

When Light Gets Pushy: Radiation Pressure Effects in Interferometric Gravitational Wave Detectors

Anna Catriona Green

A thesis submitted to the University of Birmingham for the degree of
DOCTOR OF PHILOSOPHY

Institute of Gravitational Wave Astronomy
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
May 2018

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

Advanced LIGO made the first direct observation of a gravitational wave in 2015. This signal, the largest measured so far, had a peak strain amplitude of 1×10^{-21} and frequency range 35-250 Hz.

LIGO's sensitivity was achieved after decades of development. Suspended optics made from low-loss materials are employed in an optical configuration with multiple coupled cavities to increase the circulating power and shape the frequency response of the detector. This design results in radiation pressure that substantially alters the detector's behaviour.

One challenging consequence is parametric instabilities, an unstable coupling between the optical field and mechanical resonances in the mirrors. My detailed investigations of parametric instabilities in both LIGO and ET, a planned next-generation detector, show that the severity of the instabilities depends on the complete optical configuration. With this model, an optimal operating point for LIGO can be determined and ET design choices can be weighed against potential instabilities.

An 'optomechanical filter' has been proposed that uses radiation pressure to enhance the detector bandwidth. I test the analytical models against numerical simulations and outline an experimental research programme that will implement a trial filter.

Statement of Originality

This thesis reports the research I conducted during my PhD at the University of Birmingham between September 2014 and March 2018.

The work reported in chapter 3 was led by myself with assistance from Daniel Brown and Miguel Dovale-Álvarez. It has been published in

AC Green, DD Brown, M Dovale-Álvarez, et. al. ‘The influence of dual-recycling on parametric instabilities at Advanced LIGO’. Classical and Quantum Gravity, 34(20):205004, 2017 [1].

Acknowledgements

Firstly, a massive thank you goes to my supervisor, Andreas, for introducing me to gravitational waves and helping me see my potential. Thank you for all your guidance and encouragement, for giving me the opportunity to travel to both conferences and to work as a visitor at Hanford, and for teaching me what it means to be a researcher in our field. Over the course of my PhD the group has grown and changed a lot, but you've always managed to keep it dynamic and fun. I wish you and the group all the best for the future.

Thank you to Conor, my secondary supervisor, for all of your help and support, as well as many an engaging discussion at work, Pub Club, and beyond. Thanks also to Miguel, for all of your COMSOL-ing efforts, and of course a big thanks goes to Dan Brown, without whom much of my work wouldn't have been possible. Thank you both for answering all my questions and requests! To Chris, Artemiy, Haixing, and Denis, thank you for all your help in navigating the many aspects of the unstable filter project. It's been a very exciting period, quite the rollercoaster(!), and I look forward to seeing what happens next. Thank you also to my examiners, Kate and Vincent.

A huge thanks goes to all my colleagues here in Birmingham. Special thanks of course go to all our ifolab crew for your input on my projects, for letting me get involved in yours, and for your friendship throughout: Dan, Haoyu, Daniel, Miguel, Sam, Aaron, Phil, Sam, Joe, Chiara, Chris, Artemiy, Kasuhiro, Andreas, Conor, Haixing and Denis. To all my friends in the ASR group and beyond, there's too many to name - thank you for so many conversations, office distractions, and all our celebrations, trips to the pub, and travels together. Thanks to Alberto for leading the group and supporting all of us. Extra thanks to Hannah, Sam, and all who helped with the Thinktank/Royal Society project. I think we succeeded making the shiniest Michelson, and far more. You guys made it a really great experience.

Thank you to Maryanne and Katie, for your unwavering friendship and for much needed escapes

from the physics. You're both doing fantastically and I can't wait for the next adventure with each of you. To all in the Focolare Movement, you are my family. Thank you for always supporting me, and for sharing in my enthusiasm. Thanks so much to Anja, Eleanor, and all the Gen, you are amazing and can always recharge me. Thanks for being my personal cheerleaders!

Finally thank you to my family. I wouldn't have made it this far without you. To Mum and Dad for always being there for me, helping me find my own path, and for giving me the tools to get where I am today. Thanks also to my wider family, my aunts, uncles and cousins, for your support and a healthy dose of reality— you help keep me grounded. To those who've gone on already—my grandparents, Aunty Jane and Uncle John—this is for you too, thank you for everything you are.

Contents

1	Introduction: Radiation Pressure In Gravitational Wave Detectors	1
1.1	Thesis Overview	7
2	Optomechanical Systