this work in context.

In Chapter 4 the design of the experiment is presented. This includes the initial simulations undertaken to inform key experimental parameters, and subsequent mechanical modelling, design, construction and testing of optic suspensions performed by the author. The implementation and testing of a laser amplifier to provide adequate input power for the experiment is described, and the optical system is characterised in terms of the mirror and cavity parameters.

In Chapter 5 the control systems are characterised. In particular, a novel technique is described for acquiring lock of a high-finesse cavity using digital feedback of limited speed. This technique also allows each cavity to be detuned in the creation of optical springs much further than is conventionally possible, allowing a wide range of optical spring strengths to be explored. This technique is shown to agree with simulations performed by the author, and its dynamic performance is characterised.

In Chapter 6 a number of measurements are presented of optical springs. These range from high-precision characterisation of single optical springs to show the potential of the measurement system to studies of the complex interactions between coupled optical springs and the control loops which maintain stability. These measurements will be of particular importance for large-scale gravitational wave detectors intending to operate with multiple optical springs.

The results presented in this thesis are indended to inform the implementation of a multiple-optical-spring system in an interferometric gravitational wave detector. Some current proposals for such systems are explained herein, and where possible the results throughout are explained in the context of these proposals.

Theoretical Foundations and Sources of Gravitational Radiation

1.1 INTRODUCTION

Τ

As early as 2016, a full century after Einstein published his General Theory of Relativity, a new door may be unlocked in our understanding of the universe. In developing his Special Theory of Relativity [1], the groundbreaking concepts of a finite speed of light in vacuum and equivalence of physical laws in any inertial reference frame were introduced. At the same time, however, it became clear that the existing theory of gravity, for centuries considered an instantaneous force between objects with mass, was lacking. In his General Theory [2], Einstein extended the framework of relativity to incorporate gravity as a geometrical property of the universe, a curvature of space-time in the presence of mass or energy with the result that the straight lines of particle trajectory in free space become curved *geodesics*. Further, Einstein predicted that perturbations to space-time in the form of accelerating mass would radiate energy in the form of *gravitational waves*. Advanced detectors around the world may be on the brink of the first direct detection of this gravitational radiation, opening up an exciting new facet to cosmology and a new means by which the properties of astronomical bodies and their interactions may be studied.

In this opening chapter a description of the generation of gravitational radiation is derived from Einstein's equations, and brief introductions to some of the expected sources of detectable gravitational waves given.

1.2 GENERAL RELATIVITY AND GRAVITATIONAL WAVES

The four-vector (x^0, x^1, x^2, x^3) describes position in generalised coordinates in spacetime. In Cartesian coordinates, the four-vector components are (t, x, y, z), where x, y and z are position coordinates and t time relative to some origin. The Einstein field equations in General Relativity are described by the Einstein tensor,

$$G^{\alpha\beta} = 8\pi T^{\alpha\beta} \,, \tag{1.1}$$

where the superscript denotes a tensor of order $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$, and $T^{\alpha\beta}$ is the stress-energy tensor, an expression of the energy and momentum flux and density in space-time; specifically, the flux of the α^{th} momentum component through a surface of constant coordinate x^{β} .

To derive the effect of gravitational radiation from General Relativity, a model of spacetime must be considered which is dynamic and subject to time-variance. To achieve this, we consider the simple case of "nearly flat" space-time, in which a coordinate system may be identified for which the *metric*, defining the distance between points in this system, has components according to

$$g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta} \,, \tag{1.2}$$

where $\eta_{\alpha\beta}$ is the *Minkowski metric*, commonly used in the context of Special Relativity to describe a flat space-time, and $h_{\alpha\beta}$ represents a small perturbation to this flat metric. For a suitable re-scaling of these metric perturbations, $\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \eta_{\mu\nu}h/2$, and application of the Lorentz gauge coordinate transform, Einstein's equations may be expressed in terms of the metric perturbations, $\bar{h}_{\mu\nu}$ [3]:

$$-\bar{h}^{,\alpha}_{\mu\nu,\alpha} = 16\pi T_{\mu\nu} \,. \tag{1.3}$$

In free space, the solutions to Eq. 1.3 take the form

$$\bar{h}^{,\alpha}_{\mu\nu,\alpha} \equiv \eta^{\alpha\alpha} \bar{h}_{\mu\nu,\alpha\alpha} = 0.$$
(1.4)

In the common notation of General Relativity, the "comma derivative" is taken to represent a partial derivative with respect to the coordinate following the comma; e.g.

$$f_{,\alpha\beta} = (f_{,\alpha})_{,\beta} = \frac{\partial}{\partial x^{\beta}} \left(\frac{\partial f}{\partial x^{\alpha}}\right) \,. \tag{1.5}$$

2

With this in mind, and noting that in the Minkowski metric $\eta^{00} = -1/c^2$, we find that the free-space solution to Einstein's equations in nearly-flat space-time takes the form of a wave equation, with the wave propagating at the speed of light:

$$\left(-\frac{\partial^2}{\partial t^2} + c^2 \nabla^2\right) \bar{h}_{\mu\nu} = 0.$$
(1.6)

The simplest solution to Eq. 1.6 is that of a plane wave, of form

$$\bar{h}_{\mu\nu} = \operatorname{Re}\left(A_{\mu\nu}e^{ik_{\alpha}x^{\alpha}}\right) \tag{1.7}$$

for wave amplitude $A_{\mu\nu}$ and vector k_{α} . Performing another convenient gauge transformation, termed the *transverse-traceless gauge*, we find that wave amplitudes $A_{\alpha z} = 0$ (i.e. the wave contains only components in the transverse coordinates) and $A_{\alpha t} = 0$ (and so $A_{\alpha\alpha} = -A_{\alpha\alpha}$; hence "traceless") for all α . Thus the wave amplitude components in the transverse-traceless gauge reduce to

$$A_{\mu\nu}^{TT} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$
 (1.8)

and so from Eq. 1.7 we find that

$$\bar{h}_{\mu\nu}^{TT} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \bar{h}_{xx} & \bar{h}_{xy} & 0 \\ 0 & \bar{h}_{xy} & -\bar{h}_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} .$$
(1.9)

Thus, we may express the wave solution in terms of A_{xx} and A_{xy} only:

$$h_{xx}^{TT} = A_{xx} \cos\left(\omega(t-z)\right), \qquad (1.10)$$

$$h_{xy}^{TT} = A_{xy} \cos\left(\omega(t-z)\right), \qquad (1.11)$$

and so we identify two distinct polarisation states for gravitational waves: "+ polarisation", for which $A_{xx} \neq 0$ and $A_{xy} = 0$, and "× polarisation", for which $A_{xx} = 0$ and $A_{xy} \neq 0$. The effect of each of these polarisation states on a ring of particles transverse to the direction of propagation is presented in Figure 1.1.



Figure 1.1: Exaggerated effect of a gravitational wave of + (top) and \times (bottom) polarisation on a ring of particles arranged in the transverse plane.

1.3 Sources

Thanks to the relative weakness of gravitational radiation, detectable sources must be massive astronomical bodies or cosmological in nature. Sathyaprakash and Schutz [4] make the comparison with masses of order 10^3 kg, separated by 10 m and rotating to produce gravitational waves of frequency 20 Hz. With a detector situated on the opposite side of the Earth, approximately one wavelength away, they estimate a strain amplitude of $h \simeq 5 \times 10^{-43}$. Clearly this is undetectable. However, when one considers for example supermassive black holes with masses many thousands that of our Sun, or pulsars with surface velocities up to a non-negligible fraction of the speed of light, the predicted strain amplitudes become more realistic for detection on Earth. Here we provide a very broad outline of the types of sources of gravitational waves which we may be able to probe in the near future.

Typically, sources of gravitational radiation are broadly characterised by their strain amplitude, event rate and frequency. These are estimated where possible.

1.3.1 GRAVITATIONAL COLLAPSE

The collapse of a highly-evolved stellar core or of a white dwarf that has accreted mass to reach or exceed the Chandrasekhar limit [5] can result in the formation of a black hole or neutron star. Any inhomogeneity in the collapse, caused for instance by a strong angular velocity component, could result in some angular momentum and energy being radiated gravitationally. These are often accompanied by the production of various supernovae and gamma-ray bursts. Observations of these "electromagnetic counterparts" using traditional astronomical methods can in future be used to pinpoint the angular position of a detected gravitational wave source upon the sky; further, they provide a rough estimate of the event rate of such processes. Type II supernovae are thought to occur with a rate of 0.1 - 0.01 per annum in an equivalent galaxy to our own. Extrapolating to the local group, a very rough rate of 30 per year can be predicted, with estimated frequency range $\sim 100-1000$ Hz. The strain of such an event, when compared to expected values for mass energy, frequency, etc. in our own galaxy, may be estimated as [4]

$$h \sim 6 \times 10^{-21} \left(\frac{E}{10^{-7} M_{\odot}}\right)^{1/2} \left(\frac{1 \,\mathrm{ms}}{T}\right)^{1/2} \left(\frac{1 \,\mathrm{kHz}}{f}\right) \left(\frac{10 \,\mathrm{kpc}}{r}\right) \,.$$
 (1.12)

1.3.2 Pulsars

Some sources may be considered continuous. This would include, for example, a spinning neutron star with a deformation upon its surface, resulting in asymmetric mass distribution and so loss of energy to gravitational waves. Through extended observation period, limits can be placed on the potential asymmetry in known pulsars; Initial LIGO has achieved this for the Crab and Vela pulsars [6]. Strain estimates are as follows:

$$h \sim 10^{-22} \left(\frac{f}{100 \,\mathrm{Hz}}\right)^2 \left(\frac{1 \,\mathrm{kpc}}{r}\right) \left(\frac{\epsilon}{10^{-6}}\right) \,, \tag{1.13}$$

where ϵ represents the fractional asymmetry resulting from the surface deformation.

1.3.3 Compact Binary System Inspiral

Binary star systems present a further potential source of gravitational waves, provided that they exist in tight enough orbit and are of appropriate mass to collapse to a compact object in a scenario termed "compact binary coalescence", or CBC. This is typically estimated for pairs of neutron stars or black holes, or combinations of the two. As the binary system approaches coalescence, by emission of gravitational radiation resulting in a degradation of orbit, the amplitude and frequency of waves increases, forming what is known as a "chirp" signal that ends either when the pair merge or, if particularly compact objects, their orbits reduce to beneath the "last stable orbit". For a binary neutron star system, gravitational wave strain can be estimated to be [7]

$$h \sim 1.7 \times 10^{-23} \left(\frac{15 \,\mathrm{Mpc}}{r}\right) \left(\frac{1 \,\mathrm{day}}{\tau}\right)^{1/4} \left(\frac{M}{1.4 \,\mathrm{M_{\odot}}}\right)^{5/4}$$
. (1.14)

1.3.4 Stochastic Backgrounds

Further to these specific, isolated sources of gravitational radiation, a very weak random background field is predicted to exist with its genesis as early as 10^{-30} s after the Big Bang – this can be thought of as a primordial gravitational wave counterpart to the Cosmic Microwave Background Radiation of the electromagnetic spectrum. It is thought that identification and study of this stochastic field could provide further insight into the laws of physics at the incredibly high energies of the Big Bang and the process of cosmic inflation. Since the field in this case is random, the stochastic background may be expressed in terms of a statistical energy density, $S_{gw}(f)$. This is then expressed as a fraction of the cosmological density predicted by the Hubble constant ($\rho_c = 3H_0^2/8\pi$):

$$\Omega_{gw}(f) = \frac{10\pi^2}{3H_0^2} f^3 S_{gw}(f) , \qquad (1.15)$$

where existing models of the Big Bang and observations limit Ω_{gw} to $< 10^{-5}$.

II DETECTION OF GRAVITATIONAL WAVES

2.1 A Brief History of Gravitational Wave Detectors

The first attempts to directly observe the strain effect of gravitational radiation were performed by Joseph Weber in the 1960s, by pursuing the excitation of normal modes of the Earth predicted by gravitational waves of low frequency [8]. The data available for these studies proved sufficient only for setting an approximate upper limit on the influence of gravitational radiation. Accordingly, studies progressed to examine strain effects in aluminium bars. Over subsequent studies, these bar detectors became better isolated from local seismic and acoustic effects and cryogenically cooled to reduce thermal fluctuations. Although there were indications that Weber's original bar detectors may have shown coincident tidal excitations in keeping with a detection [9], these are viewed as unlikely to be attributable to gravitational waves as their sensitivity was several orders of magnitude below the predicted detection level of today's models, and they were ultimately not reproducible in subsequent bar experiments [10].

Over the last thirty years, laser interferometry has overtaken the resonant bar detectors as the preferred detection method. This is largely due to their peak sensitivity having now exceeded those of the cryogenic bar detectors by several orders of magnitude, as well as being sensitive to a much wider range of frequencies. The first generation of large-scale interferometric gravitational wave detectors reached design sensitivities around 2005 [11]. These are now coming to the conclusion of several years of upgrades to the second generation, some of which are expected to begin data-taking within the next year. The design of such detectors is covered in detail in Section 2.3.

2.2 INDIRECT OBSERVATION

The Hulse-Taylor binary pulsar is doubtless the most famous physical manifestation of gravitational radiation. Identified by Hulse and Taylor in 1975 [12], PSR B1913₊16 is a system containing two neutron stars, one a pulsar acting as a regular frequency by which the system may be fully characterised in terms of masses and orbits. With the system understood in terms of these parameters, observations taken over almost twenty years of the shift in orbital phase as energy is lost to gravitational radiation match very well with the predictions of General Relativity [13]. Similar systems have been identified in the years since, with one example consisting of two orbiting pulsars soon expected to overtake the Hulse-Taylor result as the most accurate confirmation of gravitational radiation from General Relativity [14].

2.3 LASER INTERFEROMETRY

For *direct* detection of gravitational radiation, laser interferometry has become the preferred method. Interferometry uses light as a meter, firstly by splitting a laser beam by a beamsplitter into two orthogonal paths; both beams travel some distance, often referred to as the *arm length* of the interferometer, reflect from an optic and return to recombine at the beamsplitter. Any differential phase shift accumulated between the two beams by the point of recombination is then manifest at the output port. This operation makes the interferometer an ideal probe for gravitational radiation. If the two arms of the interferometer are optimally aligned with the gravitational wave polarisation, it is optimally sensitive to the strain effects shown in Figure 1.1. Figure 2.1 shows a simple sketch of an interferometer set up in this fashion.

For triangulation of the source position on the sky, a network of detectors spread around the world provides optimal coverage. There are currently four large-scale detectors across the world, in operation since the late 1990s and reaching design sensitivity at various points in the last two decades. These are the two 4 km arm length LIGO detectors in the USA, situated in Louisiana and Washington [15], the 600 m GEO 600 detector in Hannover, Germany [16] and the 3 km VIRGO detector in Pisa, Italy [17]. Various upgrades to these are on-going, with both Advanced LIGO [18] detectors expected to begin data-taking this year, and Advanced VIRGO [19] to follow soon after. These advanced detectors are expected to reach a factor of ten increase in sensitivity over the first-generation. The completed GEO-HF [20] upgrade has been operating in "astro-watch" for the last few years, ensuring that one detector remains operational while the rest are improved. Also



Figure 2.1: Diagram of a simple interferometer. Light from the source laser is split via a beamsplitter. Each beam then travels some distance and reflects from a mirror before the two are recombined back at the beamsplitter and the resulting signal detected at a photodiode.

under construction is KAGRA, a new facility situated in the Kamioka mine in Japan [21].

Since the gravitational wave is measured as a *strain* between two masses, the detected signal will increase linearly with interferometer arm length for systems in which the light travel time is less than a half-period of the wave in question. Since ground-based detectors are fundamentally length-limited by a great many factors, from very large beam sizes requiring even larger optics to financial, infrastructural and geographical concerns, it is impractical to build simple interferometers as in Figure 2.1 of ever-increasing size. The resultant accumulation of phase with distance can alternatively be achieved by employing a *Fabry-Pérot cavity* in each arm. Inside such an optical cavity each photon circulates many times, a number set by the parameters of the two bounding optics, thus accruing extra phase shift as a result of the passing wave. Figure 2.2 shows an updated sketch of a basic interferometer to include these arm cavities, achieved by the inclusion of a partially-transmissive optic in the beam path following the beamsplitter. For a single cavity, this optic is often termed the *Input Test Mass* (ITM), and that at the end of the arm the *End Test Mass* (ETM).

Although the largest fringe gradient in a simple interferometer is produced when half of the light is directed to the photodetector, it is much more efficient to operate the interferometer at "dark port", whereby destructive interference is observed at the output



Figure 2.2: Representative design for an advanced gravitational wave detector. In contrast with the simple interferometer (Figure 2.1), this design incorporates Fabry-Pérot arm cavities, both of length L, between ITM and ETM in the x and y arm. Also included are the Power Recycling Mirror (PRM) and Signal Recycling Mirror (SRM).

for normal operation, and constructive interference indicates a differential arm length change. In particular, this avoids saturation of the signal photodetectors at the output and improves the dynamic range for signal detection. Supplementary optics may then be included in the system to take advantage of the light at each port; this is illustrated by the two extra optics shown in Figure 2.2. These are the Power Recycling Mirror (PRM) and Signal Recycling Mirror (SRM). Both of these are used to create supplementary cavities with the interferometer arms. In the case of the PRM, light that would ordinarily be reflected back towards the laser when the arms are tuned for dark port will instead get recycled back towards the cavities. This serves to increase the optical power in each cavity without requiring an increase in power at the source laser, in turn improving the shot noise of the detector by providing more photons. As we shall see, however, the benefits of higher power in the arm cavities must always be weighed against the corresponding increase in radiation pressure noise resulting from more photons imparting momentum to each optic.

The purpose of the SRM is to enhance at particular frequencies signals that exit the interferometer through the dark port. Precise positioning of the SRM and choice of optical transmission may be used to tune the signal response of the interferometer output in such a way as to resonantly enhance signals at particular frequencies. Once again, this must be a trade-off; a high-reflectivity SRM will result in greater enhancement of the

output signal by way of recycling each photon for longer, but also means a higher-quality factor, and hence diminished bandwidth, resonant response.

2.4 Sources of Noise

When constructing a laser interferometer for gravitational wave detection, a number of different noise sources must be considered. In order to improve the sensitivity of the instrument to the desired level, each must be reduced. Without applying any reduction techniques, the noise is generally dominated by seismic noise at the lower-frequency end of the spectrum, and laser noise at the higher end. Beneath these, but still of highmagnitude, are such sources as acoustic noise, thermal noise in the optics and quantum noise on the light itself.

Over the course of Sections 2.5 to 2.8 we examine a number of the most prominent of the noise sources that contribute to the overall sensitivity limit of a large-scale interferometer. We begin in Section 2.5 with a discussion of the infrastructure requirements of such an instrument with regard to vacuum systems for mitigation of acoustic and scattered light noise and the suspension of optics as pendulums to mitigate seismic noise. This is followed in Section 2.6 by an introduction to the sources of thermal noise in each optic and suspension. Section 2.7 considers the frequency and amplitude noise imposed upon the laser light by the conditions of the source, and Section 2.8 examines the fundamental statistical noise introduced by the use of photons as the meter. Figure 2.3 then presents expected sensitivity limits for each of these noise sources in the context of the Advanced LIGO design.

2.5 INFRASTRUCTURE, SUSPENSIONS AND SEISMIC ISOLA-TION

Careful choice of site and design of vacuum systems are vital if an advanced large-scale interferometer is to be operated at the required sensitivity to detect gravitational waves. As has already been stated, long arms composed of Fabry-Pérot cavities are necessary to improve the potential signal size obtained from an observation. Once an appropriate site for such apparatus has been located, a great many sources of noise must be considered and accounted for if the instrument is to achieve sensitivities necessary for signal detection.

The entire arm and main optical components of a ground-based detector must be com-

pletely enclosed in an ultra-high vacuum system. This serves to reduce the phase noise resulting from light scattering from residual gas molecules, as well as acoustically isolating the optics from the environment. In Initial LIGO, the vacuum system was designed to reach pressures of $10^{-8} - 10^{-9}$ Torr for this purpose, throughout both 1.2 m diameter, 4 km long tubes enclosing the arms. The diameter of vacuum tube is set by scattered-light constraints, and may also be related to the need for multiple beams in a single tube [22].

While the low-pressure environment may isolate the cavity and supplementary optics from acoustic noise, they remain subject to local Newtonian gravitational couplings and seismic motions. Seismic noise is dominant below a few tens of Hz in the large-scale detectors, and below approximately 10 Hz local fluctuations in density from either artificial or environmental factors are predicted to result in perturbations to the local gravitational field from which the optics cannot be shielded. This imposes the low-frequency limit on detector sensitivity, although techniques are being developed for evaluation of the local field fluctuations with arrays of seismic instruments and correction of the data stream from the instrument accordingly, or active "feed-forward" to the optics for cancellation purposes [23]. Some upcoming detector designs, such as KAGRA and the proposed third-generation detector Einstein Telescope [24], will be built in underground tunnels to mitigate the effect of surface seismic waves, which fall off exponentially with depth.

Above 10 Hz, seismic noise is predominantly a result of local man-made sources. The influence of these effects upon the interferometer is reduced by suspending the optics as multiple-stage pendulums. If we consider a simple pendulum formed by a rigid mass, m, suspended by an ideal elastic wire of stiffness k, length ℓ and damping constant b, the transfer function to mass motion, x_m , from ground motion, x_g , may be expressed as [25]

$$\frac{x_m}{x_g} = \frac{\omega_0^2}{\sqrt{(\omega_0^2 - \omega^2)^2 + \omega^2 \gamma^2}},$$
(2.1)

where $\omega_0 = \sqrt{k/m} = \sqrt{g/\ell}$ is the fundamental resonant frequency of the pendulum and $\gamma = b/m$. At low frequencies, this reduces to unity; at higher frequencies, and for low levels of damping, it reduces to ω_0^2/ω^2 . Thus a single pendulum stage may be considered to provide $1/(\text{frequency})^2$ isolation from ground motion above its fundamental resonant frequency. Each subsequent stage will provide the same factor, with the result that the advanced suspension systems, typically incorporating several stages of pendulum isolation, as well as inverted pendulum, blade spring and active damping components, attain significant seismic isolation above the Newtonian noise limit.

2.6 Optics and Thermal Noise

The suspended optics in the current generation of detectors are formed of fused silica substrates coated with multiple dielectric stack layers to facilitate Bragg reflection. The substrate, coating and further suspension parts are subject to thermal fluctuations, which manifest in the interferometer output as phase (equivalently, displacement) noise. These fluctuations arise from mechanical modes of the suspension, internal modes of the masses, violin modes and deformations in the suspension wires and surface or bulk imperfections in the optic coatings and substrates. We briefly examine the largest contributors of these in this section.

Typically, silica fibres are used to suspend the lowest suspension stage (i.e. the test mass). Suspension thermal noise exists as a result either of Brownian noise, resulting in mechanical dissipation in the various suspension materials, or of thermoelastic noise. The Fluctuation-Dissipation theorem expresses the equivalent displacement noise from these effects above the fundamental resonant frequency of the suspension as [26]

$$x_{\text{suspension}}^2 = \frac{4k_B T \omega_0^2 \phi_{\text{total}}}{m \omega^5} \,, \tag{2.2}$$

where T is temperature, ϕ_{total} the mechanical loss of the suspension, m the suspended mass, ω_0 the resonant frequency of the mode in question, and ω the observation frequency. The mechanical loss is proportional to the inverse of the quality factor, Q, of the suspension mode being studied. The total mechanical loss of the monolithic fused silica lower stage can be separated into the contributions from surface and thermoelastic loss.

Surface losses from the suspended fused silica substrate and bonded fibres arise from imperfections in the material surface [27, 28], while thermoelastic noise arises from deformation (e.g. bending of suspension fibres near the break-off point at the test mass), leading to local heating and dissipation across the resulting temperature gradient [29]. The losses from these two effects can be expressed as [25, 30]:

$$\phi_{\text{surface}} = \frac{4d_s\phi_s}{r}, \qquad (2.3)$$

$$\phi_{\text{thermoelastic}} = \frac{YT}{\rho C} \left(\alpha - \frac{\sigma\beta}{Y} \right)^2 \left(\frac{\omega\tau}{1 + (\omega\tau)^2} \right) ,$$
 (2.4)

where d_s is the depth over which surface losses are believed to act, ϕ_s is the mechanical loss of the material surface, r is the radius of the fibre, Y is the Young's modulus, C the specific heat capacity, ρ the density, α the coefficient of thermal expansion, σ the stress placed upon the fibre by the suspended mass and $\beta = (dY/dT)/Y$ is the thermoelastic coefficient.

The thermal noise in the multilayer optic coatings may be reduced by careful choice of coating material, and by using a large beam diameter at each optic to average the contributions across a larger optic surface. The coating Brownian noise contribution to displacement noise may be expressed as [31]

$$x_{\text{coating}}^2 = \frac{2k_B T \phi_{\text{eff}}(1-\sigma)}{\pi^{3/2} f \omega_m Y}, \qquad (2.5)$$

where we have defined

$$\phi_{\text{eff}} = \phi + \frac{d}{\omega_m \sqrt{\pi}} \left(\frac{Y}{Y_\perp} \phi_\perp + \frac{Y_\parallel}{Y \phi_\parallel} \right) \,, \tag{2.6}$$

where ϕ is the mechanical loss, and Y again the Young's modulus, of the substrate. The terms have been separated into their appropriate perpendicular and parallel components within the coatings.

2.7 LASER NOISE

Noise related to the laser source can be separated into contributions from amplitude and frequency noise on the light itself, and fluctuations in the geometry (i.e. shape and directionality) of the beam.

Amplitude fluctuations in the source laser will result in fluctuations in the radiation pressure on the suspended optics, and thus in noise on the relative optic positions that couples directly into noise on any detected signal at the interferometer output [32]. Further, if not operated at the dark fringe this amplitude noise will couple directly into the interferometer output; this is mitigated in part by operation at the dark fringe, but any fluctuations from this condition continue to couple noise to the output. This results in a strain sensitivity limit imposed according to

$$h_{\text{Laser amp.}} = \frac{\Delta P}{P} \frac{\delta L}{L} \,, \tag{2.7}$$

where L is the length of the interferometer arm, δL the offset from dark fringe and $\Delta P/P$ the fractional noise on the laser power. Typically, this noise can be reduced by feeding back a signal to the laser to compensate power fluctuations.

Variations in the length of the laser cavity or refractive index of the laser medium manifest as frequency noise upon the output laser field. This will couple to the output photodetector of the instrument due to offsets from the ideal relative effective length of each arm of the interferometer, resulting in a limit to the strain sensitivity of

$$h_{\text{Laser f.}} = \frac{\Delta f}{f} \frac{\delta L}{L} \,, \tag{2.8}$$

where $\Delta f/f$ is the fractional frequency noise on the laser light of frequency f. This noise can be reduced by stabilising the laser frequency to some well-defined frequency reference (e.g. by using the Pound-Drever-Hall technique, covered in Section 5.2). Feedback to the laser is typically applied in different ways depending on the frequency range – for example, an NPRO laser may be frequency-stabilised by feeding back to the crystal temperature at low frequencies (typically < 10 Hz) and to a piezo-transducer to squeeze the crystal at higher frequencies. Further details can be found in Section 5.3.1.

Fluctuations in beam pointing can arise from various sources; commonly these are seismic and acoustic effects at the optical bench upon which the laser source rests, as well as effects related to processes intrinsic to the laser [10]. This can result in higher-order laser modes coupling into the interferometer output. It is possible to reduce the influence of this effect by using optical elements which reduce or remove higher-order beam couplings, such as mode-cleaning cavities [33, 34] or single-mode optical fibres.

2.8 QUANTUM NOISE

Quantum noise arises as a consequence of using light as the meter in a gravitational wave detector. Thanks to the quantum nature of photons, two distinct effects are observed, the combination of which comprises a fundamental (although malleable) limit to the sensitivity of an interferometric detector in which the optics are suspended as free masses.

Photon counting statistics (i.e. the statistical uncertainty in the number of photons detected at a particular time) will lead to *shot noise* at any photodiode. This statistical noise thus contributes at the readout of the interferometer, and is inversely proportional to the square root of the optical power incident upon the photodiode. As a result of the same Poisson distribution of photon population with time that generates shot noise, at any optic with which the light field interacts there will exist an uncertainty in the photon momentum imparted to that optic. This results in a fluctuating force on the suspended optic, giving rise to *radiation pressure noise*. This is commonly referred to as a *back-action noise*, as it arises as a consequence of the measurement itself.

Noting that an advanced interferometer design (similar to that sketched in Figure 2.2) can be described by an equivalent Fabry-Pérot cavity, as determined by Mizuno [35] and Buonnano and Chen [36], the contribution of each of these noise sources to strain sensitivity at frequency f can be expressed in terms of common parameters for such an equivalent cavity as [37]

$$h_{\rm SN}(f) = \frac{1}{L} \sqrt{\frac{\hbar c \lambda}{2\pi P_C}}, \qquad (2.9)$$

$$h_{\rm RP}(f) = \frac{1}{\mu f^2 L} \sqrt{\frac{\hbar P_C}{2\pi^3 c \lambda}}, \qquad (2.10)$$

where P_C is the intra-cavity power and μ the reduced mass of the cavity optics. Immediately obvious is the fact that increasing the cavity length improves the sensitivity of both. However, once a cavity length and source laser wavelength are set, shot noise is only reduced by an increase in optical power, for which radiation pressure noise increases equivalently, and vice versa. Increasing optic masses will reduce radiation pressure noise, but as for an increase in cavity length this is only an option until infrastructural or manufacturing limitations make it unreasonable. The limit imposed upon a free-mass gravitational wave detector, or an equivalent Fabry-Pérot cavity, by the interplay between these two effects is termed the *Standard Quantum Limit* (SQL), which is derived in detail for a generic quantum measurement scheme in Chapter 3.

Figure 2.3 presents the expected noise budget for Advanced LIGO [18], including the contributions from each prominent noise source discussed in this chapter.

2.9 CURRENT STATUS AND FUTURE UPGRADES

As noted previously, Advanced LIGO is expected to begin data-taking in the near future, having in the first few months of the year successfully locked both arm cavities in both detectors for a number of hours each. Both Advanced LIGO interferometers incorporate power and signal recycling techniques and multi-stage suspended optics with monolithic fused silica lower stages, with the thermal noise level expected to be reduced



AdvLIGO Noise Curve

Figure 2.3: Expected sensitivity plots for Advanced LIGO indicating contributions from prominent noise sources including those described in Sections 2.5 to 2.8 and combined sensitivity level. This plot was produced using GWINC, software written to simulate noise performance in advanced interferometers.

to beneath the Standard Quantum Limit, such that quantum noise becomes the limiting noise source across the bulk of the detection bandwidth (roughly 10 Hz to 1 kHz) [18]. In order that sensitivities beneath the SQL be obtained, thus increasing the volume of space that can be probed by each detector, new techniques must be applied or topologies designed in order to transform or surpass this limit. These technologies broadly combine under the label of "quantum non-demolition" (QND) techniques.

Optical springs, the main focus of this thesis, are one example of an effect the exploitation of which has been proposed for transforming the standard quantum limit. As we shall see, optical springs in a Fabry-Pérot cavity have the potential to rigidly couple the cavity optics together using the radiation pressure transfer between them, in such a way as to transform the dynamics of each suspended test mass from free masses to harmonic
oscillators, for which a different form of quantum limit appplies. This effect is derived and various applications are explored in Chapter 3.

Some proposals favour using a completely different approach that isn't subject to the problems associated with the position measurements of Michelson interferometers. The "speed-meter" style detector chooses as its measurement observable the test mass *veloc-ity*, as opposed to position. Since velocity commutes with time, it can be measured in theory to arbitrary precision without violation of the Uncertainty Principle. This method has been shown to have the potential to cancel radiation pressure noise at lower frequencies in the SQL [38], but is some distance from application in a large-scale detector. A proof-of-principle experiment is currently being installed in the Glasgow 10 m prototype laboratory [39].

Beyond these techniques, designs are already being proposed for the third generation of gravitational wave detectors. Design is underway for the Einstein Telescope, a conceptual large-scale detector aiming for a factor ten increase in sensitivity over the second generation [24]. Current designs aim to achieve this ambitious aim by constructing the entire detector some distance underground to mitigate seismic noise, and by incorporating larger beam and optic sizes, higher laser power, higher-order suspension stages and improved optic coatings and cryogenics. It is likely that such a detector would employ advanced optical techniques, such as some of the QND methods covered briefly here and later in this thesis, or operating in "xylophone" format whereby two detectors run in parallel; one optimised for low-frequency sensitivity (e.g. low-power beams and heavier masses to improve radiation pressure noise at the expense of shot noise) and one for highfrequency sensitivity (e.g. high-power beams, amidst other optimisations).

Further proposals exist for the operation of space-based detectors, for which low-frequency Newtonian and seismic noise sources that impose the lower-frequency limit for groundbased detectors wouldn't be problematic. Since infrastructure issues are redundant beyond the housing of optics and sensing and control technologies, such detectors can be operated with arm lengths of millions of kilometres. The Laser Interferometric Space Antenna, or eLISA [40], aims to achieve this at some point over the next few decades. The technological challenges of such an instrument are significant; for this reason, the LISA Pathfinder satellite [41] is expected to be launched in 2015 for testing and proof-ofprinciple experiments.

III Theory and Experimental Concepts of Opto-Mechanical Rigidity

3.1 Operating at the Standard Quantum Limit

An interferometric gravitational wave detector is set up to detect the phase modulations imposed upon the circulating laser light as a passing gravitational wave differentially alters the relative lengths of the two interferometer arms. Therefore, the measurement of phase fluctuations must be as sensitive as possible. We show that quantum measurement theory predicts the existence of a fundamental limit to this sensitivity, a combination of the quantum fluctuations on the measured signal, i.e. *shot noise* on the detected photons at the output port, and the back-action noise imposed upon the detector by the measurement, i.e. *radiation pressure* from the probing photons interacting with each mirror.

3.1.1 SIMPLE QUANTUM METERS

Consider a quantum measurement as carried out by the linear measurement scheme of Figure 3.1. Quantum measurement theory predicts a fundamental limit to this sensitivity that depends on the dynamics of the detector "probes" – the interferometer test-masses – being measured. To begin with, we will consider linear measurement of a quantum system with particular relevance to gravitational wave detection techniques. This discussion follows much of the quantum measurement framework and terminology laid out by Danilishin and Khalili [37] and Chen [42].



Figure 3.1: Linear measurement scheme, where G is a classical force acting on the probe, \hat{F} is the back-action noise imposed upon the probe by the meter, \hat{O} is the readout observable of the meter, \hat{x} is the displacement of the probe and Δx the measurement noise [43, 37].

3.1.2 DISCRETE LINEAR POSITION MEASUREMENT

Consider a simple system in which we start out by representing a continuous laser beam as a series of discrete pulses. These are used to measure the displacement, x, of a mirror, M. For particular application to a gravitational wave detector with arm cavities, we will suppose that the incident laser light, which probes the mirror position periodically with time interval τ , reflects from the test mass a number of times $\mathcal{F} > 1$ before being measured. We may thus use this system as an analogue for the light probing the position of a mirror at the end of a Fabry-Pérot cavity in one arm of a gravitational wave detector. Each pulse will experience a phase-shift proportional to the position of the mass upon reflection [44, 37], expressed in terms of the appropriate quantum mechanical operators as

$$\hat{\phi}_j^{\text{refl}} = \hat{\phi}_j - 2\mathcal{F}k_0 \hat{x}(t_j) \,, \tag{3.1}$$

for pulse j, where $k_0 = \omega_0/c$ and ω_0 is the laser frequency, and $\hat{\phi}_j$ is the initial phase. Each j^{th} pulse is measured by a phase detector, assumed to introduce a measurement error negligible in comparison to the root-mean-square uncertainty of all phases, $\Delta \phi = \sqrt{\langle \hat{\phi}^2 \rangle - \langle \hat{\phi} \rangle^2}$.

If we assume this uncertainty $\Delta \phi$ to be the only error in the measurement, then the position measurement uncertainty is

$$\Delta x_{\text{measured}} = \frac{\Delta \phi}{2\mathcal{F}k_0} \,. \tag{3.2}$$

This then allows us to infer the independent random position values of the mirror with RMS uncertainties $\Delta x_{\text{measured}}$ to be

$$\hat{x}_{\rm fl} = -\frac{\hat{\phi}_j}{2\mathcal{F}k_0}\,.\tag{3.3}$$

20

The equivalent displacement of the mirror itself may be expressed as

$$\hat{x}_j = -\frac{\hat{\phi}_j^{\text{refl}}}{2\mathcal{F}k_0} = \hat{x}(t_j) + \hat{x}_{\text{fl}}(t_j) \,. \tag{3.4}$$

Now we may introduce the *back-action* noise – this arises from the momentum transferred to the mirror by each pulse as it probes the position. We express this back-action momentum as $a \tau$

$$\hat{p}_j^{\text{BA}} = \hat{p}_j^{\text{after}} - \hat{p}_j^{\text{before}} = \frac{2\mathcal{F}}{c} \hat{W}_j \,, \tag{3.5}$$

where \hat{W}_j is the energy of the j^{th} pulse. The main contribution to this back-action noise is the classical radiation pressure, obtained by replacing \hat{W}_j in Eq. 3.5 with the mean pulse energy, W. This component can be compensated easily during or following the measurement, since it may be considered a constant pressure; the small part that cannot be compensated depends on the fluctuation to the mean pulse energy:

$$\hat{p}^{BA}(t_j) = \frac{2\mathcal{F}}{c}(\hat{W}_j - W), \qquad (3.6)$$

with RMS uncertainty proportional to RMS pulse energy uncertainty, ΔW :

$$\Delta p_{\rm BA} = \frac{2\mathcal{F}\Delta W}{c} \,. \tag{3.7}$$

Since the energy and phase of each pulse are canonically-conjugate observables, they obey the Heisenberg uncertainty relation, such that

$$\Delta W \Delta \phi \ge \frac{\hbar \omega_0}{2} \,. \tag{3.8}$$

Rearranging Eqs. 3.2 and 3.7 for $\Delta \phi$ and ΔW respectively, then substituting into Eq. 3.8, we find that the position measurement uncertainty and momentum perturbation due to back-action of the measurement also follow the uncertainty principle:

$$\Delta x_{\text{measured}} \Delta p_{\text{BA}} \ge \frac{\hbar}{2}$$
 (3.9)

3.1.3 Continuous Measurement

To extend the conclusions of Section 3.1.2 to a continuous measurement of mirror position, we may simply choose some time interval T that is both of sufficient length to comprise a large number of total measured pulses and short enough for mirror position x to not vary greatly during this time. This yields a number of measurements $N = T/\tau$, where τ is the time between each measurement, which allows us to reduce the uncertainty in the measurement of x by \sqrt{N} while simultaneously increasing the back-action momentum transfer by the same factor, due to accumulation of all individual interactions:

$$\Delta x_T = \Delta x_{\text{measured}} \frac{1}{\sqrt{N}} = \Delta x_{\text{measured}} \sqrt{\frac{\tau}{T}}$$
$$\Delta p_T = \Delta p_{\text{BA}} \sqrt{N} = \Delta p_{\text{BA}} \sqrt{\frac{T}{\tau}}.$$
(3.10)

To arrive at a continuous measurement of mirror position, then, we consider the time interval between each successive measurement to be small: $\tau \to 0$.

Following this definition, it is perhaps easier to consider the double-sided power spectral densities of the measurement noise and back-action force noise, S_x and S_F :

$$S_x = \frac{S_\phi}{4\mathcal{F}^2 k_0^2}, \quad S_F = \frac{4\mathcal{F}^2 S_\mathcal{I}}{c^2}$$
 (3.11)

(c.f. Eqs. 3.2 and 3.7 describing discrete position measurement noise and back-action noise), where

$$S_{\phi} = \lim_{\tau \to 0} (\Delta \phi)^2 \tau \,, \quad S_{\mathcal{I}} = \lim_{\tau \to 0} \frac{(\Delta W)^2}{\tau} \,. \tag{3.12}$$

Comparing Eqs. 3.11 and 3.12 with Eq. 3.8, we find that

$$S_x S_F = \frac{S_\phi S_\mathcal{I}}{k_0^2 c^2} = \frac{S_\phi S_\mathcal{I}}{\omega_0^2} \ge \frac{\hbar^2}{4} \,. \tag{3.13}$$

This relation describes a general model of the uncertainty in a continuous linear quantum measurement, for the specific case where the measurement noise and back-action are uncorrelated. We will refer to a linear measurement system which obeys these constraints as a *Simple Quantum Meter* (SQM). While not a general theory, this relation is very appropriate when characterising the Standard Quantum Limit of an interferometric gravitational wave detector, where the uncorrelated measurement noise and back-action force arise from shot noise on the detected photons and radiation pressure interactions between photons and mirrors.

3.1.4 Standard Quantum Limit

The Standard Quantum Limit (SQL) describes the lower limit of the sum quantum noise power spectral density – the combined power spectral density of all quantum noise sources – that a gravitational wave detector (an example of a "Simple Quantum Meter", as defined in Section 3.1.3) can have at any given frequency, Ω . This limit can be found by setting the measurement noise and back-action forces, $x_{\rm fl}$ and $F_{\rm BA}$, to the minimum combined values allowed by quantum measurement theory. With reference to Eq. 3.13, this can expressed as [44, 45]

$$S_x(\Omega)S_F(\Omega) = \frac{\hbar^2}{4}.$$
(3.14)

Further, if we define the susceptibility as relating force and displacement of the test mass ¹ via $x(\Omega) = \chi F(\Omega)$, then the sum quantum noise power spectral density of the SQM may be expressed as

$$S(\Omega) = \frac{S_x(\Omega)}{|\chi(\Omega)|^2} + S_F(\Omega).$$
(3.15)

3.1.5 Free Mass Standard Quantum Limit

We now have the theoretical framework to examine the SQL in the form most appropriate for conventional gravitational wave detectors: that for a free mass. The test-masses (i.e. the cavity optics) of a gravitational wave detector are suspended as multi-stage pendulums to isolate as much as possible from environmental factors which may increase the displacement noise, such that high above the mechanical resonant frequencies of the suspensions they can be considered to a very good approximation as free masses.

To evaluate this free mass model, we define a new parameter, Ω_q , to quantify the *precision* of measurement:

$$\Omega_q = \left(\frac{S_F}{M^2 S_x}\right)^{1/4}.$$
(3.16)

From Eq. 3.14, it follows that

$$S_x = \frac{\hbar}{2M\Omega_q^2}, \quad S_F = \frac{\hbar M\Omega_q^2}{2}. \tag{3.17}$$

Thus, for higher measurement precision (i.e. higher Ω_q), we see smaller measurement noise spectral density (i.e. smaller S_x) and larger back-action noise (i.e. larger S_F).

For a free mass M, force and distance are related by $F = -M\Omega^2 x$, such that the susceptibility is simply

$$\chi(\Omega) = -\frac{1}{M\Omega^2} \tag{3.18}$$

 $^{^1\}mathrm{The}$ susceptibility is thus a linear response function between force applied to and position of the test mass

for mass M. Thus, for the case of a free mass, the sum quantum noise spectral density (Eq. 3.15) becomes [46, 37]

$$S_{\text{F.M.}}(\Omega) = M^2 \Omega^4 S_x + S_F = \frac{1}{2} \hbar M \Omega_q^2 \left(\frac{\Omega^4}{\Omega_q^4} + 1\right) \,. \tag{3.19}$$

Finally, if we choose the measurement noise and back action noise to contribute equally to the sum noise, and in doing so optimise for their combined minimum, the SQL, then $\Omega_q = \Omega$ and we find that

$$S_{\rm SQL}^{\rm F.M.}(\Omega) = \hbar M \Omega^2 \,. \tag{3.20}$$

This relation indicates that, for the case of a free mass, the SQL defines a "border" of maximum sensitivity, below which the sum noise cannot drop.

3.1.6 HARMONIC OSCILLATOR SQL

From Eqs. 3.15, 3.18 and 3.20, it is clear that an object with larger susceptibility, $|\chi|$, at a particular frequency will have a smaller SQL. For a harmonic oscillator with resonant frequency Ω_0 , the susceptibility becomes

$$\chi_{\rm HO}(\Omega) = \frac{1}{M(\Omega_0^2 - \Omega^2)},$$
(3.21)

with the result that the sum noise power spectral density (Eq. 3.15) becomes [46, 37]

$$S_{\text{SQL}}^{\text{H.O.}}(\Omega) = M^2 (\Omega_0^2 - \Omega^2)^2 S_x + S_F = \frac{\hbar M \Omega_q^2}{2} \left[\frac{(\Omega_0^2 - \Omega^2)^2}{\Omega_q^4} + 1 \right].$$
 (3.22)

As the observation frequency approaches the resonant frequency of the mechanical oscillator, the first term in Eq. 3.22 tends to zero, corresponding to a reduction in measurement noise. Measurement precision Ω_q affects the Q of the resonant increase in sensitivity about this frequency, with a weaker measurement (i.e. smaller Ω_q) having reduced bandwidth but increased sensitivity. Optimising once again as for the free mass SQL, the harmonic oscillator SQL is

$$S_{\text{SQL}}^{\text{H.O.}}(\Omega) = \hbar M |\Omega_0^2 - \Omega^2|. \qquad (3.23)$$



Figure 3.2: Total quantum noise power spectral densities for the free mass (thick black line) and harmonic oscillator (colour) regimes, plotted at frequencies relative to the resonant frequency of the simulated oscillator. Each harmonic oscillator plot uses a different measurement precision, Ω_q , with a weaker measurement (or lower Ω_q) corresponding to a higher-Q resonant dip beneath the SQL. The free-mass standard quantum limit is represented by the thin black line.

Expressing the force SQL for both free mass and harmonic oscillator regimes as a strain sensitivity, the noise spectral densities are [47, 37]

$$S_{\text{SQL, h}}^{F.M.}(\Omega) = \frac{2\hbar}{ML^2\Omega_q^2} \left[1 + \left(\frac{\Omega_q}{\Omega}\right)^4 \right], \qquad (3.24)$$

$$S_{\text{SQL,h}}^{\text{H.O.}}(\Omega) = \frac{2\hbar |\Omega_0^2 - \Omega^2|}{ML^2 \Omega_q^4} \left[\left(\frac{\Omega_q}{\Omega} \right)^4 \left(1 + \frac{(\Omega_0^2 - \Omega^2)^2}{\Omega_q^4} \right) \right].$$
(3.25)

Examples of these are plotted in Figure 3.2. This particular regime, in which an optical system sampling the position of a mass behaves as a harmonic oscillator, may be arrived at in a detector featuring Fabry-Pérot arm cavities by detuning of each cavity from the point of resonance, resulting in a coupling of the optical power to the mirror positions with a strong linear component. This effect is derived from the optical and mechanical contributions in the following section.

3.2 RADIATION PRESSURE AND OPTICAL CAVITY DYNAM-ICS

In this section we derive the principle of opto-mechanical coupling in a Fabry-Pérot cavity, beginning from a simple mathematical description of the resonant, input and output fields of a static two-mirror cavity. This framework is extended to include the effects of the intra-cavity radiation pressure forces acting upon freely-suspended test masses in order to derive the optical spring effect. This description is then further elaborated to include the frequency-dependence of the optical spring effect, and justification provided for the extension of this technique to include multiple optical spring effects in a single system.

3.2.1 FIELD AND POWER EQUATIONS FOR A FABRY-PÉROT CAVITY

Figure 3.3 presents a sketch of a simple two-mirror Fabry-Pérot cavity and associated amplitude fields. In this diagram, E_I represents the input optical field incident upon and E_R the field reflected from the cavity. E_R is composed of some fraction of E_I reflected directly from the input test mass (ITM) and whatever portion of the circulating intracavity field has leaked out through the ITM. The transmitted field, E_T , is composed of intra-cavity light that has leaked out through the end test mass (ETM). The circulating fields in either direction are described by E_{C1} and E_{C2} , where $E_{C1'}$ and $E_{C2'}$ are these same fields after propagating down the length of the cavity, L. The ITM and ETM have amplitude reflection and transmission coefficients $\rho_{1,2}$ and $\tau_{1,2}$, with which each field can be described in terms of the contributory fields [48, 49]:

$$E_R = \rho_1 E_I + i\tau_1 E_{C2'} = \rho_1 E_I + i\tau_1 E_{C2} e^{-i\phi}$$
(3.26)

$$E_T = i\tau_2 E_{C1'} = i\tau_2 E_{C1} e^{-i\phi}$$
(3.27)

$$E_{C1} = i\tau_1 E_I + \rho_1 E_{C2'} = i\tau_1 E_I + \rho_1 E_{C2} e^{-i\phi}$$
(3.28)

$$E_{C2} = \rho_2 E_{C1'} = \rho_2 E_{C1} e^{-i\phi} . aga{3.29}$$

Note that here we have identified the additional phase accumulated as the field propagates down the cavity length as $\phi = \frac{2\pi L}{\lambda} = \frac{\omega L}{c}$. If we now opt to consider the *static* Fabry-Pérot cavity case, in which the input field is assumed constant (i.e. $E_I(t) = E_I$) and the mirrors static, the field description may be simplified to only three expressions, for the forwardpropagating intra-cavity field, $E_C \equiv E_{C1}$, and the simplified reflected and transmitted



Figure 3.3: Schematic diagram of a simple two-mirror cavity. Identified are the incident and reflected fields, the circulating intra-cavity fields, and the transmitted field, where the symbols are as defined in the text.

fields:

$$E_C = i\tau_1 E_I + \rho_1 \rho_2 E_C e^{-2i\phi} = \frac{i\tau_1 E_I}{1 - \rho_1 \rho_2 e^{-2i\phi}}$$
(3.30)

$$E_R = \rho_1 E_I + i\tau_1 \rho_2 E_C e^{-2i\phi} = \left[\rho_1 - \frac{\tau_1^2 \rho_2 e^{-2i\phi}}{1 - \rho_1 \rho_2 e^{-2i\phi}}\right] E_I$$
(3.31)

$$E_T = i\tau_2 E_C e^{-i\phi} = -\frac{\tau_1 \tau_2 E_I e^{-i\phi}}{1 - \rho_1 \rho_2 e^{-2i\phi}}.$$
(3.32)

This then allows us to define the total reflectivity of the cavity with respect to the incident light field with the cavity reflection coefficient, ρ_C :

$$E_R = \rho_C E_I \tag{3.33}$$

$$\Rightarrow \rho_C = \rho_1 - \frac{\tau_1^2 \rho_2 e^{-2i\phi}}{1 - \rho_1 \rho_2 e^{-2i\phi}}.$$
(3.34)

Finally, using Eq. 3.30 we can express the maximum power gain in the cavity by considering the case for which the cavity is on-resonance; i.e. $L = \frac{n\lambda}{2}$, where n is an integer:

$$E_{C} = \frac{i\tau_{1}}{1 - \rho_{1}\rho_{2}}E_{I} = ig_{C}E_{I}$$
(3.35)

$$\Rightarrow g_C = \frac{\gamma_1}{1 - \rho_1 \rho_2}. \tag{3.36}$$

The full relationship between cavity offset phase, $\theta = \frac{2\pi x}{\lambda}$, and intra-cavity power scales with the Airy intensity profile [50]:

$$P = |E_C|^2 = g_C^2 |E_I|^2 \frac{1}{1 + F \sin^2 \theta}.$$
(3.37)

27



Figure 3.4: Simulation of intra-cavity powers as cavity tuning is swept through 360°, for finesse of 10 and 100, using the Airy function of Eq. 3.37. Identified are the linewidth, γ , of the Airy peak and the free spectral range, $\Delta \omega$, between each pair of resonances.

Here F is the coefficient of cavity finesse,

$$F = \frac{4\rho_1 \rho_2}{(1 - \rho_1 \rho_2)^2}, \qquad (3.38)$$

which allows the quality of a cavity to be determined in relation to its losses via the finesse, \mathcal{F} :

$$\mathcal{F} = \frac{\pi}{2}\sqrt{F} \,. \tag{3.39}$$

From Eq. 3.37, we see that maximum gain will be achieved when the Airy function approaches unity; i.e. when cavity tuning $\theta = n\pi$. The separation between each subsequent cavity resonance in the frequency domain is the *free spectral range* (FSR), defined by

$$\Delta\omega = 2\pi \frac{c}{2L} = \frac{\pi c}{L} \,, \tag{3.40}$$

where c is the speed of light in vacuum. The finesse and free spectral range are then related by the cavity linewidth, γ , defined as the half-width at half-maximum (HWHM) of the resonant Airy peak:

$$2\gamma = \frac{\Delta\omega}{\mathcal{F}} \,. \tag{3.41}$$

Figure 3.4 presents a simulation of the normalised intra-cavity power as the cavity tuning is swept through two free spectral ranges.



Figure 3.5: Schematic diagram of a simple two-mirror cavity acted upon by radiation pressure. At each optic, the radiation pressure force F_{RP} is balanced by external force $F_{ext.}$, composed of the pendulum restoring force and electromagnetic feedback used to control cavity length.

3.2.2 Optical Springs

Beginning with the descriptions of a simple static two-mirror cavity from the previous section, we are now in a position to expand upon this mathematical framework to include the dynamic interactions between the suspended optics and the light fields circulating in the system.

Figure 3.5 again illustrates a Fabry-Pérot cavity with suspended mirrors and circulating fields. In this case, we now assume the two optics to be suspended as pendulums and the cavity to be held *on resonance*, by which active feedback is applied to one or both optics in order to maintain a fixed cavity length that maintains the intra-cavity power at the maximum of the Airy function, as described by Eq. 3.37. The most common method of feedback to this end for a suspended Fabry-Pérot cavity involves derivation of control signals proportional to the cavity length by modulation of the input field and monitoring of the field reflected from the cavity; this method is known as the *Pound-Drever-Hall* technique [51], and is explained in detail in Section 5.2. The radiation pressure force imparted to each suspended optic by the circulating on-resonance optical power in the cavity is thus balanced by a combination of the restoring force of each pendulum and the feedback signal, conventionally applied using electromagnetic actuators at the rear of each optic. This electromagnetic actuation is described in Section 4.5.4. Noting that momentum must be conserved at both optics, the radiation pressure force acting upon each may be expressed in terms of vector sums of the optical fields at each point (where each field E is as defined in Section 3.2.1) [52]:

$$F_{RP_1}(t) = \frac{1}{c} \left[|E_I(t)|^2 + |E_R(t)|^2 - |E_{C1}(t)|^2 - |E_{C2'}(t)|^2 \right]$$
(3.42)

$$F_{RP_2}(t) = \frac{1}{c} \left[|E_{C1'}(t)|^2 + |E_{C2}(t)|^2 - |E_T(t)|^2 \right].$$
(3.43)

For each pendulum-suspended optic the total force is related to the equation of motion via

$$\sum F = m\ddot{x} + c\dot{x} + kx \qquad (3.44)$$

$$\Rightarrow \ddot{x}_1 + \beta_1 \dot{x}_1 + \omega_{p1}^2 x_1 = \frac{F_{RP_1} + F_{ext_1}}{m_1}, \qquad (3.45)$$

$$\ddot{x}_2 + \beta_2 \dot{x}_2 + \omega_{p2}^2 x_2 = \frac{F_{RP_2} + F_{ext_2}}{m_2}.$$
(3.46)

Here we have defined each pendulum in terms of suspended mass m, damping rate β , stiffness k and fundamental resonant frequency $\omega_p = \sqrt{\frac{k}{m}}$, and x denotes the position of the optic. If we now make the assumption that the radiation pressure force imposed on the ITM due to the field incident on the cavity is negligible when compared to the intracavity radiation pressure force on each optic (which is quite reasonable for our purposes, as gravitational wave interferometers and prototype experiments typically use ratios of intra-cavity to input powers of order some thousands [18, 16]), we are able to express the equations of motion in terms of the relative cavity length: $\xi = x_2 - x_1$. This assumption implies that the internal radiation pressure forces are of equal magnitude and act in opposite directions, such that $F_{RP_1} = F_{RP_2} = F_{RP}$ and

$$\frac{F_{RP} + F_{ext_2}}{m_2} - \frac{F_{RP} + F_{ext_1}}{m_1} = \ddot{\xi} + \beta \dot{\xi} + \omega_p^2 \xi \qquad (3.47)$$

$$\Rightarrow \ddot{\xi} + \beta \dot{\xi} + \omega_p^2 \xi = \frac{F_{RP} + F_{ext}}{\mu}, \qquad (3.48)$$

where $\mu = \frac{m_1 m_2}{m_1 + m_2}$ is the reduced mass of the system and we have made the further assumption for simplicity that the two pendulums have the same resonant frequency, ω_p .

Using the previously derived Eq. 3.35 for the static forward-propagating on-resonance intra-cavity field and 3.43 for the radiation pressure acting upon the end test mass, we can now express the maximum radiation pressure force experienced by the ETM as

$$F_{RP_{max.}} = \frac{1}{c} \left[\left(\frac{\tau_1 E_I}{1 - \rho_1 \rho_2} \right)^2 + \left(\frac{\tau_1 \rho_2 E_I}{1 - \rho_1 \rho_2} \right)^2 - \left(\frac{\tau_1 \tau_2 E_I}{1 - \rho_1 \rho_2} \right)^2 \right].$$
(3.49)

30

Rearranging Eq. 3.37 then allows us to express the maximum radiation pressure force on the ETM in terms of intra-cavity power, cavity tuning and mirror properties alone:

$$E_I^2 = \frac{\tau_1 P}{1 - \rho_1 \rho_2} \frac{1}{1 + F \sin^2 \theta}, \qquad (3.50)$$

$$\implies F_{RP_{max.}} = \frac{P}{c} \frac{(1+\rho_2^2 - \tau_2^2)}{1+F\sin^2\theta}.$$
 (3.51)

As a useful approximation, for the general case where the end optic consists of a very highly-reflective optic, such that $\rho_2 \simeq 1$ and $\tau_2 \simeq 0$, this reduces to

$$F_{RP_0} \simeq \frac{2P_0}{c} \,. \tag{3.52}$$

For this same general case, the radiation pressure force will vary with detuning according to the Airy function once again, such that for the static Fabry-Pérot cavity the equation of motion, Eq. 3.48, reduces to

$$\mu \omega_p^2 \xi - F_{ext.}(t) = \frac{2P_0}{c} \frac{1}{1 + F \sin^2 \theta} \,. \tag{3.53}$$

If we now apply a static shift to the position of one mirror of the cavity of size dx, or equivalently shift the frequency of the input field by $d\omega$, the static change in the radiation pressure force on the end optic is

$$\mathrm{d}F_{RP} = -\frac{2}{c}\frac{\partial P}{\partial x}\mathrm{d}x\tag{3.54}$$

When the cavity is on-resonance, the term $\frac{\partial P}{\partial x} = 0$; mathematically, this is the gradient of the intra-cavity power curve in Figure 3.4. If a larger static shift is applied to either the length or laser frequency, such that the power in the cavity drops by some amount, and for very short perturbations to the cavity length or laser frequency about this point, the gradient of the power curve may be considered linear to first order. In this case, Eq. 3.54 takes the form of Hooke's law for a spring: F = -kx. It then follows that the *optical spring* can be expressed as the predominantly linear coupling between cavity power and suspended optic position:

$$k_{OS} = -\frac{2}{c} \frac{\partial P}{\partial x} = -\frac{2}{c} \frac{\partial P}{\partial \theta} \frac{\partial \theta}{\partial x}.$$
(3.55)

For a high-finesse cavity the detunings applied are very small when compared to the free spectral range, so the approximation $\sin \theta \simeq \theta$ can be made. Thus the partial derivatives



Figure 3.6: Simulation of intra-cavity power (blue, according to Eq. 3.37) and optical spring constant (red, according to Eq. 3.58) while the detuning is swept from -5γ to 5γ , where γ is the linewidth of the cavity. For this simulation, 1 W of carrier power is incident upon a 10,000 finesse cavity with input coupler power transmission of 100 ppm.

in Eq. 3.55 can be evaluated as

$$\frac{\partial P}{\partial \theta} = -\frac{2P_0 F \theta}{(1+F\theta^2)^2} \tag{3.56}$$

and
$$\frac{\partial \theta}{\partial x} = \frac{2\pi}{\lambda}$$
, (3.57)

which then allows us to define the optical spring restoring force as [52, 53]

$$k_{OS} = \frac{8\pi P_0 F \theta}{c\lambda (1 + F \theta^2)^2} \,. \tag{3.58}$$

This then has associated resonant frequency of

$$f_{OS} = \frac{1}{2\pi} \sqrt{\frac{k_{OS} + k_p}{\mu}} \simeq \frac{1}{2\pi} \sqrt{\frac{k_{OS}}{\mu}},$$
 (3.59)

where we have made the assumption that the restoring force of the pendulum, k_p , is far exceeded by that of the opto-mechanics. Figure 3.6 plots Eq. 3.58 to show how k_{OS} varies with detuning and intra-cavity power. Note that the maximum spring constant, and hence maximum observed associated resonant frequency, is achieved when the gradient of the power curve, equivalent to $\frac{\partial P}{\partial \theta}$, is a maximum. This occurs at a detuning of $\delta \simeq 0.57 \gamma$, where δ can equivalently be expressed as a frequency shift of the input field or length change of the cavity.

3.2.3 FREQUENCY-DEPENDENT OPTICAL SPRING

Note that two key assumptions were made in the derivation of the optical spring effect (Eq. 3.58): that the mirror motion is static, and that the fields circulating in the cavity respond instantaneously to any force. However, an optical cavity responds to motion on a time-scale consistent with the inverse of its linewidth, γ ; when dynamic elements are introduced at a frequencies of order γ , the frequency response of the system becomes very important in deriving the true opto-mechanical response [54, 55, 56].

In deriving the full expression for the optical spring, it is useful to turn to the *two-photon* formalism, originally developed by Caves and Schumaker [57]. This method decomposes the signals from an interferometer into the pair of sidebands imposed around the carrier field upon parametric modulation of the arm lengths - equivalently, upon modulation of the position of one optic in a single Fabry-Pérot cavity. These pairs of sidebands occur at frequencies $\omega_{1,2} = \omega_0 \pm \Omega$, where ω_0 is the carrier frequency [58]. This also allows decomposition of light noise into contributions from amplitude, which affects both sidebands at $\pm \Omega$ equally, and phase, which affects each sideband oppositely. While not widely adopted, this approach is ideal for gravitational wave applications, in which conventionally there is a single carrier frequency and inputs and outputs are dependent on phase modulation of this carrier, either by motion of the optics or the effect of passing gravitational waves. As a result, the two-photon formalism has been used extensively in the field. Each element of a gravitational wave interferometer or related system may then be described in this framework using *input-output relations* which relate incoming and outgoing fields and the effects of passing gravitational waves.

A full derivation of the frequency-dependent optical spring using this formalism has been performed by S. Danilishin and F. Khalili [37], and similar approaches taken in [59] and [60]. Rather than cover the extensive mathematical derivation here, the pertinent points as relating to the basics of the formalism and the operation of the experimental system are identified and the frequency-dependent spring constant inferred therefrom.

The formalism begins by defining the annihilation and creation operators of the optical mode at frequency ω as \hat{a}_{ω} and $\hat{a}^{\dagger}_{\omega}$. Accordingly, the modulation sideband amplitude operators may be defined as

$$\hat{a}_{+} = \hat{a}_{\omega_0 + \Omega} \tag{3.60}$$

and
$$\hat{a}_{-} = \hat{a}_{\omega_0 - \Omega}$$
. (3.61)

These may then be combined to yield the two-photon cosine and sine quadrature amplitude operators:

$$\hat{a}_{c}(\Omega) = \frac{\hat{a}_{+} + \hat{a}_{-}^{\dagger}}{\sqrt{2}}$$
(3.62)

and
$$\hat{a}_s(\Omega) = \frac{\hat{a}_+ - \hat{a}_-^{\dagger}}{i\sqrt{2}}.$$
 (3.63)

Combining these into a single vector, the two-photon quadratures are represented as

$$\hat{\mathbf{a}} \equiv \begin{bmatrix} \hat{a}_c \\ \hat{a}_s \end{bmatrix} . \tag{3.64}$$

These operators, \hat{a}_c and \hat{a}_s , are chosen to represent the single-photon sine and cosine quadrature amplitudes of a laser mode, such that the electric field operator of a propagating laser mode may be expressed in terms of these operators as

$$\hat{E} = u(x, y, z)\mathcal{C}_0[\hat{a}_c(t)\cos\omega_0 t + \hat{a}_s(t)\sin\omega_0 t].$$
(3.65)

where u(x, y, z) represents the spatial mode shape and $C_0 = \sqrt{\frac{4\pi\hbar\omega_0}{Ac}}$, where \mathcal{A} is the cross-sectional area of the beam. Now appying this to a Fabry-Pérot cavity, which allows us to make the simplification that a single mode propagates, for particular examination of the radiation-pressure interactions with each of the two optics, we define the intra-cavity quantum field amplitudes impinging upon each optic (and thus analogous to the classical fields $E_{C1'}$ and $E_{C2'}$ in Figure 3.3) to be:

$$\hat{\mathbf{e}}(\Omega) = \frac{1}{\sqrt{\tau}} \mathbb{L} \left[\sqrt{\gamma_{1,2}} \hat{\mathbf{a}}_{1,2}(\Omega) + \sqrt{\gamma_{2,1}} \hat{\mathbf{a}}_{2,1}(\Omega) + \hat{\mathbf{X}}(\Omega) \right]$$
(3.66)

where the subscripts 1 and 2 denote the field impinging upon each optic, $\tau = (2f_{\text{FSR}})^{-1}$ (where $f_{\text{FSR}} = \Delta \omega / 2\pi$ is the free spectral range in Hz), $\hat{\mathbf{X}}(\Omega)$ is an operator relating the intra-cavity electric field strength quadratures, \hat{E} , and the variation in cavity length, $\hat{x}(\Omega)$, i.e.

$$\hat{\mathbf{X}}(\Omega) = \hat{E} \frac{k_0 \hat{x}(\Omega)}{\sqrt{\tau}}, \qquad (3.67)$$

where $k_0 = \omega_0/c$, $\gamma = \gamma_1 + \gamma_2$ is the cavity linewidth and

$$\gamma_{1,2} = \frac{T_{1,2}}{4\tau} \,, \tag{3.68}$$

 $T_{1,2}$ is the transmission of each optic,

$$\mathbb{L} = \frac{1}{\mathcal{D}(\Omega)} \begin{bmatrix} \gamma - i\Omega & -\delta \\ \delta & \gamma - i\Omega \end{bmatrix}, \qquad (3.69)$$

and $\mathcal{D}(\Omega) = (\gamma - i\Omega)^2 + \delta^2$ provides the cavity pole. Here, δ is the detuning from the carrier frequency in Hz and Ω is the frequency of observation. Noting that for the single-mode approximation used here we can make the same assumption that was used to arrive at Eqs. 3.48 in the derivation of the static optical spring – that the internal radiation pressure forces upon each optic are of equal magnitude – we find that the back-action force on each optic is

$$\hat{F}_{1,2}(\Omega) \equiv \hat{F}_{BA}(\Omega) = 2\hbar k_0 \hat{E} \hat{\mathbf{e}}(\Omega) \,. \tag{3.70}$$

Separating the total back-action radiation pressure force into the sum of random and dynamical terms,

$$\hat{F}_{BA}(\Omega) = \hat{F}_{BA}^{(0)}(\Omega) - K(\Omega)\hat{x}(\Omega), \qquad (3.71)$$

allows us finally to separate out the random and dynamical (i.e. opto-mechanical) components as

$$\hat{F}_{BA}^{(0)}(\Omega) = \frac{2\hbar k_0 \hat{E}}{\sqrt{\tau}} \mathbb{L}(\Omega) \left[\sqrt{\gamma_1} \hat{\mathbf{a}}_1(\Omega) + \sqrt{\gamma_2} \hat{\mathbf{a}}_2(\Omega)\right]$$
(3.72)

and
$$K(\Omega) = \frac{2\hbar k_0 \hat{E}}{\sqrt{\tau}} \mathbb{L}(\Omega) \hat{\mathbf{X}}(\Omega).$$
 (3.73)

Upon simplification, this reduces to

$$K(\Omega) = k_{OS} \frac{1 + \delta_{\gamma}}{(1 + i\Omega_{\gamma})^2 + \delta_{\gamma}}, \qquad (3.74)$$

where we have defined $\delta_{\gamma} = \delta/\gamma$ and $\Omega_{\gamma} = \Omega/\gamma$, and k_{OS} is the static (i.e. DC) optical spring constant derived in Eq. 3.58. Noting the complex term, we can separate Eq. 3.74 into its real and imaginary parts:

$$K_{OS}(\Omega) = K(\Omega) - i\Gamma(\Omega). \qquad (3.75)$$

Since the second term is proportional to $i\Omega$, it can be thought of as a viscous damping term (alternatively, as an "optical friction" [37]) which is always of opposite sign to the restoring force, K. Crucially, this indicates the optical rigidity is *unstable*. This is illustrated in Figure 3.7, in which the real and imaginary parts of Eq. 3.74 (equivalently, the restoring force and damping of the optical spring) are plotted for varying detuning and observation frequency. This means that for a positively-detuned cavity exhibiting positive optical rigidity there will exist an anti-damping effect, and for a negatively-detuned cavity there will exist positive damping but an anti-restoring force.

3.2.4 Multiple Optical Springs

The potential benefits available to a system exploiting opto-mechanical effects for narrowband enhanced quantum sensitivity may be extended with the application of multiple detuned optical fields. As has been noted, a single optical spring is inherently unstable. This is ordinarily corrected by the application of electronic feedback systems shaped in such a way as to keep the system from going unstable; however, with a second optical spring of appropriate detuning the system may be stabilised purely optically [61]. Configurations involving two or more optical springs can also allow re-shaping of the noise spectral density of the system.

When two optical springs are combined in a cavity, their spring constants sum. If the desired effect is a stable opto-mechanical system, a high-power, highly-detuned positive spring, with a large restoring force and small anti-damping, can be combined with a lower-power smaller-detuned anti-spring, with small anti-restoring force and larger damping. The total optical rigidity is then arranged to have positive real and imaginary parts in Eq. 3.74 [60]:

$$K_{\text{total}}(\Omega) = K_{OS,1}(\Omega) + K_{OS,2}(\Omega)$$

= $k_{OS,1} \frac{1 + \delta_{\gamma,1}}{(1 + i\Omega_{\gamma,1})^2 + \delta_{\gamma,1}} + k_{OS,2} \frac{1 + \delta_{\gamma,2}}{(1 + i\Omega_{\gamma,2})^2 + \delta_{\gamma,2}}$ (3.76)
= $(K_1 + K_2) - i(\Gamma_1 + \Gamma_2)$.

Figure 3.8 depicts this combined rigidity as a vector sum in the $K-\Gamma$ plane. This stable trapping of an optic has been demonstrated in the past, typically for light masses (of order 1 g) in short, tabletop experiment-scale cavities. These and other applications of one or more optical springs are discussed in Section 3.2.6.



Figure 3.7: Variation of opto-mechanical restoring force (a) and damping (b) with detuning and observation frequency. For these simulations, 1 W of optical power in incident upon a Fabry-Pérot cavity of length 10 m, finesse 10,000 and input coupler power transmission 100 ppm.



Figure 3.8: Graphical demonstration of the stable optical trap method. In the red and blue curves, detuning is swept from $-\infty$ to $+\infty$ for a high power carrier (blue) and low power sub-carrier (red). Choosing a particular positive detuning of the carrier (solid arrows) and negative detuning of the sub-carrier (dashed) results in a vector sum of K and Γ values for which both are positive.

3.2.5 MOTIVATION FOR THIS THESIS

Proposals have been made for the inclusion of multiple-spring systems in advanced gravitational wave detectors. One such example put forward by Rehbein *et al.* [61] proposes employing a second optical spring in an interferometer incorporating detuned signal recycling, with the result that a stable double-optical-spring (DOS) combination is achieved. This in turn can reduce the undesirable complexity in the control systems, and accompanying classical noise, associated with stabilisation of a single optical spring. In such a setup, the supplementary optical spring is achieved by injecting a second laser beam into the interferometer via the bright port which is resonant in the arm cavities and detuned in the signal recycling cavity. By appropriate choice of detunings in the signal recycling cavity, stability of the opto-mechanical system can be achieved, along with a potential improvement in sensitivity.

Rehbein *et al.* identify two distinct schemes of double-optical-springs: *weak stabilisation*, in which (as outlined in Section 3.2.4) a weak anti-spring may be used to stabilise a strong spring, and *annihilation*, whereby the optical springs are formed of two carriers of equal power and opposite detuning, with the result that the net opto-mechanical rigidity cancels. Simulating the effect of these schemes as an upgrade to the Advanced LIGO detectors, they draw a number of important conclusions. Firstly, when compared to the quantum noise limit of a single-optical-spring interferometer, the inclusion of the detuned sub-carrier generally results in a weaker net opto-mechanical rigidity, with the result that as well as yielding increased stability the sensitivity of the detector is improved below the associated resonance. They further show that, compared to the full Advanced LIGO noise budget when optimised for detection of binary neutron star systems, the event rate observed by the instrument can be improved by up to 36% when compared to the single optical spring case. Under the assumption of some reasonable future improvements to the classical noise floor of the detector, the event rate is further improved by 110% when compared with an equivalent single optical spring system. Finally, they note that future reductions in the levels of classical noise would result in further areas of the spectrum in which sensitivity could be improved by exploiting more than two carrier fields.

The simulations performed by Rehbein *et al.* in particular motivate the experimental work described in the following chapters. This experiment is designed to explore the practical concerns of implementing and controlling such a multiple-spring system in a large-scale detector; the measurements presented in this thesis seek to examine the interactions and couplings between each detuned field, the optics of the cavities and the control loops which maintain stability. As is shown in Chapter 4, this experiment takes the form of a three-mirror coupled-cavity system capable of sustaining several optical springs. This is implemented in the Glasgow 10 m Prototype Laboratory, and is intended to represent a system dynamically-similar to that proposed by Rehbein *et al.* A simple schematic of the system, termed the Multiple Optical Spring Experimental System (MOSES), is shown in Figure 3.9.

In Table 3.1 some key parameters of the Rehbein *et al.* proposal are compared to those of the experiment depicted in Figure 3.9. In this table, we have chosen to represent parameters from [61] optimised for Advanced LIGO narrowband neutron star binary detection mode, as this was identified as the configuration in which the greatest sensitivity improvement could be achieved over the single optical spring case. Comparing the two, the system described in this thesis will provide a suitable test-bed for investigating a dynamically-similar system to the DOS proposal, but on a smaller scale. The system has further been designed to be very flexible, with the potential to implement a wide range of combined opto-mechanical rigidities. This will allow investigation of further optical spring configurations beyond the application of the DOS scheme to the detuned signal-recycled



Figure 3.9: Schematic sketch of the experimental system described in detail in Chapter 4. Three mirrors (Left Test Mass, Central Test Mass and Right Test Mass) form a threemirror coupled-cavity system capable of sustaining several optical springs. In the first instance, as shown, a single carrier frequency is incident upon each cavity, such that each cavity may support a single optical spring and the two couple mechanically via the opaque CTM.

Parameter	DOS Value	MOSES Value
Optic masses	$40\mathrm{kg}$	$100\mathrm{g}$
Total intra-cavity optical power	$\leq 800 \mathrm{kW}$	$\leq 4 \mathrm{kW}$
Cavity linewidth, γ (approx.)	$\leq 750\mathrm{Hz}$	$750\mathrm{Hz}$
Detuning range (approx.)	$+2.5\gamma,\!-0.65\gamma$	$\pm 5 \gamma$

Table 3.1: Comparison of key parameters for the proposed double-optical-spring (DOS) upgrade to Advanced LIGO and the multiple optical spring apparatus (MOSES) at the Glasgow 10 m Prototype Laboratory. The DOS values presented here are taken from [61] and assume Advanced LIGO operating in narrowband mode optimised for improved neutron star binary event rates. The total intra-cavity optical power in both cases is split between the two carrier fields. The DOS linewidth quoted is the estimated effective linewidth after incorporating detuned signal recycling into the existing Advanced LIGO designs [18]. More detail regarding the design parameters of the MOSES apparatus are located in Chapter 4.

interferometer.

Another proposal for implementation of optical springs in a large-scale detector is the "optical bar" scheme. With this method, the end mirrors of the two interferometer arms are coupled to a central optic *via* optical springs. In the presence of a passing gravitational wave, differential motion of the arms results in local motion of the central mirror below the opto-mechanical resonance frequencies, which can be read out by a supplementary measurement system [62, 25]. The MOSES apparatus of Figure 3.9 displays obvious parallels with this system, and so allows investigation of the coupling of two optical springs to a shared optic and the resulting effects on the control systems governing the overall stability of the three-mirror coupled cavity.



Figure 3.10: Simulation of quantum noise performance of a coupled cavity system with an optical spring in each cavity, performed using Optickle [81]. Quantum noise levels are shown for each cavity and compared to the equivalent isolated, un-coupled cavity cases. The cavity parameters (e.g. finesse, optic masses, etc.) are similar to those obtained for the experimental system described in Chapter 4. Note in particular the summed quantum noise of the entire system, which is observed to dip beneath the free mass SQL over a period of almost 400 Hz, far wider than is possible for either individual coupled or isolated spring.

The monitoring of the shared optic in the optical bar system itself comprises a quantum non-demolition technique, in which the low-power readout system enables the coupling between the optics to be monitored without perturbation of the optical and mechanical systems. This technique, commonly referred to as "local readout", was the subject of a previous investigation at the Glasgow 10 m prototype by J. Macarthur [53]. Rehbein *et al.* [63] also propose using this method to improve the low-frequency sensitivity of an interferometer incorporating detuned signal-recycling by monitoring of the position of the input test mass to each arm cavity. The MOSES apparatus of Figure 3.9 can also be operated as an analogue of the local readout technique, whereby one of the two cavities is operated on-resonance to monitor the position of the Central Test Mass whilst it is rigidly coupled to the input test mass of the adjacent cavity via opto-mechanical rigidity.

The quantum noise spectrum of a Fabry-Pérot cavity may also be re-shaped with the application of optical springs, as was shown for a single spring in Figure 3.2. If multiple springs are employed in a single system, the quantum noise may be further re-shaped. In

particular, the increase in sensitivity beneath the free mass SQL has the potential to be wider than for a single-spring system. Figure 3.10 presents a simulation of the quantum noise for the MOSES apparatus as shown in Figure 3.9, in which the coupling optic is assumed to be opaque and optical springs are present in both cavities. The quantum sensitivity of each coupled cavity is compared to the equivalent noise limits for the case of a single isolated cavity with one optical spring. Note the perturbation to the equivalent un-coupled optical spring transfer functions below the opto-mechanical resonance resulting from the adjacent cavity. This divergence from the "single spring" case is due to the coupling of the two opto-mechanical responses; equivalently, this can be thought of as the cavity susceptibility (Eq. 3.21) being modified as the effective mass distribution of the system changes upon rigid coupling of input test masses to the shared mass. In particular note that in the quantum noise figure the summed quantum noise dips below the SQL over a bandwidth of several hundred Hz. As we show in Chapter 4, simulations such as this further motivate the design of the experiment that forms the remainder of this thesis.

The optical and mechanical design of the MOSES apparatus is described in the following Chapter 4. Chapter 5 covers the design of the control systems, including implementation of a novel scheme facilitating the detuning of each cavity much further than is conventionally possible. Chapter 6 then covers several measurements performed with the complete system. In this thesis, the measurements are designed to characterise the potential of the apparatus for the accurate measurement of optical spring transfer functions and to investigate couplings between the various opto-mechanical effects and the control loops. Follow-up experiments are anticipated that proceed to investigate stable opto-mechanical systems similar to those envisioned by Rehbein *et al.* as potential upgrades to the Advanced LIGO detectors.

3.2.6 Previous Applications

Opto-mechanical rigidity has been explored and exploited in a number of different ways in recent years, from proof-of-principle tests in tabletop experiments or prototype laboratories to practical application in the large-scale gravitational wave detectors. We shall briefly detail a few of these tests here.

A number of proof-of-principle optical spring tests have been performed on various scales using Fabry-Pérot cavities with mirrors either suspended or mounted upon light flexures. T. Corbitt *et al.* [56] observed and characterised an optical spring on the gram scale, and later demonstrated stable optical trapping and resulting optical cooling using two detuned optical fields on the same scale [64], as was achieved similarly by Mow-Lowry *et al.* [65]. At the Glasgow 10 m prototype laboratory, M. Edgar [52] demonstrated the effect on the larger 100 g scale, by detuning a 10 m optical cavity to observe and characterise the resulting opto-mechanical resonance, yielding an effective modified rigidity of 9.4×10^5 N/m. As was noted in the previous section, J. Macarthur demonstrated the "local readout" method in the Glasgow laboratory [53], by which a secondary low-power readout is used to identify opto-mechanical coupling by monitoring the position of one optic without disturbing the quantum state of the system.

Conventionally, studies of optical springs observe the opto-mechanical interaction to be *dispersive*, such that a shift in optic position results in a corresponding shift in the resonance frequency due to the changing radiation pressure force. However, it has also been shown that *dissipative* coupling can also exist, whereby the *width* of the cavity resonance is affected by a mechanical displacement [66]. This effect can be exploited to approach ground-state cooling of the mechanical oscillator. This dissipative opto-mechanical coupling has been observed in recent years by Xuereb, Schnabel and Hammerer [67] using a Michelson-Sagnac interferometer with a movable micromechanical membrane, with which strong cooling of the mechanical motion was characterised. The same apparatus has been employed by Tarabrin *et al.* [68] to show that when operated off-dark-port a stable single optical spring is possible, as was predicted by Vostrosablin and Vyatchanin [69]. More recently, dispersive and dissipative opto-mechanical coupling has been characterised in this apparatus by Sawadsky *et al.* [70], in particular with observation of optical cooling through dissipative coupling.

As has been touched upon in Chapter 2, an interferometer can be operated with signal recycling if a mirror is inserted between the beamsplitter and output port. This facilitates recycling of the signal sidebands leaving the interferometer, and thus enhancement of the signals over a particular bandwidth dictated by the signal recycling mirror (SRM) reflectivity [71, 72]. The centre frequency of this band is determined by the tuning of the cavity formed between the SRM and the rest of the interferometer [73, 10]. Thus for any particular detuned signal recycling configuration to enhance sensitivity over a particular frequency range, opto-mechanical effects will be important. Compensation of these effects either in the control loops or using supplementary detuned fields motivates a large portion of the work undertaken towards this thesis.

The optical spring effect has also been investigated outside the gravitational wave com-

munity. The effect may be observed in microresonators, in which light is guided around a fabricated resonator, typically a microdisk, microsphere or microtoroidal resonator, with each having particular advantages as regards optical and mechanical quality [74]. These structures make observation and characterisation of optical springs much simpler, due to their small size and and efficient opto-mechanical coupling [75]. The resulting opto-mechanical frequencies range from the MHz (e.g. Hofer, Schliesser and Kippenberg [76]) to GHz (e.g. Ding *et al.* [77]) scales. Such structures can be applied to a diverse range of fields, from quantum optical devices and optical fibre technologies to biological and chemical sensing systems.

In more general applications, K. Agatsuma *et al.* developed a method for precise measurement of laser power using an opto-mechanically-coupled system [78], and A. Mizrahi and L. Schächter have demonstrated the optical trapping effect between two dielectric slabs forming a waveguide [79].

IV CONCEPTUAL EXPERIMENTAL DESIGN AND IMPLEMENTATION

The investigation described in this thesis aims to take advantage of the theory of optomechanical coupling introduced in Chapter 3 to produce and characterise doubly-resonant optical spring structures in a high-finesse coupled-cavity system. Further, the practical concerns associated with operating a prototype-scale interferometer under such a regime will be detailed.

In this chapter the design and construction of most of the experimental system is described. Sections 4.1 and 4.2 present an overview of the experiment and the Glasgow Prototype Laboratory. Section 4.3 covers the investigative simulations performed to inform the basic design of the system. Section 4.4 describes the construction and characterisation of a laser amplification system required to amplify the existing laser source in the laboratory to provide adequate input powers for this experiment. Section 4.5 describes the further optical and mechanical design work necessary prior to commissioning of the system; specifically, this section covers the design of seismic isolation systems and optic suspensions.

The subject of feedback and control of such an experiment will take some time to detail; for that reason, the following Chapter 5 is dedicated to this particular facet of the commissioning of the experiment. The measurements and results from this experiment are then presented in Chapter 6, with conclusions following in Chapter 7.

4.1 OVERVIEW

The principal motivation presented in Chapter 3 was to improve the quantum-limited displacement sensitivity of a Fabry-Pérot cavity by transforming the dynamics of the system into those of a harmonic oscillator. This is accomplished using one or more detuned optical fields yielding increased displacement sensitivity about the associated optomechanical resonances. A design study was embarked upon towards an experiment to demonstrate the feasibility of such a system on a larger scale and with greater flexibility than had been attempted previously. Importantly, this experiment also requires design of innovative feedback and control methods in order to maintain stability – this topic in particular has implications for large-scale interferometers operating with optical rigidity, as well as to the broader world of high-power, high-finesse optical resonators operating with digital feedback.

The resulting Multiple Optical Spring Experimental System (or MOSES) was designed to be a flexible test-bed for various regimes of opto-mechanical coupling between suspended optics. It consists of two identical optical cavities, each 10 m in length, mechanically coupled together by way of housing the end mirrors of both cavities back-to-back within a single suspended optic mount. Each cavity may be held on-resonance or detuned therefrom via feedback to the position of its input mirror.

With this very general description of the experimental concept in place, we may set out the objectives:

- 1. to construct reliable control strategies for radiation-pressure-dominated optical cavities, with the intention of stably detuning each cavity to several times its linewidth in the creation of opto-mechanical rigidity,
- 2. to characterise opto-mechanical rigidity in a single high-power optical cavity to high precision,
- 3. to explore the interactions between opto-mechanical effects and digital control loops maintaining cavity stability,
- 4. more specifically, to characterise the interactions between coupled opto-mechanical effects and control loops in a doubly-resonant coupled-cavity system.

4.2 The Glasgow 10 m Prototype

Several prototype interferometer laboratories exist across the world, intended as stepping stones between tabletop experiments and technological upgrades to full-scale gravitational wave detectors (e.g. the AEI Prototype [80]). In these facilities new technologies can be tested in systems dynamically-similar to the detectors (i.e. operating with suspended optics under vacuum), but in which such studies won't impact upon the time available for data-taking science runs.

The Glasgow Prototype Laboratory contains provision for several 10 m optical cavities housed in a sizeable vacuum system, as well as space for smaller in-vacuum and external experiments. Figure 4.1 presents a schematic of the vacuum system layout. The single 10 m cavity at the bottom of this figure is used for frequency-stabilisation of the main laser, itself situated on the optical bench external to the vacuum system. This is described in detail in Section 5.3.1.



Figure 4.1: Schematic of the Glasgow 10 m Prototype Laboratory vacuum system showing the main components of relevance to the work of this thesis. The upper 10 m arm contains the multiple spring experiment, while the lower arm contains a 10 m long reference cavity that is employed to stabilise the laser frequency.

4.3 MODELLING AND DESIGN STUDY

The design study confirmed the potential of the proposed doubly-resonant optomechanical system for the observation and control of coupled optical springs. The process began by modelling a pair of colinear 10 m Fabry-Pérot cavities coupled *via* a shared optic, termed the "central test mass" or CTM. This setup allows for a number of interesting regimes to be explored: in the first instance, the purely mechanical coupling between the detuned optical fields *via* their common interaction with the shared CTM, equivalent to the case of a single cavity with two carrier fields of different frequency (e.g. [64]); secondly, the more complex case where optical transmission through the CTM is permissible and leakage fields from one cavity can couple into the adjacent; and thirdly, the case where one or more sub-carriers of varying frequencies are injected into either cavity, forming a system of numerous coupled optical springs in which a great number of stability and sensitivity regimes may be explored. The author took responsibility for the design and implementation of the system, and investigation of the first regime detailed above.

Models were constructed using Optickle [81], a set of scripts designed to be run within MATLAB[®] [82] which together may be used to build frequency-domain models of interferometers or similar optical systems. Simulations were performed to produce observable coupled opto-mechanical resonant features in the measurement region of the system, expected to run from approximately 100 Hz to 1 kHz. These limits are explained in detail in Sections 4.5 and 5.3. In these simulations, the cavity lengths were both set to 10 m in anticipation of the available laboratory space. Each of the three cavity-forming optics was modelled as a suspended test-mass, with properties constrained by further mechanical simulations, discussed in Section 4.5.2. The input power was also constrained by the operation of a laser amplifier, constructed and tested by the author and set up for earlier optical spring work in the same laboratory [53]. The design and operation of this amplifier is discussed in Section 4.4. The remaining free parameters were thus the masses and optical properties of the suspended mirrors. Examples of simulations to inform the choice of these parameters are shown in Figure 4.2. In these simulations, transfer functions are plotted from a force at the input test mass of each detuned cavity to the position of the shared CTM while varying each parameter of interest, the resulting two-dimensional contour plots giving an overview of the dynamic response of the coupled opto-mechanical resonances.

Having formed a basis for the experiment with these simulations, the practical options of its commissioning were studied. This study indicated that the most optically efficient



Figure 4.2: Examples of early-stage investigations into the most important free experimental parameters. Each plot is scaled to depict many iterations of transfer functions from ITM to CTM motion in one of the two coupled cavities; brighter regions therefore indicate the location of coupled spring and optical resonances. Unless otherwise stated, the parameters used in this particular model are as follows: Cavity 1 input power = 3 W; Cavity 2 input power = 2 W; both ITM transmissions = 600 ppm; CTM transmission = 0; all three optic masses = 100 g; both detunings = 50 pm \simeq 2 linewidths. Note that because this model assumes no optical coupling between the two cavities the CTM transmission has been set to zero and all loss within the cavities is modelled as being associated with the appropriate ITM.

way to fit the two cavities into the available vacuum system was to fold each cavity, as is shown in the schematic in Figure 4.3. Folding each cavity requires the addition of an extra optic at 45° to the circulating fields. Since the folding optics are not perfectly reflective, light leaks out of each cavity; while this increases the losses in each, it also brings the advantage of providing probes of the intra-cavity power. As is shown in Chapter 5, these probes formed a key part of our control system. Having contacted optic manufacturers to obtain predictions of coating properties (e.g. reflectivity, losses), supplementary simulations were performed to estimate the optimum balance of properties for the five optics that form the two-cavity system.

Optics were manufactured by Layertec, Berlin. Following installation of these, the entire system was characterised. During the course of this characterisation, it was found that the maximum available input power to each cavity was between 1.5 and 2 W. As is shown



Figure 4.3: Simple schematic diagram of the experimental setup. A single 1064 nm source laser illuminates two 10 m optical cavities, folded into a single 10 m length of vacuum tube. The two cavities are formed by the Left Test Mass (LTM), Right Test Mass (RTM) and coupling Central Test Mass (CTM), with cavity-folding optics termed the Left and Right Steering Mirror (LSM and RSM). Pound-Drever-Hall (PDH) error signals are obtained from the reflected optical signal from each cavity and used to apply a feedback signal for precision control of each cavity length. This process is covered in detail in Section 5.2.

in Section 4.4, the laser amplifier is operated with approximately 5 W total output power, of which approximately 20% is picked off for frequency-stabilisation of the laser source, and the remainder split into the two MOSES input paths. Losses from components in the beam path prior to each cavity (due, for example, to imperfections in the optics and beam polarisations, beam clipping, mode-mismatch, etc.) then limit the power to the measured values. The final system parameters are shown in Table 4.1.

In Section 3.2.4 a simulation is presented of the expected quantum noise performance of a system similar to the MOSES experiment, indicating the potential for achieving a wider sub-SQL quantum sensitivity than would be possible with a single optical spring. However, due to the dominance of various other noise sources (with the limiting noise proving to be residual frequency noise on the source laser, as shown in Section 5.3.1) the quantum sensitivity of the MOSES experiment cannot be directly measured. Instead, measurements are performed inside the control loops of each cavity in order to confirm and characterise the presence of coupled opto-mechanical rigidity.

Parameter	Value	Description
$ au_{LTM}^2$	100 ppm	LTM power transmission [*]
$ au_{RTM}^2$	$100\mathrm{ppm}$	RTM power transmission [*]
R_{LTM}	$15\mathrm{m}$	LTM radius of curvature [*]
R_{RTM}	$15\mathrm{m}$	RTM radius of curvature [*]
\mathcal{F}_L	10220	Left-hand cavity finesse
\mathcal{F}_R	9155	Right-hand cavity finesse
γ_L	$734\mathrm{Hz}$	Left-hand cavity linewidth
γ_R	$820\mathrm{Hz}$	Right-hand cavity linewidth
P_{in}^L	$1.65\mathrm{W}$	Left-hand cavity input power
P_{in}^R	$1.85\mathrm{W}$	Right-hand cavity input power

Table 4.1: Updated system parameters after further modelling and refinement over longterm operation of the system. Suspended mass parameters (i.e. total mass, moments of inertia, etc.) may be found in Table 4.4. Values marked * are inferred from simulations; others are measured directly.

4.4 LASER AMPLIFIER IMPLEMENTATION

In order to achieve the high optical power required for the proposed investigations of optical springs, it is necessary to construct a power amplifier for the input laser field of the system. This section comprises a brief discussion of the theory behind laser amplification in terms of the underlying atomic processes; the construction of the amplifier in the laboratory; and characterisation of the typical gain and noise performance expected of such an amplifier.

4.4.1 LASER AMPLIFICATION

Transitions between different atomic energy levels may be stimulated by the application of an optical signal, subject to the requirement that the signal frequency be tuned to the desired transition frequency. For a simple system consisting of many identical atoms across two energy levels, this will cause atoms initially located in the lower state E_1 to jump up to upper state E_2 , and atoms initially in E_2 to fall down to E_1 . In such a system, the net rate of energy transfer to the atoms is as follows [83]:

$$\frac{\mathrm{d}U}{\mathrm{d}t} = Kn(t) \times [N_1(t) - N_2(t)] \times \hbar\omega , \qquad (4.1)$$

where n(t) is the number of photons per unit volume in the optical signal, N_2 the upperstate population, N_1 the lower-state population, $\hbar\omega$ the energy transferred from signal to atoms and K some constant determining the strength of the stimulated response. Accordingly, should the lower-state contain a larger number of atoms than the upper



Figure 4.4: Sketch of a simple three-level lasing atomic system. The pump beam excites the $E_1 - E_3$ transition. Atoms from E_3 then relax into E_2 , displaying population inversion with respect to E_1 . A signal beam tuned to the $E_2 - E_1$ transition frequency stimulates emission and is thus amplified upon propagation through the medium.

(i.e. $N_1 > N_2$), net flow of atomic population is to the upper-state and so energy is absorbed from the applied signal beam. If we wish to observe the opposite effect of net energy flow from the atoms to the signal beam - i.e. amplification of the optical signal the system must be arranged in the regime where $N_2 > N_1$. This is termed *population inversion*. This situation may be orchestrated by applying a secondary "pumping" beam to stimulate atoms into an upper energy level and hence to populate it more heavily than the lower level. Obviously this will require at least a three-level atomic system, such that the illuminating pump beam is of the correct frequency to stimulate atoms into an upper level (E_3 , say) from which they will relax over some characteristic time-scale into the intended upper lasing level E_2 . This lasing level then displays population inversion with respect to the lower level E_1 , allowing a signal beam tuned to the frequency of the $E_2 \rightarrow E_1$ transition to be amplified upon propagation through the atomic medium. This is illustrated for a simple three-level system in Figure 4.4.

Upon transit through a gain medium of length ℓ , the intensity of the signal beam is amplified by the *gain factor*, G [84]:

$$G = e^{g\ell} = \exp\left(\frac{g_0\ell}{1 + I/I_S}\right),\tag{4.2}$$

where $g_0 \ell$ is the *small-signal gain*, which determines the amplification of a beam the intensity of which, I, is much lower than the saturation intensity, I_S , of the gain medium (i.e. $I \ll I_S$). The small-signal gain can be calculated *via*

$$g_0 \ell = \frac{\eta_{\text{excit.}} P_{\text{pump}}}{AI_S} \,, \tag{4.3}$$

where P_{pump} is the optical power of the pump beam, A is the cross-sectional area of the active gain medium and $\eta_{\text{excit.}}$ the excitation efficiency, given as the ratio of available power in the upper-state of the system to the pumping power.

When constructing a physical amplifier in the laboratory with, for example, a single- or double-pass configuration, several other factors beyond these brief theoretical foundations will limit the practical gain that can be extracted from the system. The gain medium will have certain restrictive dimensions, along which the active medium as defined by the pump beam and the signal laser to be amplified must be precisely aligned. Further, many solid-state gain media (e.g. doped crystals or glasses) exhibit birefringence and polarisation-dependent gain, such that maximum amplification is only achieved when the polarisation of the signal laser is aligned with the correct axis when propagating through the gain medium. The mode sizes of both the pump beam and signal beam must also be optimised for factors such as overlap and relative size (ensuring that the signal beam effectively samples the active medium illuminated by the pump beam) and intensity (ensuring that saturation of the gain medium does not become a problem).

4.4.2 Amplifier Design

It was estimated that 5 W would be adequate total laser power to split as required between the two coupled cavities of the multiple spring experiment, as well as with other ongoing experiments in the same laboratory that share the same source laser. In order to achieve this, an optical amplifier was designed to amplify the existing 1 W neodymiumdoped yttrium aluminium garnet (Nd:YAG) seed laser¹, operating at a wavelength of 1064 nm. The gain medium chosen for this purpose was a birefringent, cylindrical crystal of neodymium-doped yttrium orthovanadate (Nd:YVO₄) of length 17.2 ± 0.2 mm, diameter 3 mm and with a $3.3 \pm 0.4^{\circ}$ wedge between the two faces aligned with the horizontal axis. This gain crystal is pumped with a diode laser operating at a peak wavelength of 814 nm with up to 100 W total output power. An image of the gain crystal atop the main optical bench in the laboratory is presented in Figure 4.5.

A sketch of the optical setup used for evaluating the amplifier gain is presented in Figure 4.6. The Faraday isolators and half-wave plates were arranged to ensure that the signal beam polarisation was horizontal through the crystal, to match the gain axis of the medium. This was confirmed by rotation of the second half-wave plate along the beam path through 90°; due to the birefringence of the crystal and the small wedge between the

¹Strictly, the source laser operates with 2 W output power; however, after losses in the necessary injection and beam-handling optics this is reduced to about 1 W input power to experiments.


Figure 4.5: Photograph of amplifier crystal on the main optical bench, set up for doublepass gain, whereby the beam passes through the gain crystal twice. The amplified beam is then extracted from the rear reflected port of a Faraday isolator back down the beam path.

front and rear surfaces, rotation of the wave-plate first results in two beams exiting the crystal at the diode-pumped end (one consisting of horizontally-polarised light and one of vertical) and then a single vertically-polarised beam, distinguishable from a horizontally-polarised beam by the relatively low gain applied to it upon propagation through the amplifier.

For double-passing of the signal beam through the amplifier, the "pump-through" mirror at the rear of the crystal is set normal to the beam, such that it reflects and passes a second time through the active medium, stimulating more emission at the lasing frequency and further amplifying the beam (up until saturation of the medium whereupon the upper-state is depleted). Passing a second time through the Faraday isolator adjacent to the gain crystal, the polarisation of the return beam ends up at 90° relative to that of the input beam, such that it is rejected from a polarising beamsplitter and directed onto a broadband photodetector for analysis. By measurement of the resulting photocurrent on the photodetectors (directly proportional to the power in the beam) before and after amplification, the gain of the amplifier may be determined at different pump diode powers.

During gain analysis the waist sizes and positions relative to the crystal of the pump and signal beams were optimised. The pump beam was chosen to have a waist of radius $560 \,\mu\text{m}$ at the closest face of the crystal, while the signal beam size was adjusted to be $325 \,\mu\text{m}$ for the seed laser, at a position as close to the most active section of the gain



Figure 4.6: Optical setup for double-passed gain measurements of the signal laser through the diode-pumped Nd: YVO_4 crystal. The signal beam passes twice through the gain crystal and is detected upon rejection from a Faraday Isolator. The half-wave plate and polarising beamsplitter at the output of the signal laser are used to reduce the input power to the amplifier if required, and the isolators and second half-wave plate are adjusted to ensure that the signal beam polarisation and gain axis of the crystal are aligned.

medium (around the pump waist) as possible. Figure 4.7 displays example measurements of the beam sizes. Note that the signal beam divergence is easily low enough to avoid clipping on the edges of the gain crystal.

4.4.3 GAIN MEASUREMENTS

Using the setup as described in Figure 4.6, the gain was measured by comparing the photodetector signals as measured on an oscilloscope with and without pumping of the gain crystal.

The overlap of pump and signal beams was optimised by setting the pump diodes to some low output power and making fine adjustments to the 45° steering mirrors along the signal beam path and to the pump diode fibre mount to maximise the signal observed at the photodetector. Once this was achieved and the other amplifier parameters (mode sizes, etc.) set up as described above, the pump power was increased in small steps, optimising the pointing of the pump fibre module at each step, and the signal reaching the photodiode noted and compared to the unamplified signal to measure the gain.

Gain measurements for the input signal beam operating with both $140 \,\mathrm{mW}$ and $1 \,\mathrm{W}$ input powers, and the corresponding output power of the signal beam after gain has been



Figure 4.7: Measured beam radii of pump and signal beams through the amplifier crystal, with theoretical Gaussian beam fits to the data yielding waists of approximately $560 \,\mu m$ for the pump beam and $325 \,\mu m$ for the signal beam. The signal beam fit is included in the upper plot for comparison with the pump beam waist size. In both plots, the crystal face is situated at the zero position and extends $17.2 \,mm$ to the right.

applied, are shown in Figure 4.8. The lower-power 140 mW measurements were included in order to evaluate the relative saturation of the gain medium with the intended 1 W source.

As shown in Figure 4.8, gains of 12.8 and 6.3 were achieved with input powers of 140 mW and 1 W respectively, when operating with 40 W of pump power. Operating at a pump power of about 40 W would then give approximately 6 W of output laser power for 1 W input, which ought to be sufficient input power for the optical spring experiments. These relative gains indicate that saturation of the gain medium is beginning to influence available amplification at these signal beam power levels. For an Nd:YVO₄ crystal, the saturation intensity is $I_s = 1260 \,\mathrm{Wcm}^{-2}$ [84]. For a beam waist in the crystal of ~ $300 \mu\mathrm{m}$, the signal beam intensities are as follows:

$$I_{140mW} = 50 \text{ Wcm}^{-2}$$

and $I_{1W} = 354 \text{ Wcm}^{-2}$. (4.4)



Figure 4.8: Double-passed gain and corresponding output power for 140 mW and 1W seed laser input to amplifier.

Looking at Eq. 4.2, if $I \ll I_s$ then G approaches the exponential of the small-signal gain, $g_0\ell$. However, as we increase the input power up to 1 W, the optical intensity incident on the crystal, I, becomes of comparable size to the saturation intensity, I_s , and so a lower gain will be achieved.

Also of note is the shape of the gain curves. From Eqs. 4.2 and 4.3, we would expect an exponential growth of gain with pump power. This effect is observed, but only at lower pump powers, whereas at higher powers the trend becomes more linear. It is thought that this could be a result of the doping level of the crystal, which influences the absorption of pump light. A brief gain measurement with a lower-doped crystal indicated that better higher-power performance could be achieved. However, the required 5 W of optical power can be adequately provided with the existing crystal and 1 W seed laser. It is in this setup that the amplifier was operated for the duration of this experiment.

4.5 Seismic Isolation and Suspension Design

4.5.1 Dynamics of Suspended Optics

Objects suspended as multi-stage pendulums, as the optics of a gravitational wave interferometer and those of our experiment are, benefit from seismic isolation of $1/f^2$ per stage over the range of frequencies between the rigid-body modes of the pendulum system and the internal modes of the constituent masses and springs. Care must be taken with their design in order that the rigid-body modes are limited to a lower-frequency band and the internal modes to higher frequencies, to create a resonance-free band in which measurements can be performed without excitation of these modes.

Prior to describing the suspension designs chosen for the optical spring experiments, we will briefly present the equations of motion that govern a generic suspended optic in the three major coordinates: longitudinal position x, pitch θ and yaw ϕ . These are the three dimensions with which the new suspension systems will be characterised, and in which they will be actuated upon.

Figure 4.9 depicts an optic suspended as a pendulum from four wires. The panels of the figure show how the coordinates are defined in order to derive the Lagrangian of the system. Figure 4.9a illustrates the three translational degrees of freedom (x, y, z) and two angular: pitch, θ , and yaw, ϕ . Figure 4.9b shows the same mirror in the x-z plane being acted on in the positive x direction with the result that it shifts from the equilibrium position x_{sp} to position x. The suspension wires, of length ℓ , tilt at angle α to the vertical when the longitudinal displacement acts. The break-off height of these wires above the centre of mass of the optic is denoted b. Figure 4.9c then shows the optic in the x-y plane being acted on in yaw, with R_1 and R_2 denoting the radii of the suspension and break-off points respectively from the centre of mass, with each of these having been rotated by ϕ_{sp} and ϕ respectively. Using these definitions, the Lagrangian for the system may be expressed as [50]:

$$\mathcal{L} = \frac{1}{2}m\dot{x}^{2} + \frac{1}{2}I_{\omega}\dot{\theta}^{2} + \frac{1}{2}I_{\phi}\dot{\phi}^{2} - \frac{mg}{2\ell}(x - x_{\rm sp} - b\theta)^{2} - \frac{1}{2}mgb\theta^{2} - \frac{1}{2\ell}mgR_{1}R_{2}(\phi - \phi_{\rm sp})^{2}.$$
(4.5)



Figure 4.9: (a) shows a schematic diagram of a mirror suspended from four wires, indicating the three translation coordinates (x, y, z) and the two important angular ones (θ, ϕ) . Note that rotation about the x axis has little (or, in the ideal case, no) effect on the light beam and so is dropped here. (b) shows the motion of the mirror when acted upon in the longitudinal (x) direction, such that the pendulum is shifted from its equilibrium position x_{sp} to position x and the suspension wires shift to angle α at the suspension point. The wire break-offs are separated from the centre of mass (denoted by a red circle) by distance b. Note that longitudinal motion naturally couples in pitch motion with angle ϕ . (c) shows the motion of the mirror when acted upon in the yaw degree of freedom. Suspension points have rotated through angle ϕ_{sp} and break-off points by ϕ . Radius of suspension point and break-offs from the centre of mass are denoted by R_1 and R_2 respectively.

The individual equations of motion in each of the three degrees of freedom considered may then be determined using the Euler-Lagrange identity,

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{q}_i} \right) = \frac{\partial \mathcal{L}}{\partial q_i} \,, \tag{4.6}$$

where each q_i represents one of the coordinates x, θ or ϕ . Including a viscous damping term, $\zeta_i \dot{q}_i$, to represent the friction present in each degree of freedom, the equations of

motion are obtained:

$$\ddot{x} + \zeta_x \dot{x} + \omega_x^2 x = \omega_x^2 (x_{\rm sp} + b\theta) \tag{4.7}$$

$$\ddot{\theta} + \zeta_{\theta} \dot{\theta} + \omega_{\theta}^2 \theta = \frac{\omega_{\theta}^2}{l+b} (x - x_{\rm sp})$$
(4.8)

$$\ddot{\phi} + \zeta_{\phi} \dot{\phi} + \omega_{\phi}^2 \phi = \omega_{\phi}^2 \phi_{\rm sp} \,. \tag{4.9}$$

Here, each ω_i is the natural mode frequency of each q_i :

$$\omega_x^2 = \frac{g}{\ell} \tag{4.10}$$

$$\omega_{\theta}^2 = \frac{mg}{I_{\theta}\ell}b(\ell+b) \tag{4.11}$$

$$\omega_{\phi}^2 = \frac{mg}{I_{\phi}\ell} R_1 R_2 . \qquad (4.12)$$

These describe the rigid body mode frequencies in these three key degrees of freedom for a single pendulum-suspended mass. For a two-stage suspension, there are twice as many modes (and, indeed, n times as many modes for an n-stage suspension). Clear derivations of these, as well as the further suspension modes in such degrees of freedom as roll about the x-axis and vertical which will be enumerated in this thesis, may be found in [50] and [85], amongst many others. A list of mode frequencies for the suspensions described in the following chapter is given in Table 4.3.

4.5.2 Design of Suspensions and Pre-Isolators

Figure 4.10 shows a measurement of the background seismic noise in the laboratory. The design of new suspensions began with the requirement that the test masses be sufficiently isolated from this seismic motion at 100 Hz, taken to be the lower limit of the measurement band for the experiment (based upon previous experience with low-noise measurements in the lab). This imposes the further requirement that the rigid-body modes of the suspensions ideally be limited to below 100 Hz. A common approach is to use some passive "pre-isolation" stages in conjunction with the minimum number of suspension stages required to meet the desired isolation at the lower end of the measurement band, since seismic noise falls with frequency. As the experimental design requires several suspensions to be housed in single vacuum tanks, they must be reasonably compact. Several possible designs of suspension and pre-isolation stages were considered, with two suspension stages with pre-isolation proving a good compromise between performance, size and complexity. Frequency noise of the free-running laser is stabilised to give approximately $10^{-14} \text{ m}/\sqrt{\text{Hz}}$ through the measurement region, such that seismic noise must be



Figure 4.10: Seismic noise in the laboratory in the horizontal plane, measured and calibrated using geophones (blue) and filtered as $1/f^4$ from 0.69 Hz as per the predicted pendulum isolation (red).

below this level at 100 Hz. The lower spectrum in Figure 4.10 shows the expected seismic noise at the test-mass after filtering by a two-stage pendulum, indicating that it will be comfortably below this level. Pre-isolation is included to achieve some extra margin.

As described in Section 4.5.1, a single pendulum stage acts as a $1/(\text{frequency})^2$ lowpass filter above its resonant frequency for transmission of displacement from the pinning point to the bottom of the pendulum. This indicates that a two-stage pendulum, which thus provides $1/(\text{frequency})^4$ isolation from seismic noise, would be adequate to meet our beginning constraints. The decision was also made to build every suspension to the same basic design: a stainless steel base and top plate separated by four aluminium supporting legs. A heavier steel plate anchors the suspension wires, itself separated from the top plate by another small layer of rubber stacks for further isolation. The pendulums inside each 1 m diameter vacuum tank are situated as required atop a 60 cm square optical breadboard. This is raised to appropriate height above the tank floor for ease of access when positioning in the beam path and connecting signal cables and to situate each optic so that it is vertically centred relative to the vacuum tube connecting adjacent tanks. Between the tank floor and breadboard are two layers of passive seismic isolation, formed



Figure 4.11: Sketch of the seismic isolation apparatus in each vacuum tank. The twostage pendulum structure sits atop a 600×600 mm breadboard, itself isolated from ground motion by two layers of rubber stacks.

by supporting rubber blocks, that will provide a little extra damping of seismic motion above their resonant frequencies. A sketch of the full seismic isolation apparatus is shown in Figure 4.11.

The rubber stacks were separated into two standard sizes based upon available height between isolated layers; the lower two stages beneath the optical breadboard are formed of a number of individual blocks of size $16 \times 16 \times 10$ mm, while those supporting the top anchor plate are a smaller $9 \times 9 \times 6$ mm. Tests were performed on each block size with various mass loads to determine the resulting resonant frequency, which then allows rough tuning of the resonant frequency of each layer by careful choice of the total mass load and number of individual stacks per layer. Figure 4.12 shows the results of these tests to determine the vertical resonant frequencies of each block. These tests were performed by loading sets of four stacks of each size with various weights and exciting the structure while reading out the voltage from an accelerometer adhered to the topmost weight with



Figure 4.12: Tests of rubber stack resonant frequencies. Upper figures show accelerometer spectra for various mass loads upon sets of four sets of (a) smaller and (b) larger stack sizes. (c) and (d) display approximate trends between resonant frequency and mass load for each stack size respectively.

a spectrum analyser (Stanford Research model SR785). Resonant frequencies were then estimated from the peaks in the resulting voltage spectra, and approximate trends were evaluated to determine the dependence of resonant frequency on mass load.

Table 4.2 displays the final stack numbers and mass loads per layer in each vacuum tank. The total mass figures include estimates for the suspensions. The final mass of each suspension depends on the design details discussed in Section 4.5.3. Note, however, that the heaviest parts are the frame and test-masses themselves, which do not vary significantly.

Structure	Isolation Layer	Number of	Mass Load per	Vertical Mode
		Stacks	4 Stacks (kg)	Frequency (Hz)
LTM, RTM and	Upper	4	2.8	29.5
CTM suspensions				
LSM and RSM	Upper	4	3.2	29
suspensions				
Input Tank	Middle	8	12.4	25
	Lower	6	20.9	21.5
End Tank	Middle	8	14.9	24
	Lower	6	24.2	20.5

Table 4.2: Rubber stack configuration and approximate vertical resonant frequencies in the two cavity-housing vacuum tanks. Each isolation layer may be identified with reference to the sketch in Figure 4.11.

4.5.3 SUSPENSION DESIGN AND CHARACTERISATION

As described in Section 4.5.2, the five suspensions of the MOSES experiment are required to be two-stage pendulums to meet our seismic noise requirements. These suspensions should be relatively simple in design, and of a convenient size that they can be re-purposed for future work. With reference to Eq. 4.10, a total pendulum length of 60 cm gives each suspension a suitably low fundamental longitudinal mode frequency of 0.69 Hz, as well as allowing for ease of transport and positioning throughout the lab. The mass of each component suspended by a double-pendulum is determined in Section 4.3, yielding 100 g for the three principal cavity-forming optics (these being the LTM, RTM and CTM, with reference to Figure 4.3) and a heavier 300 g for the two intra-cavity folding optics (LSM and RSM). These designs were then realised in SolidWorks[®] (3D computer aided design software) to give accurate estimates of total masses and moments of inertia about the three dimensions of each mass.

With the pendulum lengths and masses now fixed, simulations were performed using Simulink[®] in MATLAB[®] to determine the number, length, material, thickness and position of suspension wires. This results in appropriate positioning of suspension modes in frequency. These simulations were adapted from an existing dynamical model of a three-stage pendulum developed by C. Torrie [85] and are based upon the pendulum dynamics introduced in Section 4.5.1. The important suspension parameter decisions arising from this modelling are presented over the following tables: Table 4.4 shows the total masses and moments of inertia for both the upper and lower masses in each of the five cavity-forming suspensions, the wire separation in the x (perpendicular to the optic face) and y (parallel, horizontal) axes for each mass, and the suspension wire properties decided

Suspension	Mode	Frequencies (Hz)	
	Longitudinal	0.69, 1.69	
	Pitch	3, 20.9	
CTM	Yaw	1.18, 3	
	Vertical	33.1, 94.9	
	Roll	43.8, 128	
LTM & RTM	Longitudinal	0.69, 1.69	
	Pitch	2.96, 20.7	
	Yaw	0.9, 2.3	
	Vertical	33.2, 95.8	
	Roll	56, 164	
LSM & RSM	Longitudinal	0.69, 1.71	
	Pitch	2.93, 12.2	
	Yaw	1.07	
	Vertical	19, 54.7	
	Roll	29.8, 89.2	

Table 4.3: Mode frequencies in five degrees of freedom for each of the five cavity-forming suspensions. The missing degree of freedom, transverse horizontal motion, does not couple to the measurements. Its mode frequencies are approximately those for the longitudinal degree of freedom.

upon; Table 4.3 indicates the calculated suspension mode frequencies for each of the three models of cavity-forming suspension.

After construction of the suspension structures, pendulum modes could be examined physically in the lab and compared to the predictions of the models in Table 4.3. This was achieved by injecting a band-limited white noise signal on to the suspension using electromagnetic actuators, the design and operation of which are covered in Section 4.5.5. The longitudinal response of the pendulum to the injected noise was read out using a Polytec OFV-505 vibrometer. The vibrometer provides velocity information, which must be integrated to produce the resulting displacement response spectrum. Figure 4.13a displays this measurement for the central test mass (CTM), together with the predicted response from the Simulink[®] modelling. Note the coupling of angular modes into the longitudinal response; the extent to which this occurs will depend on the position of the vibrometer probe beam relative to the centre-of-mass of the optic, although as identified in Section 4.5.1 pitch modes will couple into longitudinal modes. The model considers longitudinal to pitch coupling only; any yaw coupling will be evident only in the measurement and ought to arise solely from vibrometer alignment. The angular response of the pendulum was measured by reflecting a low-power laser from the surface, while once again exciting pendulum modes with white noise injected *via* actuators, and reading out

Part	Parameter	Value
	Material	Stainless steel AISI 302
~	Diameter	$100 \mu\mathrm{m}$
Suspension Wire	Wires in upper-stage	2
	Wires in lower-stage	4
	Upper-stage length	$0.25\mathrm{m}$
	Lower-stage length	$0.35\mathrm{m}$
	Mass	98 g
	I_x	$35.1 \mathrm{kg/mm^2}$
LTM/RTM Upper Mass	I_{η}	$21.1 \mathrm{kg/mm^2}$
	I_z	$20.2 \mathrm{kg/mm^2}$
	Wire separation (y)	$50.8\mathrm{mm}$
	Mass	99.3 g
	I_x	$38.9\mathrm{kg/mm^2}$
LIM/RIM lest Mass	I_y	$23.1 \rm kg/mm^2$
	I_z	$23.9\mathrm{kg/mm^2}$
	Wire separation (x)	$7\mathrm{mm}$
	Wire separation (y)	$50.8\mathrm{mm}$
	Mass	98 g
CTM Upper Maga	I_x	$35.1\mathrm{kg/mm^2}$
CTM Opper mass	I_y	$21.1 \mathrm{kg/mm^2}$
	I_z	$20.2\mathrm{kg}/\mathrm{mm}^2$
	Wire separation (y)	$50.8\mathrm{mm}$
	Mass	$102\mathrm{g}$
CTM Test Mass	I_x	$37.5\mathrm{kg}/\mathrm{mm}^2$
	I_y	$22.0\mathrm{kg/mm^2}$
	I_z	$22.7\mathrm{kg/mm^2}$
	Wire separation (\mathbf{x})	$7\mathrm{mm}$
	Wire separation (y)	50.8 mm
	Mass	301.7 g
LSM/RSM Upper Mass	I_x	$11.1 \mathrm{kg/mm^2}$
	I_y	$68.8 \mathrm{kg/mm^2}$
	I _z	$58.1\mathrm{kg/mm^2}$
	Wire separation (y)	50.8 mm
	Mass	302.7 g
LSM/RSM Lower Mass	I_x	$12.5 \mathrm{kg/mm^2}$
	I_y	$84.8 \mathrm{kg/mm^2}$
		$77.6 \mathrm{kg/mm^2}$
	Wire separation (x)	$7\mathrm{mm}$
	Wire separation (y)	$50.8\mathrm{mm}$

Table 4.4: Wire and suspended mass properties required to calculate suspension modes. Note that for the test masses, calculated values include the fused silica optics housed in each suspension (with two housed in the CTM test mass). Moments of inertia and masses are taken from the SolidWorks[®] design models and rounded to three significant figures.

differential signals in the x- and y-axis using a quadrant photodiode. Fig 4.14b presents the angular response of the CTM in yaw and pitch compared to the predictions of the model. Again, extra mode couplings not present in the model are due to beam alignment in the experiment.

The mechanical transfer functions for the other two cavity optics, the LTM and RTM, are assumed to be the same as those for the CTM to the precision that we require. The two heavier cavity folding optic suspensions were modelled and tested, with a threefold increase in suspended mass and slightly modified dimensions, yielding very similar results. With the cavity optics suitably characterised, we now consider the control methods for each individual suspension and for each cavity as a whole.

Figure 4.14 shows photos of the completed suspensions and parts following installation inside the vacuum system.

4.5.4 Actuation and Damping

Actuation is applied at two levels in each of the three cavity-forming suspensions. At the intermediate mass, control signals are applied to adjust the position of that mass, and hence the test mass suspended therefrom, in pitch and yaw; this is termed *local control*, used for independent precision alignment of each suspension. *Eddy current damping* of suspension motion is also applied using the local control apparatus, as is described in the following section. The actuation at the test mass itself forms a crucial part of the feedback system used to lock each cavity to resonance. Here, control signals are applied longitudinally to adjust the overall length of each cavity; this is termed *global control*, and is the method by which the cavity length is maintained. The two cavity folding suspensions (the LSM and RSM) and all of the beam steering suspensions external to the cavities are subject to local control actuation only.

In both stages, actuation is achieved electromagnetically. Four coils are mounted on the suspension frame, directly behind the mass being actuated upon. Four magnets are bonded onto the rear of each mass in the same formation, such that a current driven through each coil will impart a force to the adjacent magnet.

With the coils and magnets arranged in a diamond formation, actuation may be applied in one of three degrees of freedom: pitch, by differentially actuating upon the top and bottom magnets; yaw, by differentially acting upon the left and right magnets; and



Figure 4.13: Mechanical transfer functions and corresponding models for three degrees of freedom of the CTM: (a) longitudinal; (b) pitch and yaw.

CONCEPTUAL EXPERIMENTAL DESIGN AND IMPLEMENTATION



(c)





(e)



Figure 4.14: Pictures of suspensions: (a) Model of the CTM suspension in SolidWorks[®]; (b) the LSM and CTM being installed in the vacuum system; (c) top steel plate at which suspension wire is anchored; (d) intermediate mass with local control coils and eddy current dampers; (e) test mass with global control coils; (f) view of the LSM through vacuum system porthole.



Figure 4.15: Left: Force per unit current produced between a single 3 mm diameter by 4 mm thick cylindrical magnet and 80-turn, radius 3.5 mm coil. Displacement is measured from the centre of the 4 mm-long magnet. Right: Coil former dimensions.

longitudinal, by acting upon all four in phase. A Mathematica[®] [86] model written by M. Barton to numerically evaluate the magnetic field at arbitrary distances from a coilmagnet system was used to calculate the variation of force with distance for given coil and magnet dimensions. The magnets used are cylindrical with diameter 3 mm by 4 mm thick, with adjacent coil of 80 turns, inner radius 3.5 mm and length 4 mm; using these parameters in the model, we arrive at the force-displacement plot shown in Figure 4.15. This simulation places the point of optimum (i.e. maximum) force at approximately 1 mm in front of the coil edge, where the force per unit current is given as 44 mN/A. With reference to the coil-former dimensions provided in Figure 4.15, this means optimum magnet position is lined up with the front edge of the coil-former. At this point, we expect minimal transmission of actuator noise to the masses, as the gradient of force with displacement is close to zero near the maximum of the force.

4.5.5 LOCAL CONTROL

Since in the MOSES suspensions local control actuation is used for DC alignment of each suspension (further to the passive damping described below), it does not form part of the feedback loops to each cavity. Rather, it is subject only to the requirements that (a) no alignment change be applied instantaneously, as a large impulse would excite most of the suspension modes to some degree; (b) they be stable in applied current, low in noise and reliable in maintenance of alignment; and (c) that they actuate to adequate angular resolution for precision alignment of both cavities.

With each of the three test masses being nominally 5 cm in diameter and situated 10 m apart, the angular adjustment range was chosen to be $2\xi = 10 \text{ mrad}$, where ξ is the horizontal or vertical angle away from the cavity axis, yielding $\pm 5 \text{ cm}$ spot positioning range at 10 m. This provides ample range to scan completely across the diameter 2.54 cm optic surface and to steer completely off the mass if necessary (for example, if complete misalignment of one cavity is desired).

With this range in mind, we can calculate the force required from the two yaw actuators to balance the restoring force of the pendulum in that degree of freedom:

$$F_{\rm yaw} = \frac{mgt\xi}{\ell} \,, \tag{4.13}$$

where m is the total suspended mass (since local control is applied at the intermediate stage), t is the radius of actuation from the centre-line of the suspension and ℓ the suspension wire length. Then, with the knowledge that the actuator force per magnet is 44 mN/A at the optimum coil-magnet distance, the required current is calculated to be $I_{\text{yaw}} = 11.3 \text{ mA}$. Repeating this process for the pitch axis, we find that $I_{\text{pitch}} = 16.8 \text{ mA}$.

The coil driver control signals are supplied by an Arduino microcontroller [87], such that alignments can be specified *via* a web server that instructs the Arduino to apply appropriate output signals. These signals are applied using *pulse width modulation* (PWM), a technique that allows analogue signals to be mimicked by way of modulation of the width of a square-wave pulse. Appendix A.2 presents a circuit diagram of the coil driver between Arduino outputs and local control coils. Using two different PWM channels allowed both a "coarse" and "fine" control to each degree of freedom.

When applying some constant current to maintain alignment, care must be taken that the current driver to each coil includes a low-pass filter of appropriate order and corner frequency to ensure current noise resulting from the noise on each Arduino PWM output doesn't impair the displacement sensitivity of the cavity in the measurement frequency range above 100 Hz. The voltage noise on a single Arduino PWM output was measured to be approximately $0.8 \,\mu\text{V}/\sqrt{\text{Hz}}$ at 100 Hz, and assuming that each coil has a nominal resistance of $4 \,\Omega$, this translates into a current noise of $0.2 \,\mu\text{A}/\sqrt{\text{Hz}}$ and hence force noise at the optimum coil-magnet separation of $8.8 \,\text{nN}/\sqrt{\text{Hz}}$. Converting into an equivalent displacement due to a single coil and filtering by the pendulum stage, this translates into $1.6 \times 10^{-17} \text{ m/}\sqrt{\text{Hz}}$ at 100 Hz (being the lower bound of the measurement region). Thus if a third-order low-pass filter is included in the coil driver circuitry at 10 Hz, this further reduces to $1.6 \times 10^{-20} \text{ m/}\sqrt{\text{Hz}}$. As we shall see in Section 5.3.1, the system is designed to be limited by frequency noise of the laser to an equivalent displacement noise level of approximately $10^{-14} \text{ m/}\sqrt{\text{Hz}}$, putting the Arduino output contribution to displacement noise comfortably below this limit. The 10 Hz limit to alignment signal frequency allows for a short ramp between each subsequent suspension alignment, ensuring that modes aren't excited to a prohibitive level.

The pitch and yaw ranges were confirmed by reflecting a visible laser source from a suspended optic, adjusting its alignment *via* the local control and observing the spot position at 10 m away. This indicated full coarse alignment ranges of ± 6.1 mrad in pitch and ± 4.3 mrad in yaw. Testing the fine control using a quadrant photodetector to ascertain accurate beam position, we note an alignment resolution limit at 10 m of 5 μ m in pitch and 7 μ m in yaw. These ranges and resolutions are more than adequate for precision alignment of both cavities during the experiment.

As noted previously, the same apparatus used for the local control is also used for damping of suspension noise. The coil-former used in the intermediate stage coils is made of copper; when the associated magnet moves through the coil former, the Lorentz force acts upon electrons in the surrounding copper to generate eddy currents. These current loops themselves induce a magnetic field which tends to oppose that which created them, resulting in a repulsive force and hence viscous damping of magnet motion. This effect is termed *eddy current damping* [88] and may be characterised by observation of any of the pendulum resonances shown in Figure 4.14b. When characterising the damping observed, we compare to the results of finite element analysis simulations for coil-formers of the same dimensions for use in the AEI Prototype system given in a technical paper by W. Cunningham [89]. The damping is expressed in the form of the damping coefficient, κ , corresponding to the viscous term in the system's equation of motion:

$$m\frac{d^2x}{dt^2} + \kappa\frac{dx}{dt} + kx = 0.$$
 (4.14)

The fundamental frequency ω_0 and damping radio ζ are given by

$$\omega_0 = \sqrt{\frac{k}{m}} \tag{4.15}$$

and
$$\zeta = \frac{\kappa}{2\sqrt{mk}} = \frac{1}{2Q}$$
, (4.16)

where Q is the quality factor of the resonance, corresponding to the ratio of the fundamental frequency and full-width at half-maximum. In the angular transfer functions given in 4.14b, the yaw mode has Q = 21.5. Rearranging for the damping coefficient, we find that

$$\kappa = \frac{m\omega_0}{Q} = 34.6 \times 10^{-3} \,\text{kg/s per coil.}$$
(4.17)

This compares favourably with the result of $\sim 0.03 \text{ kg/s}$ per actuator (or $\sim 0.12 \text{ kg/s}$ total) given for the AEI Prototype coil modelling in [89].

4.5.6 GLOBAL CONTROL

The global control uses the same coil-magnet actuators as the local control, but without the inclusion of damping, which is not required at the lower stage as it may increase the thermal noise and, more importantly, "short-circuit" the isolation of the suspension by coupling the mirror to the support. For this reason, the coil-formers are made of moulded plastic rather than copper. Three four-channel drivers are provided for the global actuators – one for each of the cavity optics, with one channel per coil – although in the course of feeding back signals to lock each cavity only the relevant input coupler (i.e. the LTM or RTM) will have global control signals applied.

A diagram of the driver circuitry is provided in Appendix A.3. Since shaping of the feedback signal is achieved elsewhere in the loop (see Section 5.3.5), this circuit is predominantly designed for routing of signals to appropriate coils. The output voltage of the coil driver is applied to each individual coil in series with a 500 Ω resistor to provide the coil current. This resistor value was chosen as the final-stage amplifiers in the circuit can can drive a maximum voltage of ± 15 V and maximum current of $30 \text{ mA} - a 500 \Omega$ resistance thus provides the maximum dynamic range between output voltage and current. Note that as well as longitudinal control, the circuit also allows for the application of pitch and yaw signals to the actuators at the test mass. These paths are not used for locking purposes, since for cavity lock we need only control the total length of each; rather, these channels can be used to apply *auto-alignment* signals, forming part of a secondary control loop to maintain good alignment of each cavity over time (see Section 5.4).

4.6 INSTALLATION

During the process of installing and commissioning, a number of measurements had to be performed to ensure proper operation of the MOSES apparatus. These are described in the following sections.



Figure 4.16: (a) shows an exaggerated sketch of the ray-optic layout for testing the wedge angle between the CTM optics. The wedge angle is determined from the pattern of spots reflected from each surface. (b) shows a section through the CTM mount with the two optics in place.

4.6.1 Determination of CTM Wedge Angle

One requirement imposed in order to simplify the experiment was that the two cavities be coupled mechanically only; i.e. resonating fields from one cavity leaking through the CTM should not be allowed to couple into the adjacent cavity. To this end, the two optics used in the CTM were both designed to include a 0.5° wedge angle between the front high-reflection (HR) coated surface and the rear anti-reflection (AR) coated surfaces. In principle this ensures that any transmitted light from one cavity will end up too far off-axis from the ideal cavity mode to resonate in either; however, in practice the particular relative orientation of the optics, as well as the tolerances on the machining of the aluminium mount in which they are housed, could result in partial optical coupling between the two cavities. After installation of the new suspensions, optics and supporting structures the total wedge angle between the two optics, and thus the relative angle between each resonating field and the leakage field from the adjacent cavity, was tested. Figure 4.16 shows a sketch of the layout of this test and the positioning of the two optics within the CTM.

The test was performed using a visible helium-neon (HeNe) laser directed normal to one side of the central test-mass, since at this wavelength (633 nm) both the HR and AR surfaces of each optic are reflective enough for measurements to be performed. Since both wedged optics are housed back-to-back in the CTM, they effectively form three short etalons; from the pattern of reflected HeNe beams observed and using simple ray optics, reasonable estimates can be made of the relative orientations of each optic in the suspension, the relative wedge between the two AR surfaces of each, and the wedge angle of the individual optics themselves.

Using this method, the wedge angle of the front optic was determined to be $(0.523 \pm 0.028)^{\circ}$. Making the assumption (based on the expected precision of manufacture) that both mirrors have the same wedge angle then allows the total vertical angle between the two high-reflective surfaces to be estimated to be $(1.275 \pm 0.040)^{\circ}$. For a beam incident normal to one optic, then, we can be quite confident that any transmitted beam will be far enough off-axis in the adjacent cavity to avoid optical coupling into the cavity mode.

4.6.2 MODE MATCHING

Optimal coupling of the input light into an optical cavity occurs when the curvature of the wavefront at the position of each mirror is equal to the curvature of that mirror. If the curvatures of the two mirrors are known, the size of the input beam can thus be designed in such a way as to couple the optimum amount of power into the resonator [90]. As a Gaussian TEM₀₀-mode beam propagates, its radius, w(z), at position z relative to the narrowest point in the beam (the beam waist, w_0) is determined by [48]

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2},$$
 (4.18)

where the beam radius is defined as the distance from the propagation axis at which the intensity has dropped to $1/e^2$ of its maximum value and z_R is the *Rayleigh length*, defined as the distance from the beam waist at which the beam radius increases by $\sqrt{2}$ (and thus the beam area doubles):

$$z_R = \frac{\pi w_0^2}{\lambda} \,. \tag{4.19}$$

At the beam waist, the wavefront is flat. At distance z, the curvature, R, of the wavefront is

$$R(z) = z \left[1 + \left(\frac{z_R}{z}\right)^2 \right] \,. \tag{4.20}$$

As shown in Section 4.3, both cavities of the multiple optical spring system have a flat end optic and R = 15 m curved input optic, with a cavity length of z = 10 m. Substituting these parameters into Eq. 4.19 and 4.20, we determine that the required beam waist at the end optic must be $w_0 = 1.54$ mm for optimal coupling of power into each cavity. To achieve this goal, it is convenient to propagate this ideal mode-matched beam



Figure 4.17: JamMt GUI displaying the final mode-matching solution that works for both cavities.

back through the vacuum system to the main optical bench in the laboratory, where a reasonable starting waist size can be defined and implemented.

The Gaussian beam modelling tool JamMt [91] was used to model the beam radius as it is propagated from the optic bench to the CTM. Since the distance the two input beams travel to the left- and right-hand cavities has been designed to be identical and the cavities are nominally the same length, a single mode-matching solution should suffice for both beams. Figure 4.17 displays the JamMt graphical user interface and mode-matching solution arrived at for each cavity. This solution requires a beam waist on the optic bench of 680 μ m at a position 1.17 m prior to injection into the vacuum system, yielding a final waist at the CTM of 1.55 mm. Translation of initial to final beam waist is achieved by locating a relatively weak lens in a vacuum tank 6.5 m after injection into the system. This solution yields a manageable and easily measured input waist which can be achieved prior to the splitting of the main beam into the two separate input paths to each cavity, such that only one waist need be defined and adjusted to match into both cavities.

A series of lenses was added to the beam path on the optical bench in order to translate from an initial beam waist of approximately $100 \,\mu\text{m}$ further back down the beam path into the desired $680 \,\mu\text{m}$ waist. A beam profiler was then used to sample the radius of the



Figure 4.18: Measured beam radii in horizontal and vertical axes, and corresponding least squares fits to the data.

beam in vertical and horizontal axes at a number of positions along the propagation axis, such that fitting an ideal Gaussian beam to the data using Eq. 4.18 and 4.19 would yield an accurate estimate of the beam waist and position. Figure 4.18 displays the results of these measurements, yielding estimates of the beam waists in the horizontal and vertical axes of $692.5 \pm 2.9 \,\mu\text{m}$ and $685.6 \pm 8.6 \,\mu\text{m}$ respectively, where the tolerances on each result are specified as the 95% confidence bounds of the least squares fit to the datasets. The zero of the x-axis scale is taken relative to a prior optic on the bench, placing the desired waist position at approximately $z = 0.2 \,\mathrm{m}$ on this scale. The measured waist sizes are adequate for our mode-matching requirements, differing only on the percent level from the desired waist size and position.

4.6.3FINESSE MEASUREMENT

Equation 3.33 describes the beam reflected from a two-mirror cavity in terms of the input field only; here the reflected beam is shown to consist of that portion of the input field reflected directly from the rear of the input test mass, plus that which has resonated inside the cavity, accruing some phase shift, before leaking out through the input mass. By careful examination of the DC signal obtained from a photodetector monitoring this reflected beam a lower limit may be placed on the finesse of each cavity.

The storage time of the cavity is given as [92]

$$\tau_s = \frac{2\mathcal{F}L}{\pi c} \tag{4.21}$$

and is defined as the time required for a decay of the cavity field to 1/e of its maximum value on resonance. As the frequency of the in-going light approaches and passes the resonant frequency of the cavity, beats at the difference frequency are observed when the reflected light is detected, partially due to the cavity field decaying with time constant τ_s beating with the directly reflected and transmitted fields. It can further be shown that these oscillations become larger if the length or frequency of the cavity is modulated by more than half of the cavity linewidth during one storage time [93, 94]. For practical purposes, this means that examination of the exponential decay of the oscillatory behaviour observed on the reflected DC from the cavity allows estimation of the cavity finesse, as per Eq. 4.21 above.

Figure 4.19 shows this measurement for each cavity. A fringe is observed on the reflected DC trace, corresponding to a drop in reflected power occurring when the cavity passes through the resonant point. An exponentially-decaying oscillation is observed as the resonance decays (often referred to as a cavity *ringdown*), to which an exponential fit is applied to ascertain τ_s and hence \mathcal{F} . These measurements are typical of several measurements performed of cavity ringdowns to estimate the finesse; averaging five measurements for each cavity yields finesse estimates of 9,155±122 for the right-hand cavity and 10,220±164 for the left-hand cavity.

As we shall show in Chapter 5, the oscillatory behaviour which is so useful for determining the finesse of each cavity can also cause difficulties in acquisition of lock when feedback is applied to each cavity. The same oscillations present on the reflected signals will also be present on the error signals used to construct the feedback to the optics for position control; these oscillations tend to act on a time-scale much faster than the digital feedback loops can act to compensate them, such that a novel modification of the acquisition systems is required in order that this problem be mitigated.



Figure 4.19: Measurements of the exponentially-decaying oscillations present on the reflected signal derived from each cavity. The raw data (purple) is first smoothed with a rolling average (blue) before peaks and troughs are identified and fitted to exponential decays (red). Exponential fits to the data in this instance indicate finesse estimates of 10,613 for the left-hand cavity (a) and 9,438 for the right-hand cavity (b).

V Digital Control Systems Design

In order to carry out the proposed investigations into characterisation and control of a coupled-optical-spring system, a measurement system consisting of two cavities with a shared end test mass was constructed, as described in Chapter 4. Both cavities of the system must be locked to the point of resonance, such that a large amount of optical power (of order some kilowatts) builds up between the optics that form each cavity. Each is then individually detuned from this locked state to create the opto-mechanical coupling that we intend to measure. To achieve stable operation with a wide detuning range, robust feedback and control loops must be implemented for acquisition and maintenance of lock. It is inside these control loops that measurements to confirm the opto-mechanical couplings in the system will be performed. The principles employed here have been applied to lock cavities many times in previous experiments. Although traditionally analogue feedback circuitry has been employed, more recently this has given way to digital control systems in most of the advanced gravitational wave detectors and prototypes currently operating or upgrading (e.g. LIGO [18] and the AEI 10 m Prototype [80]).

To meet the goals of the proposed set of experiments, it was necessary to ensure that a wide range of detunings could be achieved with the cavity parameters described in the previous chapter. This required the development of new techniques, facilitated by digital control. In this chapter we discuss the operational challenges and the ways in which they have been overcome to arrive at a digitally-controlled system capable of locking each cavity and detuning over a range not before attempted for a system with suspended masses and cavity lengths on this scale. These challenges run from lock acquisition for a high-finesse system when using digital feedback of fundamentally limited bandwidth to maintenance of and modifications to lock stability when two or more coupled optical springs act within the system.

We begin with a brief discussion of control theory and its application to digital systems,



Figure 5.1: Block diagram of a generic control system, where the symbols are as defined in the text.

followed by an introduction to the particular system used in the Glasgow 10 m prototype. The lock acquisition scheme and loop gain of the final lock state for both cavities are then described.

5.1 DIGITAL FEEDBACK AND THE CDS

We begin with a short introduction to the aspects of control theory relevant to this work. This includes derivation of the open- and closed-loop gain of a feedback system and a discussion of stability criteria for that system, followed by an introduction to the Control and Data System (CDS) that was developed and is used by LIGO, and is also used at GEO 600 and at the AEI Prototype. This has advantage over traditional analogue systems in that the frequency-domain shape and gain of the servo may be augmented in real-time from a computer interface, and it is just such a system that we employ to control our coupled-cavity experiment.

5.1.1 A Brief Introduction to Control Theory

Figure 5.1 displays a block diagram of a negative feedback control system. The dynamics of this system are described by the *transfer function* G(s), itself a function of the complex variable $s = \sigma + j\omega$. This is a generic description which can apply to any linear dynamical system subject to negative feedback, in which the intended outcome is for the system output, y(t), to follow some control signal, r(t), as closely as possible, often in practice to neutralise the effect of external disturbances upon the system being controlled. This is achieved by identifying the *error signal* as the difference between the measured output and the desired control signal. A further transfer function, $\beta(s)$, very likely also to be frequency-dependent, is applied to this error signal before it is added into the control signal in such a way as to minimise the difference between the output and control signal. If the feedback signal at the summing point is here given by $b(t) = \beta(t)y(t)$, the error signal is then described by e(t) = r(t) - b(t) [95]. With reference to the block diagram, we can relate the input control signal and output *via*

$$y(s) = G(s)e(s) = G(s)[r(s) - b(s)] = G(s)[r(s) - \beta(s)y(s)] = \frac{G(s)}{1 + \beta(s)G(s)}r(s).$$
(5.1)

This then allows us to define the *closed-loop transfer function* of the system, i.e. the ratio of output to input in the presence of feedback, as

$$CLTF = \frac{y}{r} = \frac{G}{1 + \beta G}, \qquad (5.2)$$

and the open-loop transfer function as the same in the absence of feedback:

$$OLTF = \frac{b}{e} = \beta G.$$
(5.3)

In order to define and construct appropriate transfer functions for any desired negative feedback system, it is convenient to describe the closed-loop transfer function in terms of its complex *poles* and *zeros*. These are identified from the *characteristic equation*, itself constructed by setting the denominator of the CLTF (Eq. 5.2) equal to zero: $1 + \beta(s)G(s) = 0$. The values of s corresponding to roots of the characteristic equation then give the complex poles and zeros of the transfer function. Including an arbitrary gain factor K, representing any scalar gain applied to $\beta(s)G(s)$, we find that [96]

$$1 + \frac{K(s+z_1)(s+z_2)\cdots(s+z_m)}{(s+p_1)(s+p_2)\cdots(s+p_n)}s = 0, \qquad (5.4)$$

where z_m and p_n represent the *m* zeros and *n* poles of the system respectively.

From Eq. 5.4, we see that the closed-loop transfer function depends on the poles and zeros of the open-loop transfer function, βG . Thus if the OLTF can be constructed in such a way that it is *stable*, the closed-loop transfer function ought itself remain stable in the presence of feedback. The stability (or otherwise) of a system can be determined using the Nyquist criterion [97].

The Nyquist criterion determines whether a closed-loop system is stable by analysis of the frequency response and poles of the corresponding open-loop transfer function. It is of particular use in practical applications as it allows the stability of the complete closedloop system to be evaluated using graphical representations of the OLTF; e.g. Bode plots of either modelled or measured open-loop response. Formally, the Nyquist stability criterion assesses the relative stability of any system by careful analysis of its response in the complex s-plane. This criterion may be very briefly summarised as follows: plotting the open-loop transfer function $\beta(s)G(s)$ in the complex s-plane, we find in general that the system becomes more oscillatory the closer the approach of the OLTF to the $(-1 + j \times 0)$ point [96]. This allows us to very simply express the relative stability of any system in terms of two easily identified values: the phase margin and the gain margin.

The phase margin is evaluated at the frequency at which the magnitude of the OLTF, $|\beta(s)G(s)|$, becomes one - the *unity gain point* (UGP) - and is generally expressed as the phase lag required to bring the system to critical stability *via* an intersection with the $(-1 + j \times 0)$ point in the complex plane. If the phase of the OLTF at the UGP is ϕ , the phase margin γ is 180° plus this value:

$$\gamma = 180^\circ + \phi \,. \tag{5.5}$$

In a reciprocal fashion, the gain margin is evaluated at the frequency at which the phase of the OLTF is -180° (and hence the phase margin $\gamma = 0$), and is the gain required to bring the system to critical stability – i.e. the inverse of the OLTF gain at this point. Thus, if we define the -180° phase point to be at frequency $\omega_{-180^{\circ}}$, the gain margin K_g may be expressed as

$$K_g = \frac{1}{|G(j\omega_{-180^\circ})|} \,. \tag{5.6}$$

These criteria give enough information to identify a stable open-loop (and hence closedloop) transfer function. However, further constraints linking the phase and gain margins may be applied if we wish to minimise the potential for oscillation at the UGP frequency. A general guide supported by Ogata [96], amongst others, is to maintain the phase margin between approximately 30° and 60° and gain margin above 6 dB. These constraints indicate that in a Bode diagram of the OLTF the logarithm of the magnitude should have a slope with frequency shallower than -40 dB/decade through the unity gain point. Typically, this implies a convenient slope of -20 dB/decade, or $1/\omega$, through the UGP for adequate stability.

To measure the open- and closed-loop transfer functions of any system, only two measurements are required: firstly, a frequency-dependent measurement (a swept-sine response, for example) between two points in the loop with the system feeding back, or "locked", in normal operation – this is often termed the M_1 measurement; plus a second measurement performed between the same two points but with the feedback path open, or "unlocked" – this is the M_2 measurement. With these two measurements taken over the same frequency range of the system, the open- and closed-loop transfer functions throughout that range are obtained via

$$OLTF = \frac{M_2}{M_1} - 1$$
 (5.7)

and CLTF =
$$1 - \frac{M_1}{M_2}$$
. (5.8)

These measurements are useful for characterisation of the frequency-shaping of the control loops, allowing visual inspection of the resultant OLTF and CLTF to determine stability margins. As we shall see, the method of M_1 and M_2 measurement inside the control loops will also form our main method of optical spring measurement, as the influence of the higher optical gain about the associated resonance imposes obvious transformations upon both the OLTF and CLTF of each cavity system.

The conditions covered in this section allow us to design complete and stable control loops for acquisition and lock of both cavities of the multiple-spring system. The feedback servo is in practice implemented digitally, using the Control and Data System (CDS).

5.1.2 DIGITAL CONTROL

Digital control allows for very complicated frequency-shaping to be implemented in real-time *via* a computer interface, and in most cases makes stability and loop performance much easier to assess when compared to an equivalent analogue system. However, there remain a number of subtleties to digital control that must be taken into account in the design of feedback loops.

Any digital system will operate using discrete samples of the input signal of length T_s , or equivalently of sample frequency $\omega_s = 2\pi/T_s$. If a continuous signal f(t) is sampled at multiples n of the sample time, the sampled signal will be [95]:

$$f_{samp}(t) = \sum_{n=-\infty}^{\infty} f(t)\delta(t - nT_s), \qquad (5.9)$$

where δ here is the Dirac delta function, which can itself be expressed as a Fourier series via

$$\sum_{n=-\infty}^{\infty} \delta(t - nT_s) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} e^{ik\omega_s t}.$$
(5.10)

Thus we can express the sampled signal as

$$f_s(\omega) = \frac{1}{T_s} \sum_{n=\infty}^{\infty} f(\omega - n\omega_s).$$
(5.11)

This then indicates that the power spectrum of the continuous input signal $f(\omega)$ is duplicated at each multiple n of the sample frequency ω_s in the sampled signal $f_s(\omega)$. To examine the implications of this effect, consider the Nyquist frequency, defined as half the sampling frequency: $\omega_N = \omega_s/2$. If the highest-frequency component in the continuous signal is less than ω_N , the sampled signal accurately recreates the continuous input signal. However, if signals in f exceed ω_N , these higher-frequency components will contaminate the power spectrum of f_s . In the time domain, this means that any continuous signal of frequency greater than $1/2T_s$ will be mistakenly represented as a lower frequency after the digital conversion. This effect is known as *aliasing*. In general, this dictates that sampling must be achieved at a minimum rate of twice per period for a faithful digital representation of the input signal. In order that this effect is avoided, in particular for systems in the Glasgow 10 m prototype where, shown in Section 5.2, control signals passing through the digital filtering system will have components of many megahertz in frequency, an analogue low-pass filter must be added to the control loop prior to sampling by the digital system with a cut-off frequency of $\omega_s/2$ or less. As described in Section 5.1.3, our digital system incorporates analogue "anti-aliasing filters" on its inputs for just this purpose.

When an analogue signal is converted into the digital domain, as will be necessary in this experiment when interfacing between the analogue signals derived from the cavities and the digital feedback systems, some information is lost in the conversion of the continuous signal into discrete time-samples. This is *quantisation noise*, with associated noise arising from analogue-to-digital conversion (ADC) expressed as [98]

$$\epsilon_{\rm ADC} = \sqrt{\frac{V_{\rm adc}^2}{12f_N}} \,\mathrm{V}/\sqrt{\mathrm{Hz}}$$
(5.12)

in amplitude spectral density, where V_{adc} is the minimum voltage representable by the ADC process and f_N is again the Nyquist frequency. Alternatively, we may express the quantisation noise in terms of the full input voltage range of the system, V_{range} , the sample



Figure 5.2: Example of a filter bank in the Control and Data System. Each filter bank added to the Simulink servo model allows up to eight separate filter stages to be switched in and out of the signal path. Test signals can be injected into the loop via the "EXC" point identified on the filter bank, and read out via the other labelled test points (e.g. "IN1", "OUT").

time and the number of bits associated with the ADC, b:

$$\epsilon_{\rm ADC} = \frac{V_{\rm range}}{2^b} \frac{T_s}{6} \, \mathrm{V}/\sqrt{\mathrm{Hz}} \,. \tag{5.13}$$

5.1.3 The CDS

The Control and Data System (CDS) [99], developed predominantly for use at the LIGO sites but also adopted by a number of prototype laboratories, is composed of a combined hardware and software system for real-time digital control and data acquisition, storage of data and automation of experimentation.

Digital servo design is achieved through Simulink[®] and MATLAB[®], and custom software is used to interface with live experiments in real-time. This software allows arbitrary filters to be applied to an input signal, which are represented graphically in "filter banks" like that shown in Figure 5.2. Here filters can be switched in and out of the signal path, and "test points" are identified at which signals can be either injected or read out, allowing measurements to be performed in the active control loops as described in Section 5.1.1. Figure 5.3 shows an example Simulink model detailing graphically the signal paths from inputs to outputs. Each filter bank is represented by a green block in this model.

For the multiple optical spring experiment, the designs of the CDS input and output channels have important implications for the control loops. There are thirty-two input channels incorporating anti-aliasing filtering, and sixteen output channels with corresponding antiimaging filters, to remove spurious frequency components above the Nyquist frequency. The CDS operates at a maximum sample rate of 65536 Hz, such that any unfiltered signals of frequencies higher than 32.5 kHz will result in aliasing. Accordingly, the anti-aliasing



Figure 5.3: Section of the Simulink model used to route signals from CDS inputs (left-hand side) to outputs (right-hand side) and apply mathematical and logical operations for the multiple spring experiment. Note that each green block in the diagram represents a filter bank like the one shown in Figure 5.2.

and -imaging filters consist of third-order Butterworth low-pass filters with corner frequencies of 9 kHz as a compromise between attenuation at 32.5 kHz, phase lag below 9 kHz and circuit complexity. Figure 5.9 displays a measured transfer function through both the in-built anti-aliasing and anti-imaging filters.

The design of these input and output filters places a limit on the maximum speed of feedback. In Section 5.3 we explain how this affects our feedback systems design and describe an alternative filter design to facilitate slightly faster feedback.

The ADC noise of the CDS has been measured to a level of $1.8 \times 10^{-6} \text{ V}/\sqrt{\text{Hz}}$. From Eq. 5.13, this agrees with the CDS ADC effectively operating at 13 bits.



Figure 5.4: Sketch of the Pound-Drever-Hall locking scheme for a simple two-mirror cavity. The electro-optic modulator (EOM) imposes sidebands upon the carrier beam at frequencies set by the local oscillator (LO). The fields returning from the cavity are detected at the photodiode (PD), the RF output of which is fed into a mixer. Here the signal is multiplied by the LO frequency in the mixer to obtain an error signal. This error signal, proportional to the length of the cavity, is then fed-back via a servo to either the laser frequency or optic positions to lock the cavity to the resonant point.

5.2 The Pound-Drever-Hall Technique

The Pound-Drever-Hall technique [51] has become the standard locking scheme for optical cavities in the field of gravitational wave interferometry. Although initially intended as a method for stabilising the frequency of a laser by locking it to an external reference Fabry-Pérot cavity, the method may also be used to achieve the opposite aim: to stabilise the length of an optical cavity by "locking" it to a stable laser frequency.

The technique for locking a single optical cavity is illustrated in Figure 5.4. The carrier beam incident upon the cavity has phase sidebands imposed upon it by an electro-optic modulator (EOM) at frequencies chosen to lie outwith the bandwidth of the cavity. Accordingly, when the carrier is resonant in the cavity the sidebands will be reflected from the input test mass (ITM). The reflected beam is then composed of the sum of the fields reflected from the ITM and those leaking back out through the ITM. When the carrier is on-resonance, then, the photodiode sampling the reflected beam ("PD" in Figure 5.4) observes the beat between the carrier and the sidebands at radio-frequencies (RF). The resulting signal is then *demodulated* to yield the *error signal* which can be fed-back to either the frequency of the laser or to the position of either test-mass to hold the cavity on-resonance.

To derive this error signal more rigorously, we must consider the incident, intra-cavity and reflected electric fields. As outlined in Section 3.2.1, the incident and circulating



Figure 5.5: Schematic diagram of a simple two-mirror cavity. Identified are the incident and reflected fields, the circulating intra-cavity fields, and the transmitted field. Label definitions are given in the text.

fields for a simple two-mirror cavity, as shown in Figure 5.5, can be described by [49]

$$E_I = E_0 e^{i\omega t} \tag{5.14}$$

$$E_{C_1} = i\tau_1 E_I + \rho_1 E_{C_2} \tag{5.15}$$

$$E_{C_2} = \rho_2 E_{C_1} e^{-2i\phi} , \qquad (5.16)$$

where $\omega/2\pi$ is the frequency of the carrier light and E_0 its initial electric field amplitude, ρ and τ are the amplitude reflection and transmission coefficients of the mirrors, and $\phi = \frac{2\pi L}{\lambda} = \frac{\omega L}{c}$ is a phase term due to the propagation of the field along the length of the cavity, *L*. Now, identifying that the reflected field is composed of E_I reflecting from the ITM and E_{C_2} transmitted through the ITM, and noting that by rearranging Eq. 5.15 and 5.16 we are able to express E_{C_2} in purely in terms of E_I , we find that

$$E_R = \rho_1 E_I + i\tau_1 E_{C_2} = \rho_1 E_I - \frac{\tau_1^2 \rho_2 E_I e^{-2i\phi}}{1 - \rho_1 \rho_2 e^{-2i\phi}}$$
$$= \rho_1 E_I - \frac{\tau_1^2 \rho_2 E_I (e^{-2i\phi} - \rho_1 \rho_2)}{1 + \rho_1^2 \rho_2^2 - 2\rho_1 \rho_2 \cos 2\phi}.$$
(5.17)

Similarly, we can express the transmitted field, E_T purely in terms of the input field and the mirror properties:

$$E_T = -\frac{\tau_1 \tau_2 e^{-i\phi}}{1 - \rho_1 \rho_2 e^{-2i\phi}} E_I.$$
 (5.18)

If phase sidebands are then imposed upon the beam at frequencies $\pm \Omega$ relative to the carrier, i.e. at frequencies of $(\omega \pm \Omega)/2\pi$, with modulation depth β , the incident electric
field becomes [100]

$$E_{I} = E_{0}e^{i(\omega t + \beta \sin \Omega t)}$$

$$\simeq E_{0}[J_{0}(\beta) + 2iJ_{1}(\beta)\sin \Omega t]e^{i\omega t}$$

$$= E_{0}[J_{0}(\beta) + iJ_{1}(\beta)e^{i\Omega t} + iJ_{1}(\beta)e^{-i\Omega t}]e^{i\omega t}.$$
(5.19)

The field returning from the cavity is, again, composed of the fields reflecting from the ITM and those leaking back through the ITM from inside the cavity. Since we now have phase sidebands at frequencies $\pm \Omega$ applied to the input field, all three frequencies must now be considered when evaluating the field returning from the cavity. With reference to the simpler case expressed in Eq. 5.17,

$$E_{R} = J_{0}(\beta)E_{0}e^{i\omega t}\left(\rho_{1} - \frac{\tau_{1}^{2}\rho_{2}E_{I}e^{-2i\phi}}{D}\right) + iJ_{1}(\beta)E_{0}e^{i\omega t}e^{i\omega_{+}t}\left(\rho_{1} - \frac{\tau_{1}^{2}\rho_{2}E_{I}e^{-2i\phi_{+}}}{D_{+}}\right) + iJ_{1}(\beta)E_{0}e^{i\omega t}e^{-i\omega_{-}t}\left(\rho_{1} - \frac{\tau_{1}^{2}\rho_{2}E_{I}e^{-2i\phi_{-}}}{D_{-}}\right).$$
(5.20)

Here we have explicitly separated the terms for the upper and lower sideband frequencies, at $\omega_+/2\pi$ and $\omega_-/2\pi$ respectively, allowing us to define $\phi_{\pm} = \omega_{\pm}L/c$. For convenience, the definition $D = 1 + \rho_1^2 \rho_2^2 - 2\rho_1 \rho_2 \cos 2\phi$ has been made, with corresponding terms for the upper and lower sidebands at D_+ and D_- .

This return field is then detected on the photodiode, which produces a signal proportional to the power in the optical field, $P_R = |E_R|^2$. With reference to Eq. 5.20, and avoiding some of the more extraneous algebra, this will give some terms at zero frequency (i.e. DC signals), a term derived from the resonant carrier leakage field beating with the reflected sidebands, a term derived from the resonant sidebands beating with the reflected carrier, and a great many higher-order terms (which, as we show below, can be ignored at this point).

This electronic signal from the photodiode, directly proportional to the optical power, is then passed through a *mixer* which *demodulates* it. The mixer takes the input RF signal, in this case our optical signal which has prominent components at the beat frequency between the carrier light and the sidebands (i.e. at frequency Ω), and multiplies it by the local oscillator which was initially used to define the modulation frequency *via* the electro-optic modulator in the input beam path (at frequency Ω , once again). Since this mixing can be thought of very simply as a multiplication of two sine waves



Figure 5.6: Error signal simulation performed in Optickle. For this simulation, input and end test masses are separated by 10 m, assumed lossless and of transmissions 500 ppm and 100 ppm respectively.

of frequency Ω , and since for two sine waves at frequencies Ω_1 and Ω_2 we know that $\sin \Omega_1 t \times \sin \Omega_2 t = \frac{1}{2} [\cos (\Omega_1 - \Omega_2)t - \cos (\Omega_1 + \Omega_2)t]$, the mixer output will consist of a zero frequency signal proportional to the length of the cavity, as well as signals at 2Ω which can be filtered out at this stage. The remaining DC signal is the *error signal*, a simulation of which can be found in Figure 5.6. Note that the size of this signal is subject to the relative phase of the RF signal from the photodiode and the local oscillator signal by which it is multiplied. In practice, the local oscillator phase is ordinarily optimised *in situ* to maximise the error signal.

The simplified DC error signal may be expressed mathematically as [49]

$$P_{\rm PDH} = \frac{P_{\rm demod.}}{P_0} = J_0 J_1 \rho_2 \tau_1^2 \left[\frac{\sin 2\phi_+}{D_+} + \frac{\sin 2\phi_-}{D_-} - 2 \frac{\sin 2\phi}{D} \right] + \frac{J_0 J_1 \rho_2^2 \tau_1^4}{D} \left[\frac{(\cos 2\phi - \rho_1 \rho_2) \sin 2\phi_+ - \sin 2\phi (\cos 2\phi_+ - \rho_1 \rho_2)}{D_+} + \frac{(\cos 2\phi - \rho_1 \rho_2) \sin 2\phi_+ - \sin 2\phi (\cos 2\phi_+ - \rho_1 \rho_2)}{D_+} \right].$$
(5.21)

This error signal then forms the basis of the feedback signal, which may be applied to either the laser frequency (by, for example, regulating the temperature of the laser crystal) or to the cavity length (by actuating on one or both suspended optics), with appropriate filtering such that when the feedback loop is closed the servo maintains stability. This feedback signal acts to drive the reflected signal from the cavity to the zero-crossing point of the error signal, corresponding to the maximum point on the DC power curve (also shown in Figure 5.6).

In our system, we use this same method to derive control signals for our cavities. As we explain in the following sections, the decision to use a digital system to control filtering of the error signal within the control loop offers huge flexibility both in real-time alterations to the servo design and in the number and variety of measurements that can be performed within that system. However, using such a servo for the high-power, radiationpressure-dominated optical system that we have built comes with its own challenges, which must be overcome for reliable, robust locking and detunings, as well as measurements of opto-mechanical couplings on the scales that are expected.

5.3 Control Systems Design and Implementation

For the purposes of this experiment, in which optical spring resonant features are measured from approximately 100 Hz up to around 1 kHz, we pre-stabilise the source laser by locking it, with feedback to the laser frequency, to a separate 10 m suspended-optics reference cavity. The two MOSES cavities are then locked to the laser frequency *via* feedback to the positions of the input test masses to sub-picometre precision. This feedback is controlled *via* servos implemented in the CDS, which allows very flexible filtering and signal injection capabilities when compared to the analogue systems used in the past. Since the measurement region extends up to a kilohertz in frequency, for appropriate stability of the control loops the servos are designed to give unity gain points of at least a kilohertz.

With this setup, monitoring of the reflected and transmitted DC signals from each cavity reveals fundamental cavity mode fringes of width $\sim 100 \,\mu$ s as each cavity swings freely through resonance. The CDS imposes an inherent phase delay upon any propagating control signals that makes acquisition of lock during these short fringe time-scales, as well as stable lock with unity gain frequencies above 1 kHz, very challenging indeed. This is described in detail in the following sections.

5.3.1 FREQUENCY STABILISATION

The system is seismically-isolated to a displacement sensitivity estimated to be approximately $10^{-18} \text{ m}/\sqrt{\text{Hz}}$ at 100 Hz. However, electronic noise in the control loops and thermal noise throughout the system, amongst other likely sources, limit the sensitivity of



Figure 5.7: Schematic of the experimental system, updated to include locking of the laser frequency to a separate 10 m reference cavity using the Pound-Drever-Hall method for stabilisation purposes.

each cavity to less than this. Since this experiment relies upon observation of the optical gain changes in each cavity in the presence of opto-mechanical rigidity, it must be ensured that our limiting noise source arises from the propagating laser fields.

From previous experiments performed in the Glasgow 10 m Prototype, the freely-running frequency noise on the source laser is known to behave as approximately $(10^4/f) \text{ Hz}/\sqrt{\text{Hz}}$, where f is the observation frequency [53]. This places frequency noise at 1 kHz at approximately $10 \text{ Hz}/\sqrt{\text{Hz}}$. Equivalently, this yields a displacement noise of $3.6 \times 10^{-13} \text{ m}/\sqrt{\text{Hz}}$. In order to reduce frequency motion of the laser enough to facilitate lock of both cavities used for this experiment it is stabilised using an additional 10 m reference cavity. This is achieved once again using the PDH technique, this time feeding back to the laser frequency *via* two channels: the laser PZT, acting between 10 Hz and 10 kHz, and the NPRO crystal temperature below 10 Hz. This is illustrated in the updated experimental schematic, Figure 5.7. The layout of the three cavities inside the vacuum tanks is shown in Figure 4.1.

As was confirmed several times during various laboratory experiments, the PZT actuator responds as 1.35 MHz/V and the peltier cooler temperature control as roughly 1 GHz/V. The laser itself maintains a flat frequency response over the entire bandwidth of the frequency stabilisation servo. Figure 5.8 shows the measured and simulated OLTF

of the frequency-stabilisation servo. For the purposes of this experiment, as has been noted, the pertinent frequency range is between 100 Hz and 1 kHz, through which the frequency-stabilisation servo is well-characterised by J. Macarthur [53]. This implies a reduction in frequency noise at 1 kHz of approximately a factor 30, or yielding approximately $10^{-14} \text{ m}/\sqrt{\text{Hz}}$ displacement sensitivity for our two cavities. This is confirmed in Chapter 6, in which we measure displacement sensitivities of approximately $0.9 \times 10^{-14} \text{ m}/\sqrt{\text{Hz}}$ at 1 kHz. Small discrepancies are due to manual adjustment of the servo gain between lock states to improve stability of the frequency stabilisation loop.



Figure 5.8: Modelled and measured magnitude response of the frequency stabilisation servo. The PZT response is dominant through the frequency range shown, with the temperature actuation taking over below 10 Hz. Models of the PZT and temperature actuation responses were simulated in LISO [101]. Simulations and measurements performed by J. Macarthur [53]. No data was taken below 100 Hz to avoid excitation of suspension modes.

5.3.2 CDS TIMING DELAY AND SERVO BANDWIDTH

The CDS operates at a sampling rate of 65 kHz, giving a discrete sample time of $15.26 \,\mu\text{s}$ (henceforth referred to as a *clock cycle*). Upon propagating a signal through the CDS from input to output, that signal accrues a time delay of a few clock cycles, depending on the input and output configuration used. This delay manifests as a phase lag in any propagated signal. For our purposes, this delay poses two problems: firstly, it

Configuration	Timing Delay (clock cycles)	Timing Delay (μ s)
No AA or AI Filters	3.64	55.55
Custom AA and AI Filters	5.12	78.13
In-built AA and AI filters	8.86	135.20

Table 5.1: Timing delay of various input/output configurations of CDS. The timing delay is defined here from the frequency at which the phase lag passes through 180°.

imposes a fundamental limit to the speed at which feedback can be applied. In order to lock the cavity, feedback is triggered on when a resonant fringe is detected in the reflected field (like that shown in Figure 5.10), the intention being to "catch" the cavity and drive it towards lock in the duration of this fringe, and so feedback speed is critical. Secondly, the phase lag imposes a further limit to the maximum frequency of unity gain in each servo, as at some frequency this lag will dominate and eventually lead towards instability.

This delay is measured by performing transfer functions through CDS, injecting a sweptsine signal at various input and output points and identifying the point at which the phase lag reaches 180°; from this point, we can calculate the corresponding timing delay in units of clock cycles. In its "standard" setup, each CDS input includes a hardware anti-aliasing filter, and each output an anti-imaging filter. Unfortunately, the design of these filters imposes a prohibitively large phase lag at our intended unity gain region, reaching 45° lag at approximately 1.5 kHz. An external analogue "fieldbox" has been designed, with anti-aliasing and -imaging filters of Scultete design [102] to push this lag a little higher in frequency without compromising isolation at higher frequencies. This is achieved using higher-Q poles in the filter, with the result that a peak in response is also observed at around 10 kHz; fortunately, there are no significant signal components around 10 kHz and so this response peak doesn't pose a problem for this experiment. Using this fieldbox as an alternative input stage to CDS, and bypassing the hardware filtering on the existing inputs and outputs, application of faster feedback is possible. Figure 5.9 dislays transfer function measurements through CDS using (1) just its in-built input and output filtering, (2) the custom-made fieldbox filters and CDS inputs and outputs with their hardware filters bypassed, and (3) just through the bypassed inputs and outputs alone.

The results of these measurements in terms of the observed timing delay in clock cycles and seconds are presented in Table 5.1.

With reference to these values, and noting that for the reasons laid out in Section 5.1.2 we can't realistically expect to operate without anti-aliasing and -imaging filtering, in order



Figure 5.9: Transfer functions through CDS from (1) in-built anti-aliased CDS inputs to in-built anti-imaging CDS outputs (black); (2) custom anti-aliased inputs to custom anti-imaging outputs, bypassing in-built CDS filters (red); and (3) CDS inputs to outputs, with in-built filters bypassed (blue). Note that the differences in magnitudes are due to gain differences depending on the measurement points.

to push the unity gain point appropriately high above the measurement region we have to operate by replacing the in-built filters with the custom ones in the analogue fieldbox. Even so, we remain limited to a unity gain point some distance below 2 kHz to ensure enough phase margin at this point for adequate stability.

Figure 5.10 illustrates a typical DC fringe in the reflected signal from one of the two cavities, with the corresponding Pound-Drever-Hall error signal. The typically observed duration of the DC fringe was 100 μ s, measured as the FWHM of the fringe. This means we have a time-scale of 100 μ s in which to trigger, switch on feedback, catch the mass and drive towards the centre of the fringe to acquire lock. With fast analogue feedback, this can be relatively simple; however, even using the custom anti-aliasing and -imaging filters we are limited to a timing delay of 5.12 clock cycles, or 78.13 μ s, in our digital feedback. On top of this, the finesse of each cavity is such that the DC error signal experiences sizeable "ringing" as a fringe passes, due to the decaying cavity power beating with the external fields (an explanation of this effect may be found in Section 4.6.3). This ringing causes the error signal to change sign on the μ s scale, such that even if CDS was to begin



Figure 5.10: Measured fringe in the reflected DC signal and corresponding Pound-Drever-Hall error signal, for s-polarised input light. This is a typical example of an observed DC fringe due to random uncontrolled motion of the suspended optics.

feeding back within the duration of the fringe, the feedback wouldn't be fast enough to compensate it.

5.3.3 Lock Acquisition by Finesse Transition

The first potential locking solution approached was to artificially lower the finesse of each cavity by rotating the polarisation of the input field. Since the cavity folding optics have 45° angle of incidence, they will have different reflectivities for s- and p-polarised light. The optic manufacturer did not provide estimates of the expected reflectivity for p-polarised light, but the change in finesse was evident experimentally by noting the change in width of DC fringes and drop in observed throughput power from each cavity. When the polarisation is rotated by 90° from s- to p-polarisation, the fringes widen considerably from ~ 100 μ s to ~ 2 ms. This gives ample time for the digital feedback to kick in and lock the cavity to the centre of the fringe.

With the system locking in this "low-finesse" state, the intention had been to rotate the input polarisation by 90° again back to s-polarisation, adjusting the gain of the feedback servo to keep the loop stable and transition to the full "high-finesse" lock. This would have been possible had the s- and p-polarisations locked to precisely, or very close to, the same cavity length. As it turns out, the differences in reflection depths in the coatings of the folding optics are large enough for the two polarisations to make a smooth transition from one to the other impossible. Expressed in terms of cavity tuning, where one free spectral range is 180°, the two polarisations look to be separated by roughly $20^{\circ} - 30^{\circ}$. This is illustrated on the DC throughput signal in Figure 5.11. Figure 5.11a depicts several instances of the cavity swinging freely through one free spectral range. Each plot is scaled such that one free spectral range with p-polarised light lies between 0° and 180°. Red arrows indicate the position of the decreasing s-polarisation fringes as the input polarisation is rotated by changing the angle of a half-waveplate in the beam path through 45°. Note that the plots have not been calibrated for relative velocity of the cavity optics, which accounts for the varying position of the sub-free-spectral-range features. In spite of this, it is evident that the s- and p-polarised fringes are separated in frequency by a margin significantly larger than the cavity linewidth, making any "jump" from one polarisation to the other after acquiring lock prohibitively difficult¹. This is further illustrated on a smaller scale in Figure 5.11b, wherein the scaling has changed to separate p-polarised and s-polarised fringes by the approximately 20° observed in Figure 5.11a. Note that for each of these traces the mirrors were swinging in the opposite direction to those in Figure 5.11a, hence the opposite relative positioning of s- and p-polarised fringes.

5.3.4 Lock Acquisition and Maintenance by Widening of the Error Signal

As described in the preceding section, acquisition of lock is possible with a wider DC fringe, and hence wider error signal from which feedback signals are derived. If this time-domain widening is achieved *via* manipulation of the cavity finesse *via* the input light field polarisation, however, problems arise when transitioning towards the desired full lock state. In order to lock the two cavities in this experiment, then, acquisition and maintenance of lock could feasibly be achieved using a method that widens the capture range of the servo – effectively increasing the number of clock cycles per triggered signal, and thus allowing more time in which to apply feedback – and maintains this widened state for the duration of the lock. Such a method exists, whereby the DC fringe is unaffected but the error signal itself is artificially widened, using existing signals derived from

¹This technique of "jumping" quickly between cavity tunings is successfully applied in GEO 600 to shift between tuned and detuned signal recycling configurations [103]. However, this is achieved using small steps that do not disrupt the rest of the interferometer, which is not a possibility in the work discussed here.



Figure 5.11: (a) Throughput DC signal traces, scaled so that one free spectral range with p-polarised light lies between 0° and 180°, as the input polarisation is rotated. (b) Zoomedin and rescaled. The angles specified in both plots are those of a motorised half-wave plate situated in the beam path and used to remotely adjust the polarisation of the input field, taken relative to some arbitrary zero point. Note that for each of the traces in (b) the mirrors were swinging in the opposite direction to those in (a), hence the opposite relative positioning of s- and p-polarised fringes.

each cavity. This widening effect has a number of benefits which will be covered in this section, and brings with it some subtle intricacies which affect both the operation of the entire system and the spring results derived therefrom.

The essence of this technique is as follows: in the time-domain, a "wider" error signal may be constructed if the original error signal is divided through by a signal proportional to the buildup of optical power in the cavity as it passes through resonance². Since these two individual signals will always be coincident in time, this division, if performed appropriately, results in an artificially widened error signal every time the cavity swings through resonance. In this case we use "width" to refer to the time in seconds between the two turning points of the bipolar error signal. Thanks to the shape of the power curve which divides the error signal, the effect of the division is to elongate the time between the two turning points such that they are shifted to a greater distance from the zero-crossing than their original positions, which coincide with the linewidth of the cavity.

Mathematically, we obtain the widened signal simply by dividing through the Pound-Drever-Hall error signal (Eq. 5.21) by the normalised throughput DC signal (Eq. 3.32), such that

$$P_{\rm PDH}^{W} = P_{\rm PDH} \left[\frac{|E_T|^2}{|E_{T,\rm max.}|^2} + \sigma \right]^{-1}, \qquad (5.22)$$

where σ is some small offset to avoid division by very small signals (which otherwise couple digitisation noise into the widened signal) and $E_{T,\text{max.}}$ is the on-resonance (i.e. $L = n\lambda$) throughput power:

$$E_{T,\max} = \frac{\tau_1 \tau_2}{1 - \rho_1 \rho_2} E_I = \tau_2 g_C E_I .$$
 (5.23)

Note that the size of σ directly relates to the width of the new error signal, and also affects the magnitude of the error signal; if this method is implemented in the control loops, loop gain must be adjusted accordingly to avoid saturation at any point. An example of the two initial signals and the resultant wider error signal is shown in Figure 5.12. In this example, and in the experiment itself, the power in the cavity is represented by the throughput DC signal from each cavity. Note in particular that if the DC power signal is normalised correctly prior to the division being performed, the gradient of the linear error signal ought to remain unchanged about the zero-crossing. This corresponds to no change in the *gain* of the control loop that employs this error signal. This may be of particular importance if this method is required to lock a cavity for which noise performance is critical, as it allows the widening function to be "switched off" after acquisition without

 $^{^{2}}$ This technique has been applied previously as a method of automatic gain control in the GEO 600 and Virgo detectors [104].



Figure 5.12: Optickle simulation of normalised intra-cavity power, "normal" error signal and widened error signal. As the offset added to the throughput signal, σ , is reduced towards zero, the distance between the turning points of the modified error signal increases from ± 1 up to several linewidths.

compromising the lock state.

The effect of this technique is to allow a much longer time frame after the feedback is triggered for control signals to be applied to the cavity to drive it towards the zeropoint of the error signal. It has the added benefit of reducing the influence of the ringing signal that is imposed upon the error signal as the cavity passes quickly through resonance (depicted in Figure 5.10). A physical example of this from our system is shown in Figure 5.13. Note that on the widened signal the ringing effect is greatly reduced in impact, and that the predominantly linear slope of the widened signal continues well after this effect has died out.

In practice, this wide control signal is constructed using our digital servo system (CDS), and as is demonstrated in Figure 5.13 succeeds in widening the linear capture range of the feedback loops from the roughly 100 ms of the fringe duration to greater than 1 ms. This provides ample time for feedback to be applied to the input test-masses of each



Figure 5.13: Time-series response of the "normal" and modified Pound-Drever-Hall error signal and reflected DC signal. In this instance, the approximate effective error signal width has been increased from $\sim 100 \,\mu s$ to $\sim 1.3 \,ms$. Note that the ringing effect is still observed on the modified PDH signal, but that the linear slope continues after this has died out.

cavity (the LTM and RTM respectively) to catch them and drive each towards resonance. The feedback servo shaping required to acquire and maintain lock in this fashion will be covered in more detail in Section 5.3.5.

A further advantage to the widening effect of this technique becomes important when performing optical spring measurements. Opto-mechanical rigidity is generated by adding in an offset to the error signal, effectively shifting the zero-crossing point to introduce a dependence of intra-cavity power on cavity length. With a "normal" Pound-Drever-Hall error signal shape, the DC gain of the cavity is given by the gradient of the error signal. As the cavity is detuned towards one linewidth, then, the DC gain of the system can be expected to tend to zero. This renders detuning of a cavity towards one linewidth and beyond very difficult if lock stability is to be maintained. By widening the error signal with the technique presented here, the predominantly linear slope of the new error signal extends some distance beyond one linewidth in both the positive and negative direction, allowing the cavity to be detuned by several linewidths while maintaining stability, in turn allowing a much wider range of optical spring frequencies to be explored than would be available to a normal PDH error signal.

In the course of the measurements presented in Chapter 6, use of the modified PDH technique is observed to introduce dynamic components at frequencies of order the linewidth of the cavity that become important as the detuning is increased. This effect is characterised and its implications considered in Section 6.6.

5.3.5 Servo Design

In order that digital filters may be designed such that the system is stable, the frequency response of each individual element must be understood. Here we briefly describe the response of each element of the loop from the output of CDS through to the input, before explaining the digital filtering added such that the whole is stable over the measurement bandwidth from approximately 100 Hz up to 1 kHz.

Figure 5.14 displays a simple block diagram of each of the main elements in the loop. From the output of CDS (in which the anti-imaging filtering has been bypassed), the signal passes through the fieldbox, within which the alternative anti-imaging filters are applied. This signal is then sent, after some attenuation, to the *coil driver*, a simple voltage driver that acts to inject a current through the control coils on the rear of the input test mass of the cavity in question. This then drives the pendulum and imposes motion on the cavity with the intention of minimising the error signal observed in CDS; this error signal is constructed in the RF elements following the reflected signal photodiode external to the cavity, before being sent through the fieldbox again to provide anti-aliasing prior to injection into CDS via a preamp for application of gain.

The fieldbox applies Scultete filtering, comprising a two-pole filter from 10 kHz, and does so identically in both the anti-aliasing and anti-imaging paths. A transfer function measurement through a pair of these filters is given in Figure 5.9.

The attenuator exists to reduce the signal size with a voltage divider by a factor of $-35.1 \,\mathrm{dB}$ in the left-hand cavity loop and $-23.1 \,\mathrm{dB}$ in the right. This is to compensate the pre-amp gain applied prior to the CDS input, as described below.

The coil driver circuitry is used for the application of global control, acting at the bottom stage of the input test mass pendulum to maintain the length of each cavity. Appendix A.3 describes this circuitry. Note that no filtering is applied at this stage and there is a flat



Figure 5.14: Block diagram of the main shaping elements of the control loops.

gain of 20 dB applied to the voltage. The resulting potential is then applied over each coil of resistance $\sim 4 \Omega$ in series with a 500 Ω resistor to generate the required current. This circuitry is primarily used to route signals appropriately for longitudinal, pitch or yaw feedback to the mass.

As described in Section 4.5.1, each stage of a pendulum provides seismic isolation of $1/(\text{frequency})^2$. Since the control signal is injected directly onto the test-mass at the lower stage *via* electromagnetic actuation, this imposes a resonant two-pole low-pass filter upon the signal from the fundamental longitudinal pendulum frequency at $f_0 = 0.69 \text{ Hz}$. A measurement of this response may be found in Section 4.5.3.

The cavity itself then imposes a two-pole low-pass filter from the cavity linewidth. From our finesse measurements as detailed in Section 4.6.3, we estimate the corner frequency of both cavities to be at approximately $\gamma = 700$ Hz. This two-pole filter acts upon the power in the cavity; since the Pound-Drever-Hall error signal (as defined in Section 5.2) is derived from the beat between the power resonating in the cavity that has leaked back out through the ITM and the power in the sidebands reflected directly therefrom, the cavity filtering manifests as a *single*-pole low-pass filter in the amplitude of the error signal.

After detection of the reflected field on the appropriate photodiodes (the circuitry for which can be found in Appendix A.1) the signal is passed to the RF circuitry. A schematic of this is given in Figure 5.15. This circuitry imposes a Butterworth-design low-pass filter



Figure 5.15: Block diagram of the RF setup. PS denotes a 50/50 power splitter and H a "HELA" 12 dBm RF amplifier. These and the mixers are all standard Mini-Circuits components. The 10 MHz local oscillator is phase-locked to the existing oscillator used to set the control sidebands on the input field and obtain control signals for the frequency-stabilisation loop. The Butterworth filters exist to low-pass-filter the signals to remove the 2f signals from the mixer outputs (see Section 5.2 for details).

from 60 kHz to remove the 2f signal resulting from demodulation in the mixer (derivation of which was given in Section 5.2) and maintain only the DC error signal.

After passing through the Scultete filtering in the fieldbox a second time, the signal passes through a preamplifier which applies a single-pole low-pass filter from 10 kHz. This filter is applied to provide extra isolation from higher-frequency signals – in particular, a high-quality factor resonance at approximately 7.5 kHz, presumed to be an internal mode of the test mass from its shape and high frequency – and with some 20 dB of gain. This gain is applied to ensure that higher frequencies remain shot-noise limited after the analogue-to-digital conversion has been applied with its associated digitisation noise (Eq. 5.13). After this point, the signal is applied to the analogue-to-digital converter at the input to the CDS, and the digital filtering takes over, designed to re-shape the error signal such that the closed loop is stable.

Figure 5.16a shows simulations of the gain shaping internal and external to the CDS, and Figure 5.16b shows the shaping of the entire open-loop gain, with an overlay of an on-resonance open-loop transfer function performed through the measurement region, confirming unity gain at 1.2 kHz. Measurements below 100 Hz are made difficult by the presence of many suspension modes that are liable to excitation by the measurement process.



Figure 5.16: (a) Simulated gain in the analogue (external to CDS) and digital (internal) sections of the loop. (b) Simulation and measurement of the complete open-loop transfer function for one cavity. In both (and annotated on (b)): the resonance at 0.69 Hz is the fundamental longitudinal pendulum mode; the five resonant features in the 20-60 Hz range are resonant gain filters, designed to reduce excitation of prominent suspension modes at these frequencies; the notch at 7.5 kHz exists to mitigate excitation of a (presumed) internal mode of the test mass at this frequency.



Figure 5.17: Strength of the global control actuation, calibrated in metres per volt of feedback signal, for each cavity.

5.3.6 GLOBAL ACTUATION CALIBRATION

The feedback strength of the global control actuation was obtained in terms of metres of cavity length change per volt of feedback signal at the output of the CDS. This was achieved by locking the laser frequency to each cavity using the PDH technique – i.e. the same method used to stabilise the laser frequency to the external reference cavity. This method was used as the loop gain of the frequency-stabilisation servo is very well characterised (see Figure 5.8). A swept-sine signal was injected onto the cavity *via* the global control coils upon each input test mass and the PZT response of the laser read out. Resultant calibration gives Figure 5.17, indicating that the two feedback paths have very similar actuation strengths. This information is useful for estimating the gain required in each loop. It also allows estimation of the power in each cavity by monitoring of the CTM displacement *via* the feedback signal strength as each cavity acquires or loses lock – measurements of this are presented in Chapter 6.

5.4 Auto-Alignment of Cavity Optics

In a Fabry-Pérot cavity, misalignment of the circulating or input beams can lead to several negative effects: firstly, misalignments can cause higher-order modes to couple in to the circulating beam [105]. Since these don't constitute optimal matching of the beam mode into the cavity, this leads to excess loss of power. Secondly, and particularly for systems such as these in which high input power is used with light suspended optics, small offsets of the circulating beam from the centre of mass of each optic will result in radiation-pressure-induced torques on the input and end test masses. If the resulting oscillation is large enough, this can cause the system to lose lock entirely. This effect could be particularly damaging in this system, wherein we have two high-power beams impinging upon either side of a single suspended test-mass (the CTM). This torque instability is sometimes referred to as the *Sidles-Sigg effect* [106], and we cover this briefly in the following section.

Fortunately, these angular misalignment issues can be mitigated by the maintenance of beam alignment to the centre of each optic. This can result in necessary realignment of the optics over time as conditions in the lab change (through temperature drifts, etc.) and cause misalignments over longer time-scales. However, this process can be automated to a certain extent by application of extra feedback to the cavity optics. This supplementary feedback, termed *auto-alignment*, will also be covered in the coming sections.

5.4.1 The Sidles-Sigg Instability

In a suspended Fabry-Pérot cavity consisting of two mirrors, m_1 and m_2 , with radii of curvature R_1 and R_2 , each mirror experiences a natural restoring torque to external angular forces due to their suspension wires. If there is no light circulating within the cavity, each mirror is independently subject to oscillations in yaw and pitch resulting from external excitations (e.g. seismic activity). When light is allowed to circulate within the cavity and the system is held on or close to resonance, such that a sufficiently large light field builds up, the resulting opto-mechanical system imposes two coupled modes upon the cavity. These modes of torsional stiffness are not local to either suspension, but rather are stiffnesses between the two cavity mirrors that couple them together. In this section, we will derive an expression for the torsional eigenfrequencies corresponding to these modes. Note that this mathematical description may be applied to both the yaw and pitch modes.

Figure 5.18 illustrates the two coupled torsional modes. In Figure 5.18a, radiation pressure acts to enhance the existing restoring torque due to the stiffness of the suspensions, pushing the mirrors back to their original position. Thus the mirrors will experience a stronger torque than due to the mechanical torque alone, and the eigenfrequency of this mode will be higher than the natural resonance frequency of the uncoupled pendulum suspension. This is the "stiff" torsional mode.



Figure 5.18: (a) "Stiff" tortional mode in a Fabry-Pérot resonator. In this mode, radiation pressure from the circulating light acts to enhance the existing mechanical restoring force due to the optic suspension. (b) "Soft" torsional mode. Here, radiation pressure works against the mechanical restoring force. If the intra-cavity power exceeds a critical level this mode becomes unstable and the cavity is forced away from resonance [107].

In Figure 5.18b, radiation pressure now works against the existing restoring torque. As a result, the net restoring torque is smaller than that due to the mechanical torque alone, and so when the intra-cavity power increases beyond some critical value the net restoring force will become negative. This obviously will lead to instability as the radiation pressure force overcomes the natural restoring force of the pendulum, causing misalignment and loss of lock. This is the "soft" torsional mode.

We can work out the eigenfrequencies of these two modes from the *torsional stiffness* matrix [106]:

$$\bar{\kappa} = \frac{2PL}{c(1-g_1g_2)}\underline{G}\,,\tag{5.24}$$

where P is the power circulating in the cavity, L is the cavity length,

$$\underline{G} = \begin{bmatrix} -g_2 & 1\\ 1 & -g_1 \end{bmatrix}$$
(5.25)

and the g_i are the stability parameters for each mirror *i*: $g_i = 1 - L/R_i$. The eigenvalues of \underline{G} , λ_G , are given by $\det(\underline{G} - \lambda_G \underline{I}) = 0$, i.e.

$$\lambda_G^2 + (g_1 + g_2)\lambda_G + (g_1g_2 - 1) = 0$$
(5.26)

$$\Rightarrow \lambda_G = \frac{-(g_1 + g_2) \pm \sqrt{4 + (g_1 - g_2)^2}}{2}, \qquad (5.27)$$

and so the eigenvalues of the torsional stiffness matrix, $\bar{\kappa}$, are

$$\lambda_{\kappa} = \frac{PL}{c} \left[\frac{-(g_1 + g_2) \pm \sqrt{4 + (g_1 - g_2)^2}}{1 - g_1 g_2} \right].$$
 (5.28)

The torsional eigenfrequencies are then given by the square root of the stiffness eigenvalues divided by the moment of inertia about the appropriate mass:

$$\omega_{\kappa}^{2} = \frac{\lambda_{\kappa}}{I} = \frac{PL}{Ic} \left[\frac{-(g_{1} + g_{2}) \pm \sqrt{4 + (g_{1} - g_{2})^{2}}}{1 - g_{1}g_{2}} \right].$$
 (5.29)

Finally, if we express these relative to the fundamental torsional frequency of the pendulums, ω_0 , we arrive at the expressions for the stiff and soft torsional mode frequencies ω_{\pm} shown in Figure 5.18 [107]:

$$\omega_{\pm}^{2} = \omega_{0}^{2} \pm \frac{PL}{Ic} \left[\frac{-(g_{1} + g_{2}) \pm \sqrt{4 + (g_{1} - g_{2})^{2}}}{1 - g_{1}g_{2}} \right].$$
 (5.30)

As the power in the cavity is increased, we expect the two relative torsional mode frequencies, ω_+ and ω_- , to behave differently: as P increases, so does the eigenfrequency of the stiff torsional mode, ω_+ , which will always be higher than the fundamental pendulum torsional frequency ω_0 ; however, at the same time the soft torsional mode eigenfrequency will *decrease*, until at some critical power $P_{\text{crit.}}$ the fundamental torsional mode frequency ω_0^2 is overwhelmed and ω_-^2 becomes negative, and hence the mode becomes unstable.

The MOSES experiment is susceptible to these instabilities, as indeed is any suspended optical cavity operating with high power. In this coupled-cavity system, it is further



Figure 5.19: Observation of radiation-pressure-driven instabilities in the reflected DC signals from each cavity, presented as the time series of the DC voltage obtained from each photodetector. Top: the right-hand cavity acquires lock, with the beam not centred on the CTM with the result that an oscillation is driven up on the left-hand cavity signal (blue). Note that this signal is also present on the right-hand cavity signal (red), but is less visible due to lower optical power at the detection photodiode. Bottom: some time later, the alignment of the right-hand cavity is minorly adjusted to better centre the beam on the CTM, resulting in a reduction of observed oscillation on the left-hand cavity signal.

exacerbated by the high power beams impinging upon either side of the shared Central Test Mass, such that anti-symmetric beam offsets on opposite sides result in a larger torque. This instability can be improved by careful alignment of the system, as is shown in Figure 5.19 in which the alignment of one cavity is adjusted to remove a slow oscillation from the observed reflected DC signal. In order that slow drifts in alignment over time (e.g. from thermal shifts in the laboratory) don't ring up these instabilities, it may be beneficial to include extra active feedback to automatically align each cavity to achieve optimal power coupling and beam centring.

5.4.2 Auto-Alignment

To ensure that the ability to maintain cavity lock is not compromised by the Sidles-Sigg instability, so-called "auto-alignment" systems can be included in the reflected signal path of the experimental setup. Misalignment sensing is achieved using quadrant photodiode detectors in the path of the reflected signal from the cavities. These photodiodes are able to differentially measure signals across four individual quadrants, allowing both localisation of the beam position on the photodiode and sensing of higher-order resonant cavity modes coupling into the return beam. If the soft torsional mode becomes unstable, the torque applied to the mass will cause a shift in alignment, and hence coupling of the cavity field from the fundamental TEM_{00} mode into either the TEM_{01} or TEM_{10} modes (to first order), depending on the axis of misalignment; feedback systems will then act to return the mass to its equilibrium position, by correcting the mode content detected on the quadrant photodetector.

We aim to integrate this auto-alignment method into our control loops, such that future high-power work in the same system will be less sensitive to radiation pressure torques and angular misalignments. Note that on the time-scale of the experiments documented in this thesis the auto-alignment system was not ready for application during the optical spring measurements presented in Chapter 6; however, it will become a useful addition to the control systems for future experimentation in the same system, so the basic principle of operation and installation are briefly documented here.

When a Fabry-Pérot cavity like that presented in Figure 3.3 is misaligned, the effect on the return beam can be separated into the two effects stated above: misalignment of the return beam, due to angular shifts of the ITM, and coupling of the cavity field into higher-order modes, due to the relative misalignment of the ITM and ETM. Note that the former characterises the angular alignment of the ITM with respect to the input beam, while the latter manifests as a lateral misalignment parallel to the main fundamentalmode beam and characterises the relative alignment of the cavity optics.

The n^{th} -order mode of a beam propagating in a Fabry-Pérot cavity may be expressed as [90]

$$U_n(x,z) = N_n H_n\left(\sqrt{2}\frac{x}{w}\right) \frac{w_0}{w} \exp\left(-i(kz-\phi_n) - x^2\left(\frac{1}{w^2} + \frac{ik}{2R}\right)\right)$$
(5.31)

for a beam propagating along the z axis, where x is the radial distance from that axis, w is the radius of the beam and w_0 the beam radius at the waist position, H_n is the n^{th} order Hermite polynomial, $k = 2\pi/\lambda$, N_n is a normalisation factor, R is the radius of curvature of the coupling optic and ϕ_n an additional phase term. It can be shown that at position z = 0 the fundamental and first higher-order mode can be expressed as [105]

$$U_0(x) = \exp\left(-\frac{x^2}{w_0^2}\right) \tag{5.32}$$

and
$$U_1(x) = \sqrt{2\pi} \frac{x}{w_0} \exp\left(-\frac{x^2}{w_0^2}\right)$$
. (5.33)

Further, if we express the electric field of amplitude A at z = 0 as $E = AU_0$, we can approximate the resultant field E_1 due to a misalignment of angle α in the x-direction in terms of just U_0 and U_1 :

$$E_1 \simeq A \left[U_0 + i \frac{k\omega_0}{\sqrt{2\pi}} \alpha U_1 \right] \,. \tag{5.34}$$

Therefore we see that a misalignment in either rotation or tilt of the cavity will result in coupling from the fundamental mode into the first-order mode, with a 90° phase change. At the beam waist, the misalignment may be expressed now in terms of the phase difference between perfect and measured alignments, termed the *displacement angle* [108]:

$$\Theta_0 = \arctan\left(\frac{\alpha z_R}{\Delta x}\right) \,, \tag{5.35}$$

where z_R is the Rayleigh length and Δx the lateral displacement of the misaligned beam. As the beam propagates away from the waist position, an additional phase shift is observed, corresponding to a lag in phase between the fundamental TEM₀₀ and higher-order TEM_{nm} modes; this is the *Guoy phase shift* [109, 110],

$$\eta(z) = (n+m+1)\arctan\left(\frac{z-z_0}{z_R}\right), \qquad (5.36)$$

and so we can express the displacement angle for first-order modes, i.e. TEM_{01} and TEM_{10} , as

$$\Theta_d = \Theta_0 + \eta \simeq \frac{\alpha z_R}{\phi x} + 2 \arctan\left(\frac{z - z_0}{z_R}\right) \,. \tag{5.37}$$

Evidently, measurement of Θ_d will result in a combination of lateral and angular displacements, corresponding to cavity misalignment and input coupler misalignment respectively. In order to separate these, two measurements are required: one sampling of the beam in the near field, such that $z \ll z_R$ and the measurement is most sensitive to angular misalignment; and another sampling in the far field such that $z \gg z_R$ and the measurement is most sensitive to lateral misalignments. This is achieved using two quadrant photodidoes, arranged with a suitable choice of lenses between to enhance the Guoy phase shift. Each detector then measures a linear combination of lateral and angular misalignments. For appropriate photodiode positioning and lens choices, the output measurements are orthogonal and only simple linear combinations of the two signals are required to extract angular and lateral correction signals, which are then fed-back to the appropriate cavity optic.

In reality, the error signal used to form the feedback is constructed just like the Pound-Drever-Hall signal covered in Section 5.2 Assuming a perfect TEM_{00} input field, the immediately reflected field from the input coupler, containing some fraction of the input beam and the control sidebands, maintains the perfect TEM_{00} mode shape. The portion of the reflected field composed of light leaking out of the misaligned cavity through the input coupler will not have the same profile, such that demodulation of the signal hitting each quadrant of the photodiode at the modulation frequency gives information about the distribution of the wavefront; combinations of demodulated signals then yield information about pitch and yaw misalignment of each optic.

At the time of writing, auto-alignment components were being installed in the laboratory for use in the next stage of experimentation. The author was involved in the design and construction of quadrant photodiodes for wavefront sensing and error signal derivation, and in the simulation of appropriate lens solutions to separate each detector into the near- or far-field for optimal orthogonalisation of signals pertaining to each cavity optic. It is expected that the complete and functioning auto-alignment system will be described in a future publication.

VI Multiple Optical Springs in Coupled Cavities: Results

In this chapter the methods of measurement and notable optical rigidity results are detailed. These are presented broadly in the order in which they were taken, which also conveniently arranges each experimental concept in order of complexity.

We begin with an optical spring in one cavity, and the other held at the resonance point. In this configuration, the opto-mechanical resonant behaviour of the detuned cavity is monitored as it is detuned from resonance to several times its linewidth. This allows us to assess the overall quality of the system as regards precision of measurement, investigate the relative stability of the control loops in the presence both of the opto-mechanical rigidity and the loop gain, and confirm the reliance of the optical spring on intra-cavity power and detuning from resonance. Since the adjacent cavity is being held at the resonance point with feedback, it can be used as a probe of CTM motion. This then forms an analogue of the "local readout" system introduced briefly in Chapter 3, in which the on-resonance cavity may used as a quantum-non-demolition meter to identify the presence of opto-mechanical coupling in the adjacent cavity without disturbing the quantum state of that system. This technique is demonstrated to sensitivity between two and four orders of magnitude greater than previous demonstrations in the same laboratory over the same bandwidth.

In a second set of measurements, optical rigidity is introduced into both cavities of the system, coupled *via* the shared CTM, and observations of opto-mechanical effects are performed inside the control loops of one or other cavity. With springs in both cavities, the opto-mechanical response is observed to be modified from the single-spring measurements described above. This manifests as modified stiffness and damping of the observed opto-mechanical resonance, consistent with Eq. 3.76.

In a third set of measurements, transfer functions are performed simultaneously in the control loops of both cavities in the doubly-resonant system. As well as the modification to the observed stiffness and damping of each spring noted in the previous case, more complex responses are observed resulting from various couplings of the optical springs and measurement signals to the loop gains.

Finally, the sum response of the full two-cavity system is presented, indicating that the "overall" response of the system can be modified within certain frequency bands by careful tuning of each opto-mechanical response.

6.1 Measurement Technique

As outlined in Chapter 4, the measurement sensitivity is limited in the first instance by the residual frequency noise of the laser source. In the frequency range of interest, laser frequency noise is dominant over other noise contributions such as thermal noise on the optic coatings and seismic noise of the surroundings, each of which is of higher magnitude than the quantum noise on the light that is expected to be improved or modified by the presence of optical rigidity. For this reason, the effect of opto-mechanical rigidity cannot in this system be evaluated by monitoring the total noise budget; instead, our measurement plan identifies and characterises opto-mechanical effects *via* their influence upon the control loops for each cavity. These measurements can easily be performed using the digital control system that forms the bulk of the feedback loop to each cavity, and are undertaken in the same way as those yielding the open-loop and closed-loop transfer functions of any feedback system, as discussed in Section 5.1.1.

To recap, two measurements are required in order to evaluate the OLTF and CLTF of a negative feedback system: these are the " M_1 " and " M_2 " measurements described in Section 5.1.1, whereby a transfer function measurement is performed between two adjacent points in the feedback loop, once with the system locked (i.e. stably feeding back) and once with the system unlocked. Equations 5.7 and 5.8 then show how the open- and closed-loop transfer functions may be derived from these two measurements alone.

In the CDS, the injection and readout points for such measurements are the "EXC" and "IN2" points of each Simulink filter bank, as shown in Figure 5.2, and the M_1 measurement is taken as the ratio of the signal sizes at each frequency measured at each of these points. For the most part, the injected signal was chosen to be band-limited white noise, shaped



Figure 6.1: Example of frequency shaping applied to injected noise signal in CDS (top) yielding an in-loop measurement (i.e. ratio of output signal to this excitation signal) of one cavity (bottom). The shaping of this signal has been chosen to yield a good balance of signal-to-noise ratio across the frequency range of interest. In the upper plot, the magnitude of the excitation has been increased with a resonant gain filter in the region about the spring resonance where the gain is lowered. A further resonant gain filter has been included just above 100 Hz to improve the signal-to-noise ratio at the low end of the measurement band. The spring frequency is evident on the lower plot at 200 Hz, with the peak at ~ 1.2 kHz corresponding to a resonance at the unity gain point of the servo.

to provide larger signal size in low-gain regions of the measurement (e.g. lower servo gain at the optical spring resonant frequency, at which point the opto-mechanical response of the cavity has increased). This injection method allows for simultaneous measurement of each specified frequency bin in the region of interest, with arbitrary measurement bandwidth and degree of averaging to achieve the desired precision. An example of the injected noise spectrum and corresponding in-loop spring measurement (i.e. an M_1 measurement) for an optical spring in the 200 Hz region is shown in Figure 6.1.

Since the injection and measurement points in the CDS filter bank are adjacent, the corresponding M_2 measurement can be assumed to be unity. This greatly simplifies the measurement process, and allows very quick calculation of the open-loop transfer function (OLTF) in MATLAB[®] according to Eq. 5.7.

Measurements of opto-mechanical resonances in this chapter are presented in their unal-



Figure 6.2: Coupled spring simulation including all the elements of the control loop (see Section 5.3.5 for a description of each element). This is the format that most of the spring measurements in Chapter 6 will take. This simulated transfer function is performed inside the right-hand cavity control loops, with detunings of $\delta_L = 0.3 \gamma_L$ and $\delta_R = 1 \gamma_R$. Note the pair of coupled opto-mechanical resonances at approximately 150 Hz and 285 Hz.

tered open-loop transfer function format. Fitted models must thus include all the same frequency-shaping that is applied to the real signals in each control loop, either using previous measurements (e.g. the anti-aliasing and anti-imaging filters at CDS inputs and outputs, transfer functions of which are shown in Figure 5.9) or simulations (e.g. the digital filtering in the CDS, as shown in the open-loop gain plot in Figure 5.16a) of each contributory element.

Since the measurements are being presented in the OLTF format, it is useful to include here a simulation of a potential measurement in this format for comparison. This can be considered an extension to the modelling performed in Chapter 4, but with the now well-characterised elements of the control loops included. A simulation of this is presented in Figure 6.2, showing that in the specific case where a spring is present in both cavities and the OLTF is evaluated in one of these cavities, we expect to see two coupled opto-mechanical resonances.

6.2 Precision Optical Spring Measurements

One of the key remits at the design stage of this experiment was to construct an apparatus for the measurement and control of optical springs that is mechanically and optically sufficiently stable that features of the optical spring transfer function are not smeared out due to drifts. This allows for precision determination of the frequency of associated resonances in the system. Previous investigations aiming to study opto-mechanics in coupled-cavity systems have shown this to be very difficult (e.g. [52]), in part due to difficulty in stabilising the intra-cavity power, such that the opto-mechanical resonance is observed to drift in frequency over the duration of a measurement. By avoiding optical coupling between the two cavities and having a direct monitor of the internal power *via* throughput beams, it is possible to track important parameters, such as cavity finesse and detuning, with precision.

As we show here, thanks to the stability of the optic suspensions and source laser, we observe very little drift in opto-mechanical resonance frequencies over the course of measurements lasting up to several minutes, allowing precision observations to be made of the opto-mechanical response throughout the decade-wide frequency band of measurement, from 100 Hz up to 1 kHz..

During the course of each measurement, the CDS records all of the necessary signals for characterisation of the state of the system at the point of measurement – for instance, the throughput DC signals of each cavity (these being the sampled beams exiting each cavity via the 45° folding optics), the error signals of the two locking servos, and the feedback points of both prior to injection of locking signals onto the suspensions. The throughput DC signals provide a proportional measure of the power buildup in the cavity; the same Airy peak function that describes the response of the intra-cavity power to length or frequency tuning (as depicted in Figure 3.4) applies to the DC signal recorded on the throughput photodiodes. Thus by comparing the DC level on the throughput photodiode of a detuned cavity to that observed when the same cavity is held at the resonance point, one arrives at an accurate estimate of the cavity detuning with respect to the linewidth during that time.

Using this method to calibrate the detuning of each measurement, and knowing the properties of the two cavities as laid out in Table 4.1, we are able to very precisely characterise optical spring measurements in either cavity. Figure 6.3a shows a number of opto-mechanical responses as observed in the open-loop transfer function of the left-hand cavity. Calibrating the detuning of each from the recorded throughput DC signals as described above, we find them in order of decreasing resonance frequency to be $\delta_{L1} = 0.60 \gamma_L$, $\delta_{L2} = 0.99 \gamma_L$ and $\delta_{L3} = 1.91 \gamma_L$. These yield opto-mechanical resonances at approximately $f_1 = 271 \text{ Hz}$, $f_2 = 250 \text{ Hz}$ and $f_3 = 158 \text{ Hz}$. Each transfer function was measured over the course of approximately five minutes in order to accumulate a high number of measurements over which the data could be averaged. Notably, the resonant features do not appear to have drifted obviously in frequency or amplitude over this time; any significant drift would have manifested as a wider and lower resonant response than predicted by the simulations. This indicates that the cavity is capable of maintaining a stable optical power level and tuning on this time-scale, allowing observation of the opto-mechanical response to the desired precision.

Figure 6.3b shows similar open-loop transfer functions for the right-hand cavity. Here, the detunings have been calibrated to be $\delta_{R1} = 0.62 \gamma_R$, $\delta_{R2} = 0.97 \gamma_R$ and $\delta_{R3} = 1.90 \gamma_R$. Again, the measured resonant responses are very stable when compared with simulations. The third measurement at approaching twice the linewidth has resulted in a significant drop in the loop gain, lowering the unity gain point from just over 1 kHz down to 400 Hz. This is a direct result of the power in the cavity decreasing as it is detuned further from the resonance point. In the first two measurements of each cavity, at $\delta_{L1,R1}$ and $\delta_{L2,R2}$, this change in UGP has been compensated by applying extra gain in CDS in accordance with the expected change in power.

One important practical concern to be observed from these two plots is that, even with the same nominal detuning with respect to the linewidth of the respective cavity, the resonant frequency differs in each spring measurement. This is consistent with the left-hand cavity containing $(60\pm 6)\%$ the circulating power of the right-hand cavity. The finesse of the two cavities is measured to be the same to better than 10%, so the discrepancy is understood (and shown below) to be due to a difference in the light power entering each cavity. Several effects could contribute to such a difference – e.g. asymmetric beam-splitting into the two input paths, mode-matching errors, slight misalignments of the beam causing clipping on some optics, and imperfections on optic surfaces – all of which can scatter light out of the fundamental mode, and hence reduce the power coupling into the cavity mode. This relative power level is evident throughout the remainder of the results obtained. See the following Section 6.3 for predictions of the on-resonance power in each cavity, with results in agreement with this estimate.



Figure 6.3: Open-loop transfer functions showing opto-mechanical resonances for three different values of detuning in (a) the Left-Hand Cavity (LHCav) and (b) the Right-Hand Cavity (RHCav).



Figure 6.4: Comparisons between calculated K_0 values, based on observed resonant frequencies presented in Figure 6.3 and calibrated detunings, and predictions of theory, for both the left- and right-hand cavity.

Having identified each resonant frequency, we can apply Eqs. 3.59 and 3.74 to calculate the corresponding DC (i.e. zero Fourier frequency) stiffness associated with the optomechanical up-shifting of the fundamental pendulum resonances. These are compared with theoretical models in Figure 6.4, yielding maximum DC opto-mechanical restoring forces of $K_{0,R} = (2.411 \pm 0.096) \times 10^5 \,\mathrm{Nm^{-1}}$ and $K_{0,L} = (1.561 \pm 0.062) \times 10^5 \,\mathrm{Nm^{-1}}$. These were obtained at detunings of approximately 0.6 times the cavity linewidth in both cases, and so, as indicated by the theory, are expected to be close to the maximum possible spring strengths for the input powers at the time, which ought to occur at a detuning of 0.57γ . These measurements, amongst a great many similar ones obtained over the course of data-taking, indicate that optical spring properties can be confirmed to high precision using this system. This is an important result for the characterisation of more complex optical spring couplings in the following sections, as well as for the continued operation of this system with further optical rigidity experiments in the future.

6.3 SINGLE FREQUENCY INJECTIONS AND LOCAL READ-OUT

As stated in Section 3.2.6, the *local readout* method proposes using a secondary beam to monitor the position of one optic in a detuned cavity, such that its motion at low frequencies, at which it is rigidly coupled to the other cavity optic by the optical spring effect, can be sampled without disturbing the quantum state of the entire system. The MOSES coupled-cavity system can be operated as an analogue of this method, if one cavity is locked on resonance and the other detuned therefrom; the error signal of the onresonance cavity then contains information on the motion of the shared CTM. We show this effect by injecting single-frequency signals onto the input test mass of the detuned cavity and identifying these same frequencies in the feedback to the input mass of the adjacent cavity. We confirm the effectiveness of this method to a sensitivity between two and four orders of magnitude greater than previous local readout measurements performed in the same laboratory over the same frequency band. Further, to examine the accuracy with which the optical spring is characterised, this method is employed as a means of extrapolating the buildup of optical power in each cavity from the observed motion of the CTM; the power is then estimated with an alternative and more conventional method, and the two found to agree to within experimental error.

Figure 6.5 displays the spectrum of the injection signal imposed onto the control loop of whichever cavity is detuned, while keeping the other cavity on resonance. The signal is composed of identical-magnitude single-frequency signals at 135 Hz, 390 Hz, 535 Hz, 770 Hz and 970 Hz. These are imposed onto each cavity *via* the ITM, resulting in power fluctuations on the circulating fields at each injection frequency. These fluctuations can be identified in the throughput signal for each cavity. In the presence of opto-mechanical coupling, these power fluctuations result in coherent motion of the CTM and ITM at these frequencies, such that the adjacent cavity control loops must feed-back appropriate signals at the same frequencies in order to maintain the correct cavity length. Thus each injection frequency can be identified in the feedback path to the on-resonance cavity, as in the lower panel of Figure 6.5.

The first point of note is that whichever cavity is employed as the readout, the displacement sensitivity is an approximately flat $10^{-14} \text{ m}/\sqrt{\text{Hz}}$ through the measurement region, with peak sensitivity of $9 \times 10^{-15} \text{ m}/\sqrt{\text{Hz}}$ at 500 Hz. Since the dominant noise through this region is residual frequency noise on the laser, and the cavities are both nominally



Figure 6.5: Demonstration of the "local readout" method. An injection signal (top) consisting of five separate single-frequency signals is applied to the input coupler of one detuned cavity. The adjacent on-resonance cavity is then used as a readout of the resulting CTM motion in the local frame. The feedback signal applied to the adjacent cavity then includes signals at the original injection frequency (bottom), with maximum displacement sensitivities of approximately $9 \times 10^{-15} \text{ m/}/\text{Hz}$ at 500 Hz.

10 m, the displacement noise levels of both should match. A previous measurement of this type performed in the Glasgow 10 m prototype [53] suffered from poorer low-frequency sensitivity due to the use of a low-power supplementary readout system, with approximately linear sensitivity slope of $2 \times 10^{-10} \text{ m/}\sqrt{\text{Hz}}$ at 100 Hz down to $2 \times 10^{-12} \text{ m/}\sqrt{\text{Hz}}$ at 1 kHz. Thus this system reaches sensitivities through the measurement region between two and four orders of magnitude lower than that particular previous demonstration of the "local readout" method – perhaps unsurprisingly, given that the readout system used in this case is a 10 m high-finesse optical cavity.

We now perform an experiment to demonstrate the operation of local readout, normally intended to allow measurement of low-frequency displacement signals, by employing it to measure the steady displacement of the CTM induced by radiation pressure changes resulting from the injected signal. This in turn allows the on-resonance power levels in each cavity to be deduced. These are compared with alternative estimates obtained from on-resonance measurements to confirm the effectiveness of the apparatus for the characterisation of optical springs.

Having calibrated the feedback strength of each cavity in metres per volt (see Figure 5.17), the on-resonance power in each detuned cavity can be ascertained from the feedback applied to the adjacent cavity. Since the required feedback signal is measured in CDS, the amount of motion of the CTM at each injection frequency, Δx , can be calculated. The power fluctuation at each frequency in the detuned cavity, ΔP_{δ} , can thus be estimated, recalling Eq. 3.52, to be

$$\Delta P_{\delta} = \frac{1}{2} m \omega^2 \Delta x c \,, \tag{6.1}$$

where m is the CTM mass (100 g) and ω is the frequency of motion. Since the fractional power change in the cavity at a given frequency will be the same as the fractional change observed in the throughput power of that cavity, we can estimate the DC power in the detuned cavity, P_{δ}^{DC} , to be

$$P_{\delta}^{DC} = \Delta P_{\delta} \frac{P_T^{DC}}{\Delta P_T} \,, \tag{6.2}$$

where P_T^{DC} and ΔP_T are the DC throughput signal level and fluctuation thereupon respectively. Finally, the on-resonance power level of the detuned cavity, P_0 , can be estimated from the observed fractional power change between the on-resonance and detuned throughput power, dP_T :

$$P_0 = \frac{P_{\delta}^{DC}}{\mathrm{d}P_T} = \frac{m\omega^2 \Delta x c P_T^{DC}}{2\Delta P_T \mathrm{d}P_T} \,. \tag{6.3}$$

Performing this for five different detunings in each cavity yields the power estimates in Figure 6.6. Note that some of the injection frequencies have had to be discounted in both cavities; this is either because they were affected by the opto-mechanical resonances, and as such gave an incorrect prediction of intra-cavity power, or else were too close to the unity gain point of each control loop and the associated oscillation in the closed-loop transfer function. Expressing the result as the mean of the various measurements for each cavity:

$$P_0^L = (1239.29 \pm 64.44) \,\mathrm{W}\,,$$
 (6.4)

$$P_0^R = (2093.68 \pm 108.87) \,\mathrm{W}\,.$$
 (6.5)

To confirm that this method gives reasonable results, a different technique can be utilised to estimate the power of each on-resonance cavity. By monitoring the feedback signal applied to one cavity in the time domain as the other cavity loses lock (by way of switching off the feedback to that cavity) an estimate can be obtained of the corresponding shift in longitudinal position of the CTM – again with knowledge of the feedback strength cali-


Figure 6.6: Estimates of intra-cavity on-resonance power in the left-hand cavity (left) and right-hand cavity (right) by injecting fixed-frequency signals onto the input coupler of one cavity while detuned and reading out the feedback signal applied to the adjacent on-resonance cavity. This was repeated for five different detunings of each cavity and the results combined to give the mean values quoted.

bration presented in Figure 5.17. The position of the CTM depends only on the radiation pressure force on either side and the restoring force of the suspension, allowing calculation of the power in each cavity.

Figure 6.7 displays the time-series feedback signal data for each cavity as the other is removed from lock. Estimated DC positions of the CTM before and after loss of optical power on one side are obtained by averaging of the time series data up to and following that point. These give DC shifts in CTM position of $|\delta x_1| = 4.67 \,\mu\text{m}$ and $|\delta x_2| = 8.00 \,\mu\text{m}$. Solving simultaneous equations for each of the three CTM positions, one finds the two intra-cavity powers to be

$$P_0^L = \frac{mgc}{2\ell} \delta x_1 = (1144.36 \pm 12.84) \,\mathrm{W}\,, \tag{6.6}$$

$$P_0^R = \frac{mgc}{2\ell} \delta x_2 = (1962.34 \pm 21.13) \,\mathrm{W} \,. \tag{6.7}$$

126



Figure 6.7: DC position shifts of the CTM obtained from the change in feedback signal to one cavity as the other loses lock.

These estimates are consistent with those obtained using the detuned signal-injection method. However, the two measurements were taken several days apart, over which time slow drifts in cavity alignment and input power may have had an effect on the measured power. If we make the assumption that both cavities are affected equally by such effects, we can compare the ratio of on-resonance intra-cavity powers predicted by each method:

$$\left|\frac{P_0^L}{P_0^R}\right|_{\substack{\text{signal}\\\text{injection}}} = 0.592 \pm 0.044 \tag{6.8}$$

$$\left|\frac{P_0^L}{P_0^R}\right|_{\substack{\text{CTM}\\\text{displacement}}} = 0.583 \pm 0.009.$$
(6.9)

These predictions agree to within a small error, with discrepancy likely due to alignment shifts in each cavity as a result of temperature drifts (e.g. beam clipping, mirror imperfections, etc.), indicating that the "local readout" system can be used to provide reasonable estimates of intra-cavity power when no direct measurement is available. Note also that these power predictions agree well with the estimates obtained from the relative spring frequencies given in Section 6.2. These measurements serve to indicate the accuracy with which the optical spring is characterised; using measurements of the optical spring, we are able to confirm the optical power levels in each cavity to within the experimental error of an alternative method.

The measurements detailed in this section demonstrate to higher sensitivity than previous experiments the implementation of the local readout scheme for monitoring of optic motion in the local frame. In the presence of optical rigidity, this technique allows the position of the input coupler of one cavity to be monitored by the readout cavity in the absence of any interaction between that input coupler and the readout photons. Such a system could be included in an advanced gravitational wave detector by detuning each arm cavity and monitoring the motion of each input coupler with a separate laser subsystem [53]. As noted in Chapter 3, this technique could also be employed in parallel with detuned signal-recycling in a large-scale interferometer in order to increase the lower-frequency quantum sensitivity of the detector, which is otherwise reduced below the opto-mechanical resonance frequency [63].

6.4 POWER DEPENDENCE

As evidenced by Eq. 3.58, the DC optical spring restoring force is predicted to change linearly with intra-cavity power. Accordingly, the observed resonant frequency is expected to depend on the square root of the power. An experiment was set up to verify that the measurement technique allows this basic point to be confirmed. Figure 6.8 displays two opto-mechanical resonances, corresponding to approximately the same detuning and approximately a factor two difference in input power to the cavity. Accessing the recorded throughput DC signal from that cavity after measurements have been taken, the power change and each detuning are determined to be $P_{high}/P_{low} = 2.208$, $\delta_{high} = 0.81 \gamma$ and $\delta_{low} = 0.79 \gamma$. These values are calculated again using the mean of the throughput DC signals for the duration of each measurement.

We expect the spring frequency to reduce between the two power regimes by a factor of $\sqrt{P_{high}/P_{low}} = 1.48$. However, since the detuning in each case is slightly different, we can apply each value to Eq. 3.74 to correct this factor; in so doing, the expected ratio reduces to 1.46. The frequencies of opto-mechanical resonances for high and low power are estimated from a fit to the measurements to be $f_{high} = (277\pm3)$ Hz and $f_{low} = (192\pm2)$ Hz respectively. This yields a ratio of resonant frequencies of $f_{high}/f_{low} = 1.435\pm0.03$. This compares very favourably with our estimated ratio of 1.46, confirming the behaviour of the opto-mechanical effect is in line with the theoretical dependence on intra-cavity power.



Figure 6.8: Open-loop transfer functions showing opto-mechanical resonances for two different input powers. The spring frequency is observed to reduce appropriately with reduction of the input power.

6.5 ANTI-SPRINGS AND LOOP INSTABILITY

Contrary to the positive spring measurements presented in the preceding sections, in which the cavity gain is resonantly enhanced at the corresponding opto-mechanical resonance frequencies thanks to a strong restoring force and weak anti-damping effect, when the cavity is negatively detuned to form an anti-spring with an *anti*-restoring force and *positive* damping, the effective sign of the poles in the equation of motion change, with the result that no resonant behaviour is associated with the upshifted pendulum mode. This is evident in Figure 6.9a, in which the left-hand cavity has been detuned approximately equally in both the positive and negative directions to produce a spring and an anti-spring. Note the reduced gain through the region of the spring frequency in the anti-spring case, and the increase in phase margin towards 200 Hz at the expense of this lower gain margin.

In particular, if the gain is reduced too far, by way of creating a particularly strong anti-spring effect, the gain margin approaches unity and the closed-loop gain is observed to have a tendency to oscillate at that frequency. This effect is presented in the red plot in Figure 6.9b, where the right-hand cavity has been detuned close to the maximum



Figure 6.9: (a) Open-loop transfer functions in the left-hand cavity with detunings of $\delta_L = 1 \gamma_L$ and $-0.97 \gamma_L$ to show spring and anti-spring responses for nominally equal but opposite detuning. (b) Open- and closed-loop transfer functions in the right-hand cavity with an anti-spring present, with and without the addition of a resonant gain filter at 179 Hz, with Q of 4 and height of 20 dB. The addition of the resonant gain (black plots) is observed to mitigate the damped oscillation observed in the CLTF as the anti-spring approaches unity gain (red plots). The on-resonance ($\delta_R = 0$) OLTF is included to compare with the reduced gain when an anti-spring is present.

upshifted pendulum frequency. Here, the open-loop gain is observed to approach unity at ~ 179 Hz. The companion closed-loop transfer function plot then indicates a peak in response at that frequency in the closed loop gain with magnitude > 30 dB. On detuning much further than $\delta_R = -0.4 \gamma_R$ the loop became unstable, with the resulting oscillation eventually leading to saturation of signals and loss of lock. The unstable unity gain crossing may be removed if a resonant gain filter is added to the loop at the frequency of instability. This has been achieved in Figure 6.9b, wherein the black lines document the open- and closed-loop response for the same detuning when a resonant gain filter is added at 179 Hz, with quality factor 4 and height of 20 dB. The gain margin in the OLTF is observed to increase dramatically, and the instability at 179 Hz in the CLTF is removed.

This indicates that while instabilities may arise in the closed-loop system in the presence of negatively-detuned opto-mechanical couplings, they can be mitigated with the application of narrowband gain. This is an important observation as regards the implementation of anti-springs in a large-scale detector (e.g. to bring stability to an otherwise unstable single-optical-spring scheme). Note that in a system allowing a higher unity gain point than is possible using our digital control system, the gain of the loop could conceivably be increased by a flat scalar value to achieve the same goal without the complexity of a resonant filter.

6.6 Dynamic Response of the Modified Error Sig-NAL

As shown in Chapter 5, the method of widening the error signal by way of dividing it by the normalised power in the cavity allows the cavity to be detuned further than would be possible using the "normal" error signal. As we demonstrate here, this method modifies the frequency response (magnitude and phase) of the error signals. This change has implications for the stability of the system as larger detunings are applied.

Figure 6.10a shows open-loop transfer function measurements of the right-hand cavity when detuned from approximately $1 \gamma_R$ to $5 \gamma_R$ in steps of $1 \gamma_R$, as determined again from the relative DC throughput signal sizes. Notably, the phase in each detuned signal is observed to increase by between 15° and 30° at frequencies above 200 Hz, relative to the on-resonance case, $\delta_R = 0$ (with the notable exception of the $\delta_R = 1 \gamma_R$ plot in which the phase change associated with the spring resonance is much higher in frequency). This extra phase lead, conceivably allowing the unity gain point of the system to be lowered towards 200 Hz in the presence of a highly-detuned optical spring, can be attributed to the modified error signal.

Figure 6.10b shows a simulation of the frequency response of the resulting modified PDH error signal after the division has taken place – this represents the AC counterpart to the DC widened PDH signal shown in Figure 5.12. In the simulation, the effect of detuning is included by a reduction of the power in the cavity. Recall that the frequency response of the cavity power signal has a $1/f^2$ slope above the linewidth frequency, which we estimate to be around 700 Hz, and that the PDH error signal responds with a shallower 1/f slope over the same frequencies. When the DC power in the throughput field becomes comparable to that in the error signal, then, the response of the widened error signal will change, with a peak in response appearing in the kHz range due to the 1/f slope becoming less dominant. Crucially, this extra frequency-shaping of the modified error signal yields accumulation of extra phase that increases with detuning (or equivalently, decreases with power in the throughput signal).

The dynamic response behaviour of the modified Pound-Drever-Hall signal presented in Figure 6.10b matches up well with observations during the measurements presented in Figure 6.10a. While performing preliminary measurements, an increase in response consistent with the simulation was noted in the kHz region, such that for maintenance of stability at the UGP a supplementary low-pass filter was required in the digital servo. In future, this filtering could be applied to the error signal before or after the division point automatically as detuning is increased. Further, the observed phase increase throughout the measurement region is consistent with that predicted by the modelling. To our knowledge, this is the first time this effect has been observed and characterised.

It is possible that this effect could be further exploited in order to lower the unity gain of the system still further, by way of either adjusting the division with appropriate offsets or increasing the detuning to accumulate extra phase at lower frequencies; however, lowering of the unity gain point in accordance with this would then require dynamic adjustment of the control loop gain shaping in order to maintain adequate gain margin for stability. This is anticipated as a future goal for a follow-on experiment.

6.7 DOUBLE SPRINGS: MODIFICATION OF DYNAMICS

Here we examine the more complex case where an optical spring is present in each cavity, such that the stiffness and resulting measured resonant frequency of the CTM



Figure 6.10: (a) Open-loop transfer functions in the right-hand cavity for detunings of approximately 1 to 5 linewidths. (b) Optickle simulations of the modified error signal response, predicting increased phase through the measurement region relative to the on-resonance case. At each increase in magnitude between the blue plots the detuning has been increased in the simulation.

suspension are modified by the presence of opto-mechanical effects on both sides.

Figure 6.11a shows in-loop OLTF measurements of the right-hand cavity. Throughout these measurements, the right-hand cavity maintains a detuning of $0.7 \gamma_R$. For each plotted curve, the detuning applied to the left-hand cavity, external to the control loop in which measurements are being performed, is adjusted, approaching the maximum strength coupling at $0.57 \gamma_L$. As the left-hand detuning is adjusted, the observed resonance is seen to shift in frequency by up to 10%. From examination of the recorded throughput signal of the right-hand cavity the optical power is observed not to drift appreciably over the half-hour of total measurement time; as such, we assume that the shifts in resonant frequency observed are due predominantly to the coupling of the two optical springs at the CTM.

Of particular note is the gain of the resonance, observed to increase as a spring is introduced in the adjacent cavity. With reference to Eq. 3.21, we note that this increase is in accordance with the expected increase in susceptibility as the reduced mass of the now coupled cavity system changes. The magnitude of the right-hand cavity resonance is observed to increase by 2.5 dB at the strongest left-hand cavity rigidity, or a factor of approximately 1.33. Since the CTM is now rigidly coupled to the LTM at low frequencies, its effective mass has now doubled; pushing on the CTM will rigidly push on the LTM as well. This means that the reduced mass of the right-hand cavity has increased from 0.05 kg when it is bounded by two 100 g masses up to 0.067 kg when bounded by one 100 g and one effective 200 g mass; this results in a change in susceptibility of the right-hand cavity of $\Delta \chi = 1.33$, in good agreement with the observed resonant magnitude change.

Note that only a *single* resonance is observed in this measurement, in contrast with the modelling of the doubly-resonant system presented in Figure 6.2 (inset in Figure 6.11 for comparison) which indicates that *two* coupled resonances ought to be present when the cavities are mechanically coupled and both detuned. However, model and measurement are reconciled when the loop gain of each cavity is considered – this is *not* present in the Optickle models, which consider only the circulating fields and mechanical transfer functions of the system. In the presence of the control loops, each cavity is acted upon *via* the input test mass to counteract any excess relative motion that would modulate the total length. The result of this is that *outside* each control loop the opto-mechanical resonance is not observed, because the control loop is feeding back a signal of appropriate size to compensate the effect; instead, only the extra stiffness arising from the coupling of each input test mass to the CTM is observed in the adjacent cavity. This manifests once



Figure 6.11: (a) Open-loop transfer function measurements of the right-hand cavity. The right-hand cavity is set to a detuning of $\delta_R = 0.7 \gamma_R$ while the left is detuned from $\delta_L = 0.076 \gamma_L$ to $\delta_L = 0.539 \gamma_L$, and magnitude and frequency of this right-hand opto-mechanical resonance is observed to be modified by the rigidity of the left. Inset: copy of the simulation performed in Figure 6.2, to show the two coupled resonances that are expected before the loop gain is taken into account. (b) Corresponding transfer function spectra of the left-hand cavity feedback relative to the right-hand cavity injection, normalised to the case where $\delta_L = 0$. Increased stiffness is observed beneath the resonant features of both cavities.

again as changes to the susceptibility of the coupled-spring system.

Figure 6.11b illustrates the change in susceptibility at the CTM. The plot shows transfer functions from the excitation signal injected into the right-hand cavity to the feedback signal applied to the left-hand cavity to maintain its length, taken at the same time as the in-loop measurements of Figure 6.11a. These transfer functions are presented relative to the case in which the left-hand cavity is on-resonance; as such, they can be understood to indicate the changes to the susceptibility of the CTM as the left-hand cavity is detuned from resonance, while the right-hand cavity maintains a detuning of $\delta_R = 0.7 \gamma_R$.

As the detuning of the left-hand cavity is increased, a decrease in the magnitude of the feedback signal is observed towards the low-frequency end of Figure 6.11b. This indicates a smaller feedback signal is required at the LTM to correct for CTM motion, and hence the relative motion of the two has decreased thanks to stronger opto-mechanical coupling. The increase in magnitude at higher frequencies is due to the presence of the optical spring in the left-hand cavity, resulting in larger susceptibility about the resonant frequency. The shaping of this increase in susceptibility in Figure 6.11b results in the shifting in peak frequency of the right-hand cavity resonances observed in Figure 6.11a.

These measurements confirm that the opto-mechanical rigidities couple as is expected from the associated changes to the susceptibility of each cavity. Since the measurements are performed beneath the unity gain point of each cavity, coupled resonant features are observed to be compensated by the loop gain. This is an important observation as regards the "optical bar" proposal for an interferometric gravitational wave detector, whereby an optical spring in each arm couples to either side of a shared mass, with opto-mechanical resonant frequencies likely below the unity gain points of the control systems.

6.8 DOUBLE SPRINGS: SIMULTANEOUS MEASUREMENT

In the previously presented double-spring case, the effect of the rigid coupling in the adjacent cavity is observed without any obvious coupled resonant features. Accordingly, we have developed a means of measurement in which these coupled resonant features are observed. This is based upon simultaneous injection of an identical signal into *both* cavities, resulting in simultaneous measurement of the two open-loop transfer functions. Since a coherent signal is now observed in both cavities at the CTM, shaped by the gain in the feedback loops to each cavity, the closed-loop gain shaping of one cavity, itself directly related to the opto-mechanical resonant structure of that cavity, can couple into

the other. This is due solely to motion of the CTM, applied *via* the optical coupling to each ITM, which is then read out coherently by the measurement of each cavity.

Four effects contribute to the resulting transfer function shapes, each dramatically altered from the previous double-spring incarnations in which only a single cavity was measured at a time. These are described in the following four sections.

6.8.1 Coupling of Spring Effects into On-Resonance Cavity

The first observation may be made for the simplest case, in which one cavity is held on-resonance and the other detuned therefrom and an excitation signal is injected onto both the LTM and RTM positions simultaneously. In previous single-cavity measurements, this would result in a well-defined opto-mechanical resonance measurement in the detuned cavity, and an on-resonance OLTF measurement, similar to Figure 5.16b, in the other. While the former remains true in this instance, we find that the adjacent on-resonance cavity now also displays resonant features consistent with opto-mechanical coupling. This is shown in Figure 6.12.

In Figure 6.12a, the left-hand cavity is detuned to $0.9 \gamma_R$ and the right-hand cavity initially held on resonance. This results in the familiar resonant feature in the left-hand cavity, but also an obvious drop in phase towards 100 Hz in the right-hand cavity consistent with a weak opto-mechanical effect in that cavity as well. This is further confirmed when a small negative detuning is applied to the right-hand cavity, resulting in a downwards shift below the measurement band of that phase feature, as would be expected from Eq. 3.76.

The effect is far more pronounced in Figure 6.12b, in which results are shown for the case where the right-hand cavity is detuned and the left is on-resonance, thanks to the higher power circulating in the right-hand cavity yielding a stronger optical spring effect. This effect can be attributed to the simultaneous measurement of both cavities. Below the associated resonant frequency of the detuned cavity, the CTM is coupled rigidly to the input mirror. This means that the injected white noise signal moves the CTM dynamically according to the power gain in the cavity, which is now subject to a broad resonance at that frequency. The adjacent on-resonance cavity then reads out this excess motion of the CTM coherently, resulting in a modification to the expected $\delta_{\gamma} = 0$ transfer function presented as the black plot in each.



Figure 6.12: Observation of an optical spring in one cavity coupling into the measured OLTF of the adjacent on-resonance cavity. (a) LHCav detuned, RHCav on-resonance and then with a small negative detuning applied. (b) LHCav on-resonance and RHCav detuned. The black line in each indicates the equivalent "both cavities on-resonance" plot.



Figure 6.13: Simultaneously-performed open-loop transfer functions showing apparent suppression of CTM motion as observed by both cavities (thick lines) when compared to the equivalent single-spring OLTFs (thin lines).

6.8.2 Suppression of CTM Motion

The second observation made is that, when there is a coherent signal injected onto both ITMs which couples to the CTM through optical rigidity in *both* cavities, the apparent motion of the CTM is suppressed. This was observed most prominently in the measurements presented in Figure 6.13. When compared with the equivalent single-spring measurements, for which the adjacent cavity is not detuned, we note that the observed CTM response has been suppressed in both cavities, and significantly in the left-hand cavity, with its larger detuning resulting in lower optical power and weaker opto-mechanical coupling than in the right-hand cavity. The suppressed response in the right-hand cavity is less obvious, again due to its stronger coupling and higher power dominating over the left-hand cavity effect. In particular, the OLTF of the left-hand cavity in the double-spring measurement is observed to fall below unity below 200 Hz. Were this a physical effect upon the loop gain, this ought to result in a unity gain instability at this frequency; this backs up the notion that this effect is imposed upon the measurement only due to the apparent reduced motion of the CTM, rather than on the loop gain itself. Note the similarity between this measurement and the "annihilation" double-spring regime proposed



Figure 6.14: Simultaneous measurements showing coupling of a resonant from the closedloop response of one cavity into the open-loop response of the other. Inset: copy of the simulation performed in Figure 6.2, for comparison with the red plot, wherein the loop gain has been lowered by the anti-spring feature with the result that a coherent oscillation couples from the left-hand cavity into the right-hand cavity OLTF.

by Rehbein *et al.* [61] and discussed in Section 3.2.5, whereby the two springs are tuned to "trap" the CTM longitudinally by application of equal optical force on either side.

6.8.3 COUPLING OF LOOP GAIN FEATURES

The third effect attributed to the simultaneous injection and readout measurement in both cavities is the coupling of closed-loop gain features into open-loop transfer functions. An example of this is given in Figure 6.14, in which a positive spring has been created by detuning the right-hand cavity by one linewidth. The left-hand cavity is then negatively detuned to $-0.58 \gamma_L$. This is very close to the strongest possible anti-spring effect, with the result that the loop gain of that cavity is observed to drop close to unity at 180 Hz. As observed previously (e.g. Figure 6.9b), the resulting damped oscillation is not visible in the open-loop gain of that cavity; however, it nevertheless results in motion of the CTM at the frequencies of this peak in response, such that it couples coherently into the *open*-loop transfer function of the right-hand cavity. Further, it is interesting to note that this right-hand cavity response is remarkably similar in both magnitude and phase performance to the modelled case presented in Figure 6.2 (included as an inset in Figure 6.14), where using Optickle we have simulated two *positive* springs coupling at the CTM and observed the effect on the OLTF of one cavity. In this physical case, the lower resonance of Figure 6.2 has been replaced by the observed loop gain response peak at 180 Hz to similar effect.

This effect can, as was demonstrated for the anti-spring measured in Section 6.5, be mitigated by application of electronic gain in the left-hand cavity in which the anti-spring is present. This may be required in a system such as that proposed by Rehbein *et al.* [61] in which an anti-spring is combined with a positive spring in order to form a stable net opto-mechanical effect, or similarly in the optical bar interferometer configuration [62].

6.8.4 COUPLING OF SPRING STRENGTHS

The final effect observed is not unique to the simultaneous measurement, being the expected shifting of resonant features in one cavity as the other is swept through a number of different detunings (as was touched upon already in Section 6.7). Directly responsible for this is the coupling of the two optical springs at the CTM, with the associated resonances modified as the sum of the opto-mechanical stiffnesses and damping. See Figure 6.15.

As in Section 6.7, these measurements confirm the translation of the restoring force and damping coefficients of the overall rigidity predicted by Eq. 3.76. With careful tuning of each spring, the combination could be made stable, as described by Figure 3.8. If such a scheme is to be employed, the observations of couplings between optical springs and loop gains made earlier in this chapter will be very important – for example, in the case where a strong optical anti-spring is required to move the opto-mechanical system towards stability, in so doing lowering the gain of that control loop towards unity. Further investigations into stable spring combinations in this system are planned for the future.

6.9 SUMMED RESPONSE

Figure 6.16 presents another simultaneous measurement of the two cavities with both detuned from resonance. In this instance, a third readout channel has been included in the measurement to monitor the sum response of the two cavities. This channel is the sum of the two time-domain error signals recorded by the CDS. The plot indicates the sum response OLTF achieves resonant signal enhancement of up to 10 dB larger than



Figure 6.15: Left-hand cavity OLTF measurements from simultaneous injections into both cavities. The right-hand cavity maintains a detuning equal to its linewidth yielding a positive spring in the ~ 300 Hz range that couples through to left-hand cavity as a result of the measurement process, as was shown in Figure 6.12. The left-hand cavity detuning is increased negatively, introducing an anti-spring that shifts the lower resonant feature, corresponding to the observed in-loop spring feature of the left-hand cavity, further down in frequency and ultimately below the measurement band.

either of the individual-cavity measurements in the $\sim 200-300$ Hz range. Depending on the relative phase of the two cavity signals, it is possible to increase or decrease the sum response through the measurement region.

This result has important implications as regards the quantum sensitivity of the cavity. As was shown in early simulations (Figure 3.10), for particular tunings of each optical spring the resulting sum quantum sensitivity of the coupled system can dip beneath the Standard Quantum Limit over a wider range of frequencies than is possible with a single optical spring, and potentially to higher sensitivity. We have shown the potential for employing such a scheme in a prototype-scale system; further work will be required to demonstrate this technique as a viable means of surpassing the SQL in a large-scale interferometer.



Figure 6.16: Demonstration of the summed response of both cavities to the simultaneous injection method. The sum response is observed to peak at approximately 280 Hz, with resonant response approximately 10-14 dB higher than the individual cavity measurements. The summed response for both cavities held at resonance is also given for reference.

6.10 CONCLUSIONS

Resonant features consistent with strong opto-mechanical rigidity effects in both cavities have been measured to high precision, and the corresponding behaviour with cavity power and detuning confirmed experimentally. Upon examining coupled-cavity regimes, we find that optical springs couple with the control loop gains of each cavity in complex but characterisable ways. Importantly, opto-mechanically-induced instabilities in the control loops have been demonstrated to be correctable with targeted application of gain. A simultaneous measurement scheme has been developed that further illuminates the transfer of resonant features between the circulating optical fields and the control loops *via* the purely mechanical coupling at the Central Test Mass. The implications of each of the results detailed here as regards advanced gravitational wave detectors will be discussed in Chapter 7.

VII Conclusions and Future Work

Schemes involving one or more optical spring are expected to be included in the design of advanced gravitational wave detectors. In the short term, this is likely to take the form of detuned signal recycling regimes [71, 72] as were demonstrated at GEO 600 and are included in the Advanced LIGO design. It has been proposed [61] that a second detuned carrier field is included in this scheme to stabilise the optical spring without increased complexity in the control systems or extra classical noise. This could be further enhanced using the local readout [63] or optical bar [62] schemes. Further to the advantages provided by multiple optical springs as regards stability and re-shaping of the quantum sensitivity of the instrument, it is imperative that the practical aspects of their implementation are investigated and well-understood.

We have presented the design and construction of an experiment to facilitate characterisation of a range of optical spring regimes. A number of new optic suspensions have been designed and built to seismically-isolate the experiment through the measurement region, the mechanical modes of which are measured and well-defined. The suspensions were designed to maintain stable optic alignment over the course of each measurement to minimise drifts in the intra-cavity power. Control system designs have been presented that yield reliable, repeatable acquisition of lock of a high-finesse optical cavity with feedback to the optic positions, despite speed restrictions imposed by the digital feedback system, by manipulation of the error signal. The loop gain of each locked cavity has been well-characterised, allowing optical spring effects to be identified when present. Key parameters of this system, such as detuning ranges, optical power and mirror masses, are shown to be capable of representing a scaled-down version of the double-optical-spring enhancement proposed by Rehbein *et al.* [61] as an upgrade to Advanced LIGO.

The stable suspension design and stabilisation of the laser frequency to an external 10 m reference cavity result in very little drift in cavity power level, detuning and alignment over

the course of each measurement. As a result, the MOSES apparatus has been shown to be capable of measuring optical springs to high precision. The sampling of appropriate signals from each cavity allows accurate characterisation of the state of the opto-mechanical system, such that optical spring resonant features are distinguishable from loop gain effects. This capability is very important for the optical spring measurements presented in this thesis and those that are anticipated in follow-on experiments.

The operation of the system as an analogue for local readout was studied, whereby one cavity is used as a readout scheme to track the position of the optics of the adjacent detuned cavity by interaction with the Central Test Mass alone. This was demonstrated with a measurement sensitivity several orders of magnitude better than a previous study in the same laboratory [53]. The local readout method was then employed to measure the steady displacement of the CTM, allowing deduction of the on-resonance intra-cavity power. The comparison of this work with pior demonstrations stresses the importance of the choice of readout system, should this technique be employed in a large-scale detector to enhance the low-frequency sensitivity.

The dependence of the optical spring resonance upon power and detuning was confirmed, with a factor two decrease in cavity power yielding a factor $\sqrt{2}$ decrease in the resonant frequency of the optical spring. The desired measurement precision required analysis of five minutes of data; this was only possible because the system remained stable in power and detuning over these time-scales and longer. This measurement again confirms the suitability of the apparatus for the characterisation of opto-mechanical effects in the few hundred Hz range, comparable to the region of peak sensitivity in the large-scale instruments.

The interactions between opto-mechanical effects and the control loops designed to maintain stability were studied in a number of different regimes. Firstly, strong anti-springs were shown to be capable of reducing the loop gain towards unity, such that at that frequency a peak in response arises that threatens to destabilise the loop. This was shown to be correctable with targeted application of gain to increase the gain margin through the spring region. This observation could be of particular importance for systems operating in the stable double-spring regime [61] which must necessarily include optical springs and anti-springs (or, indeed, an opto-mechanical system on any scale where an anti-spring is stabilised by electronic feedback).

The technique of widening the error signal was shown to include a dynamic response

change at frequencies above the cavity linewidth as the detuning is increased. This serves to increase the phase towards lower frequencies, such that with a larger detuning $(> 1 \gamma)$ and adaptive gain scheme the unity gain point of the loop could potentially be lowered beneath that which would be possible with an on-resonance or low-detuned cavity. This is anticipated as a future line of investigation for a follow-on experiment. To our knowledge, this is the first time this technique has been applied to facilitate wider detuning of an optical cavity in the creation of opto-mechanical rigidity than is conventionally possible, and the first time the observed dynamic response of the modified error signal has been characterised.

Coupled optical springs were observed to behave as expected. The susceptibility of the CTM was modified in a manner consistent with the shifts in spring resonant frequencies and magnitudes measured. Importantly, we confirm that resonant features from one cavity are not observed to couple into the open-loop measurements of the other cavity; rather, it is only the modified susceptibility resulting from the DC optical stiffness in each cavity that couples. This result applies in particular to the design of the optical bar detectors, whereby the end mass of each interferometer arm is coupled to a central mirror with optical springs; the opto-mechanical rigidity in each arm will be modified *via* the interaction at the shared optic, and it is likely that the optical springs will be situated beneath the unity gain point of the feedback loops for stability.

The final set of measurements was designed to probe the interactions of optical springs in both cavities with the respective control loops. To achieve this a method of simultaneous signal injections into both cavities (*via* their ITMs) was developed and applied. This scheme indicates a number of further effects exist in the couplings between optomechanical rigidity and loop gains. The first observation indicated that with this scheme the resonant response of the optical spring in one cavity couples into the on-resonance response of the adjacent cavity. Secondly, with an optical spring in both cavities the simultaneous injection was observed to result in radiation pressure couplings to both sides of the CTM, effectively suppressing its motion. A third measurement indicated that closed-loop gain features of one cavity (e.g. resonant responses close to unity gain) can couple into the open-loop gain of the adjacent cavity. Each of these effects is interpreted as resulting from the simultaneous injection of coherent signals upon the CTM from either cavity. These effects will be important for any multiple-spring system in which a coherent signal influences coupled control loops.

The investigations presented in this thesis are designed to inform the practical design and

operation of future gravitational wave detectors operating with multiple optical springs in order to re-shape the quantum noise spectrum or to increase stability without excess classical noise in the control loops. We have demonstrated complex coupling between optical rigidity and the control loops that govern the system, and methods by which stability may be improved over certain frequency bands with further detuning of the cavity or increased electronic gain.

Having characterised the dual-cavity system in terms both of its range of observable optical spring behaviour and novel modes of operation, the next step will be to examine further regimes of multiple optical rigidity. It is expected that a second source-laser will be included in the input path, allowing a sub-carrier to be injected into one or both cavities to investigate optical trapping. This technique is yet to be demonstrated on a prototype (100 g) scale. In addition, the cavities may be coupled both mechanically and optically, by way of aligning both down a common optical axis through the CTM such that transmitted fields from one cavity may couple into the adjacent. This regime brings with it added challenges in control and characterisation, the study of which will further examine the challenges faced by a large-scale gravitational wave detector incorporating optical rigidity.

i Circuit Diagrams

A.1 SIGNAL PHOTODETECTOR

Figure A.1 shows the circuit diagram for the standard photodetector design, used to sample reflected and throughput signals from each cavity. The output contains a simple high-pass filter to separate AC and DC signals.



Figure A.1: Schematic of broadband photodetector design.

A.2 LOCAL CONTROL COIL DRIVER

Figure A.2 shows the local control circuitry. This is a standard design, used to distribute signals output by the Arduino PWM channels to the alignment coils on each suspension. Each pair of output channels is connected to a pair of coils to drive either tilt or rotation.



Figure A.2: Circuit diagram for local control coil driver.

A.3 GLOBAL CONTROL COIL DRIVER

Figure A.3 shows the global control circuitry. The longitudinal feedback signal from the CDS is split and routed to all four coils with a gain of 20 dB in each path. Pitch and yaw control signals can also be added in at any time, with each signal being split into a differential pair before being summed in and sent to the appropriate pair of coils.



Figure A.3: Circuit diagram for global control coil driver.

BIBLIOGRAPHY

- A. Einstein, "Zur Elektrodynamik bewegter Körper. (German) [On the electrodynamics of moving bodies]," German, Annalen der Physik, vol. 322, no. 10, pp. 891– 921, 1905.
- [2] —, "Die Grundlagen der Allgemeinene Relativitätstheorie. (German) [The foundations of the Theory of General Relativity]," German, Annalen der Physik, vol. 354, no. 7, pp. 769–822, 1916.
- [3] B. Schutz, A First Course in General Relativity. May 2009.
- B. S. Sathyaprakash and B. F. Schutz, "Physics, Astrophysics and Cosmology with Gravitational Waves," *Living Reviews in Relativity*, vol. 12, p. 2, Mar. 2009. arXiv: 0903.0338 [gr-qc].
- [5] S. Chandrasekhar, "The Maximum Mass of Ideal White Dwarfs," Astrophysical Journal, vol. 74, p. 81, Jul. 1931.
- [6] J. Aasi, J. Abadie, B. P. Abbott, et al., "Gravitational Waves from Known Pulsars: Results from the Initial Detector Era," Astrophysical Journal, vol. 785, 119, p. 119, Apr. 2014. arXiv: 1309.4027 [astro-ph.HE].
- [7] K. Riles, "Gravitational waves: Sources, detectors and searches," Progress in Particle and Nuclear Physics, vol. 68, pp. 1–54, Jan. 2013. arXiv: 1209.0667 [hep-ex].
- [8] R. L. Forward, D. Zipoy, and J. Weber, "Upper Limit for Interstellar Millicycle Gravitational Radiation," *Nature*, vol. 189, p. 473, Feb. 1961.
- J. Weber, "Evidence for Discovery of Gravitational Radiation," *Physical Review Letters*, vol. 22, pp. 1320–1324, Jun. 1969.
- [10] M. Pitkin, S. Reid, S. Rowan, and J. Hough, "Gravitational Wave Detection by Interferometry (Ground and Space)," *Living Reviews in Relativity*, vol. 14, p. 5, Jul. 2011. arXiv: 1102.3355 [astro-ph.IM].
- [11] S. E. Whitcomb, "Ground-based gravitational-wave detection: now and future," *Classical and Quantum Gravity*, vol. 25, no. 11, 114013, p. 114013, Jun. 2008.

- [12] R. A. Hulse and J. H. Taylor, "Discovery of a pulsar in a binary system," Astrophysical Journal, vol. 195, pp. L51–L53, Jan. 1975.
- J. H. Taylor, "Binary pulsars and relativistic gravity," Rev. Mod. Phys., vol. 66, pp. 711–719, 3 Jul. 1994.
- [14] A. G. Lyne, M. Burgay, M. Kramer, et al., "A Double-Pulsar System: A Rare Laboratory for Relativistic Gravity and Plasma Physics," *Science*, vol. 303, pp. 1153– 1157, Feb. 2004. eprint: astro-ph/0401086.
- [15] B. P. Abbott, R. Abbott, R. Adhikari, et al., "LIGO: the Laser Interferometer Gravitational-Wave Observatory," *Reports on Progress in Physics*, vol. 72, no. 7, 076901, p. 076 901, Jul. 2009. arXiv: 0711.3041 [gr-qc].
- [16] H. Grote, "The geo600 status," Class. Quantum Grav., vol. 27, p. 084003, 2010.
- [17] T. Accadia, F. Acernese, M. Alshourbagy, et al., "Virgo: a laser interferometer to detect gravitational waves," *Journal of Instrumentation*, vol. 7, p. 3012, 2012.
- [18] J. Aasi, B. P. Abbott, R. Abbott, et al., "Advanced ligo," Classical and Quantum Gravity, vol. 32, no. 7, p. 074 001, 2015.
- [19] F. Acernese, M. Agathos, K. Agatsuma, et al., "Advanced Virgo: a second-generation interferometric gravitational wave detector," *Classical and Quantum Gravity*, vol. 32, no. 2, 024001, p. 024001, Jan. 2015. arXiv: 1408.3978 [gr-qc].
- [20] H. Luck, C. Affeldt, J. Degallaix, et al., "The upgrade of GEO 600," Journal of Physics Conference Series, vol. 228, no. 1, 012012, p. 012012, 2010. arXiv: 1004.0339 [gr-qc].
- [21] K. Somiya, "Detector configuration of KAGRA-the Japanese cryogenic gravitationalwave detector," *Classical and Quantum Gravity*, vol. 29, no. 12, 124007, p. 124007, Jun. 2012. arXiv: 1111.7185 [gr-qc].
- [22] S. Sunil and D. G. Blair, "Vacuum system requirement for a 5 km baseline of gravitational-wave detector," *Journal of Physics Conference Series*, vol. 114, no. 1, 012025, p. 012 025, May 2008.
- J. C. Driggers, M. Evans, K. Pepper, and R. Adhikari, "Active noise cancellation in a suspended interferometer," *Review of Scientific Instruments*, vol. 83, no. 2, p. 024 501, Feb. 2012. arXiv: 1112.2224 [gr-qc].
- [24] S. Hild, M. Abernathy, F. Acernese, et al., "Sensitivity studies for third-generation gravitational wave observatories," *Classical and Quantum Gravity*, vol. 28, no. 9, 094013, p. 094013, May 2011. arXiv: 1012.0908 [gr-qc].

- [25] G. Hammond, S. Hild, and M. Pitkin, "Advanced technologies for future groundbased, laser-interferometric gravitational wave detectors," *Journal of Modern Optics*, vol. 61, p. 10, Dec. 2014. arXiv: 1402.4616 [astro-ph.IM].
- [26] H. B. Callen and T. A. Welton, "Irreversibility and Generalized Noise," *Physical Review*, vol. 83, pp. 34–40, Jul. 1951.
- [27] S. D. Penn, A. Ageev, D. Busby, et al., "Frequency and surface dependence of the mechanical loss in fused silica," *Physics Letters A*, vol. 352, pp. 3–6, Mar. 2006. eprint: gr-qc/0507097.
- [28] V. P. Mitrofanov and K. V. Tokmakov, "Effect of heating on dissipation of mechanical energy in fused silica fibers," *Physics Letters A*, vol. 308, pp. 212–218, Feb. 2003.
- [29] G. Cagnoli and P. A. Willems, "Effects of nonlinear thermoelastic damping in highly stressed fibers," *Phys. Rev. B*, vol. 65, no. 17, 174111, p. 174111, May 2002.
- [30] G. D. Hammond, A. V. Cumming, J. Hough, et al., "Reducing the suspension thermal noise of advanced gravitational wave detectors," *Classical and Quantum Gravity*, vol. 29, no. 12, 124009, p. 124009, Jun. 2012.
- [31] G. M. Harry, H. Armandula, E. Black, et al., "Thermal noise from optical coatings in gravitational wave detectors," Appl. Opt., vol. 45, pp. 1569–1574, Mar. 2006.
- [32] B. W. Barr, K. A. Strain, and C. J. Killow, "Laser amplitude stabilization for advanced interferometric gravitational wave detectors," *Classical and Quantum Gravity*, vol. 22, pp. 4279–4283, Oct. 2005.
- [33] K. D. Skeldon, K. A. Strain, A. I. Grant, and J. Hough, "Test of an 18-m-long suspended modecleaner cavity," *Review of Scientific Instruments*, vol. 67, pp. 2443– 2448, Jul. 1996.
- [34] A. Araya, N. Mio, K. Tsubono, et al., "Optical mode cleaner with suspended mirrors," Appl. Opt., vol. 36, pp. 1446–1453, Feb. 1997.
- [35] J. Mizuno, "Comparison of optical configurations for laser-interferometric gravitaionalwave detectors," PhD thesis, Universität Hannover, 1995.
- [36] A. Buonanno and Y. Chen, "Scaling law in signal recycled laser-interferometer gravitational-wave detectors," *Phys. Rev. D*, vol. 67, no. 6, 062002, p. 062002, Mar. 2003. eprint: gr-qc/0208048.
- [37] S. L. Danilishin and F. Y. Khalili, "Quantum measurement theory in gravitational wave detectors," *Living Reviews in Relativity*, 2012.

- [38] Y. Chen, "Sagnac interferometer as a speed-meter-type, quantum-nondemolition gravitational-wave detector," *Phys. Rev. D*, vol. 67, no. 12, 122004, p. 122004, Jun. 2003. eprint: gr-qc/0208051.
- [39] C. Gräf, B. W. Barr, A. S. Bell, et al., "Design of a speed meter interferometer proof-of-principle experiment," *Classical and Quantum Gravity*, vol. 31, no. 21, 215009, p. 215009, Nov. 2014. arXiv: 1405.2783 [gr-qc].
- [40] https://www.elisascience.org.
- [41] F. Antonucci, M. Armano, H. Audley, et al., "The LISA Pathfinder mission," Classical and Quantum Gravity, vol. 29, no. 12, 124014, p. 124014, Jun. 2012.
- [42] Y. Chen, "Macroscopic quantum mechanics: theory and experimental concepts of optomechanics," *Journal of Physics B Atomic Molecular Physics*, vol. 46, no. 10, 104001, p. 104001, May 2013. arXiv: 1302.1924 [quant-ph].
- [43] A. Buonanno and Y. Chen, "Signal recycled laser-interferometer gravitational-wave detectors as optical springs," *Phys. Rev. D*, vol. 65, no. 4, 042001, p. 042001, Feb. 2002. eprint: gr-qc/0107021.
- [44] V. B. Braginsky, F. Y. Khalili, and K. S. Thorne, *Quantum Measurement*. Cambridge University Press, May 1995.
- [45] V. B. Braginsky, M. L. Gorodetsky, F. Y. Khalili, et al., "Noise in gravitationalwave detectors and other classical-force measurements is not influenced by testmass quantization," *Phys. Rev. D*, vol. 67, no. 8, 082001, p. 082001, Apr. 2003. eprint: gr-qc/0109003.
- [46] P. Leaci and A. Ortolan, "Coupled harmonic oscillators for the measurement of a weak classical force at the standard quantum limit," *Phys. Rev. A*, vol. 76, no. 6, 062101, p. 062101, Dec. 2007.
- [47] F. Y. Khalili, V. I. Lazebny, and S. P. Vyatchanin, "Sub-standard-quantum-limit sensitivity via optical rigidity in the advanced LIGO interferometer with optical losses," *Phys. Rev. D*, vol. 73, no. 6, 062002, p. 062002, Mar. 2006. eprint: grqc/0511008.
- [48] G. Gilbert, A. Aspect, and C. Fabre, Introduction to Quantum Optics. Cambridge University Press, 2010.
- [49] B. W. Barr, "Experimental investigations into advanced configurations and optical techniques for laser interferometric gravitational wave detectors," PhD thesis, University of Glasgow, 2003.

- [50] M. Rakhmanov, "Dynamics of laser interferometric gravitational wave detectors," PhD thesis, California Institute of Technology, 2000.
- [51] R. W. Drever, J. L. Hall, F. V. Kowalski, et al., "Laser phase and frequency stabilization using an optical resonator," Applied Physics B: Lasers and Optics, vol. 31, pp. 97–105, Jun. 1983.
- [52] M. P. Edgar, "Experimental investigations into diffractive optics and optomechanical systems for future gravitational wave detectors," PhD thesis, University of Glasgow, 2011.
- [53] J. Macarthur, "Towards surpassing the standard quantum limit using optical springs," PhD thesis, University of Glasgow, 2014.
- [54] F. Y. Khalili, "Frequency-dependent rigidity in large-scale interferometric gravitationalwave detectors," *Physics Letters A*, vol. 288, pp. 251–256, Oct. 2001. eprint: grqc/0107084.
- [55] V. B. Braginsky and S. P. Vyatchanin, "Low quantum noise tranquilizer for Fabry-Perot interferometer," *Physics Letters A*, vol. 293, pp. 228–234, Feb. 2002.
- [56] T. Corbitt, D. Ottaway, E. Innerhofer, et al., "Measurement of radiation-pressureinduced optomechanical dynamics in a suspended fabry-perot cavity," *Physical Re*view A (Atomic, Molecular, and Optical Physics), vol. 74, no. 2, 021802, p. 021802, 2006.
- [57] C. M. Caves and B. L. Schumaker, "New formalism for two-photon quantum optics.
 I Quadrature phases and squeezed states. II Mathematical foundation and compact notation," *Phys. Rev. A*, vol. 31, pp. 3068–3111, May 1985.
- [58] T. Corbitt, Y. Chen, and N. Mavalvala, "Mathematical framework for simulation of quantum fields in complex interferometers using the two-photon formalism," *Phys. Rev. A*, vol. 72, no. 1, 013818, p. 013818, Jul. 2005. eprint: quant-ph/0502088.
- [59] T. R. Corbitt, "Quantum noise and radiation pressure effects in high power optical interferometers," PhD thesis, Massachusetts Institute of Technology, 2008.
- [60] H. Miao, Exploring Macroscopic Quantum Mechanics in Optomechanical Devices. 2012.
- [61] H. Rehbein, H. Müller-Ebhardt, K. Somiya, et al., "Double optical spring enhancement for gravitational-wave detectors," Phys. Rev. D, vol. 78, no. 6, pp. 062003-+, Sep. 2008. arXiv: 0805.3096.
- [62] V. B. Braginsky and F. Y. Khalili, "Nonlinear meter for the gravitational wave antenna," *Physics Letters A*, vol. 218, pp. 167–174, Feb. 1996.

- [63] H. Rehbein, H. Müller-Ebhardt, K. Somiya, et al., "Local readout enhancement for detuned signal-recycling interferometers," Phys. Rev. D, vol. 76, no. 6, 062002, p. 062 002, Sep. 2007. arXiv: 0705.2987 [gr-qc].
- [64] T. Corbitt, Y. Chen, E. Innerhofer, et al., "An all-optical trap for a gram-scale mirror," Phys. Rev. Lett., vol. 98, no. 15, p. 150 802, Apr. 2007.
- [65] C. M. Mow-Lowry, A. J. Mullavey, S. Goßler, et al., "Cooling of a Gram-Scale Cantilever Flexure to 70mK with a Servo-Modified Optical Spring," *Physical Review Letters*, vol. 100, no. 1, 010801, p. 010801, Jan. 2008.
- [66] F. Elste, S. M. Girvin, and A. A. Clerk, "Quantum Noise Interference and Backaction Cooling in Cavity Nanomechanics," *Physical Review Letters*, vol. 102, no. 20, 207209, p. 207209, May 2009. arXiv: 0903.2242 [cond-mat.mes-hall].
- [67] A. Xuereb, R. Schnabel, and K. Hammerer, "Dissipative Optomechanics in a Michelson-Sagnac Interferometer," *Physical Review Letters*, vol. 107, no. 21, 213604, p. 213604, Nov. 2011. arXiv: 1107.4908 [physics.optics].
- [68] S. P. Tarabrin, H. Kaufer, F. Y. Khalili, et al., "Anomalous dynamic backaction in interferometers," Phys. Rev. A, vol. 88, no. 2, 023809, p. 023809, Aug. 2013. arXiv: 1212.6242 [quant-ph].
- [69] N. Vostrosablin and S. P. Vyatchanin, "Stable optical spring in the Advanced LIGO detector with unbalanced arms and in the Michelson-Sagnac interferometer," *Phys. Rev. D*, vol. 89, no. 6, 062005, p. 062005, Mar. 2014. arXiv: 1312.7153 [quant-ph].
- [70] A. Sawadsky, H. Kaufer, R. M. Nia, et al., "Observation of Generalized Optomechanical Coupling and Cooling on Cavity Resonance," *Physical Review Letters*, vol. 114, no. 4, 043601, p. 043601, Jan. 2015. arXiv: 1409.3398 [quant-ph].
- B. J. Meers, "Recycling in laser-interferometric gravitational-wave detectors," *Phys. Rev. D*, vol. 38, pp. 2317–2326, Oct. 1988.
- [72] K. A. Strain and B. J. Meers, "Experimental demonstration of dual recycling for interferometric gravitational-wave detectors," *Physical Review Letters*, vol. 66, pp. 1391–1394, Mar. 1991.
- [73] S. Hild, H. Grote, M. Hewtison, et al., "Demonstration and comparison of tuned and detuned signal recycling in a large-scale gravitational wave detector," *Classical* and Quantum Gravity, vol. 24, pp. 1513–1523, Mar. 2007.
- [74] K. J. Vahala, "Optical microcavities," Nature, vol. 424, pp. 839–846, Aug. 2003.

- [75] M. Hossein-Zadeh and K. J. Vahala, "Observation of optical spring effect in a microtoroidal optomechanical resonator," *Optics Letters*, vol. 32, pp. 1611–1613, Jun. 2007.
- [76] J. Hofer, A. Schliesser, and T. J. Kippenberg, "Cavity optomechanics with ultrahigh-Q crystalline microresonators," *Phys*, vol. 82, no. 3, 031804, p. 031804, Sep. 2010. arXiv: 0911.1178 [physics.optics].
- [77] L. Ding, C. Baker, P. Senellart, et al., "Wavelength-sized GaAs optomechanical resonators with gigahertz frequency," Applied Physics Letters, vol. 98, no. 11, 113108, p. 113108, Mar. 2011. arXiv: 1101.4499 [physics.optics].
- [78] K. Agatsuma, D. Friedrich, S. Ballmer, et al., "Precise measurement of laser power using an optomechanical system," Optics Express, vol. 22, p. 2013, Jan. 2014. arXiv: 1309.4779 [physics.ins-det].
- [79] A. Mizrahi and L. Schächter, "Two-slab all-optical spring," Optics Letters, vol. 32, pp. 692–694, Mar. 2007.
- [80] T. Westphal, G. Bergmann, A. Bertolini, *et al.*, "The 10m AEI prototype facility A brief overview," *ArXiv e-prints*, Nov. 2011. arXiv: 1111.7252 [physics.optics].
- [81] M. Evans, "Optickle," Technical Report T07026-00, LIGO Laboratory, Tech. Rep., 2007.
- [82] MathWorks Inc., MATLAB release R2014b. Natick, Massachusetts: MathWorks Inc., 2014.
- [83] A. E. Seigman, *Lasers*. Oxford University Press, 1986.
- [84] N. Hodgson and H. Weber, Laser Resonators and Beam Propagation: Fundamentals, Advanced Concepts and Applications, ser. Springer Series in Optical Sciences. New York, NY: Springer, 2005.
- [85] C. Torrie, "Development of suspensions for the Geo 600 Gravitational Wave Detector," PhD thesis, University of Glasgow, 2001.
- [86] Wolfram Research, Inc., Mathematica Version 8.0. Wolfram Research, Inc., 2010.
- [87] http://www.arduino.cc/.
- [88] G. Cagnoli, L. Gammaitoni, J. Kovalik, et al., "Eddy current damping of high Q pendulums in gravitational wave detection experiments," *Review of Scientific Instruments*, vol. 69, pp. 2777–2780, Jul. 1998.
- [89] W. Cunningham, "Analysis of eddy current damping for AEI 10m prototype," University of Glasgow, Tech. Rep., 2012.

- [90] H. Kogelnik and T. Li, "Laser beams and resonators," Applied Optics, vol. 5, p. 1550, Oct. 1966.
- [91] http://www.sr.bham.ac.uk/dokuwiki/doku.php?id=geosim:jammt.
- [92] M. J. Lawrence, B. Willke, M. E. Husman, et al., "Dynamic response of a FabryPerot interferometer," Journal of the Optical Society of America B Optical Physics, vol. 16, pp. 523–532, Apr. 1999.
- [93] H. Rohde, J. Eschner, F. Schmidt-Kaler, and R. Blatt, "Optical decay from a FabryPerot cavity faster than the decay time," *Journal of the Optical Society of America B Optical Physics*, vol. 19, pp. 1425–1429, Jun. 2002. eprint: physics/ 0107038.
- [94] M. Rakhmanov, R. Savage, D. Reitze, and D. Tanner, "Dynamic resonance of light in fabry-perot cavities," *Physics Letters A*, vol. 305, no. 5, pp. 239–244, 2002.
- [95] J. Bechhoefer, "Feedback for physicists: a tutorial essay on control," Rev. Mod. Phys., vol. 77, pp. 783–836, 3 Aug. 2005.
- [96] K. Ogata, Modern control engineering. 1970.
- [97] H. Nyquist, "Regeneration theory," *Bell System Technical Journal*, vol. 13, p. 200, 1932.
- [98] J. R. Taylor, "Interferometric Experiments Towards Advanced Gravitational Wave Detectors," PhD thesis, University of Glasgow, 2008.
- [99] R. Bork, "AdvLigo CDS Design Overview," LIGO Scientific Collaboration, Tech. Rep., 2010.
- [100] E. D. Black, "An introduction to Pound-Drever-Hall laser frequency stabilization," American Journal of Physics, vol. 69, pp. 79–87, Jan. 2001.
- [101] G. Heinzel, LISO http://www.mpg.de/ros/geo600 docu/soft/liso/.
- [102] M. Herpy and J. C. Berka, Active RC Filter Design, J. I. Sewell, Ed. Elsevier, 1986.
- J. Degallaix, H. Grote, M. Prijatelj, et al., "Commissioning of the tuned DC readout at GEO 600," Journal of Physics Conference Series, vol. 228, no. 1, 012013, p. 012013, May 2010.
- [104] A. Freise, M. M. Casey, S. Gossler, et al., "Performance of a 1200 m long suspended Fabry-Perot cavity," *Classical and Quantum Gravity*, vol. 19, pp. 1389–1397, Apr. 2002. eprint: gr-qc/0110020.
- [105] E. Morrison, D. I. Robertson, H. Ward, and B. J. Meers, "Automatic alignment of optical interferometers," Appl. Opt., vol. 33, pp. 5041–5049, Aug. 1994.

- [106] J. A. Sidles and D. Sigg, "Optical torques in suspended Fabry Perot interferometers," *Physics Letters A*, vol. 354, pp. 167–172, May 2006.
- [107] E. Hirose, K. Kawabe, D. Sigg, et al., "Angular instability due to radiation pressure in the LIGO gravitational-wave detector," Applied Optics, vol. 49, p. 3474, Jun. 2010. arXiv: 0909.0010 [gr-qc].
- [108] G. Heinzel, "Advanced optical techniques for laser-interferometric gravitationalwave detectors," PhD thesis, University of Hannover, 1999.
- [109] A. Rubinowicz, "On the Anomalous Propagation of Phase in the Focus," *Physical Review*, vol. 54, pp. 931–936, Dec. 1938.
- [110] S. Feng and H. G. Winful, "Physical origin of the Gouy phase shift," Optics Letters, vol. 26, pp. 485–487, Apr. 2001.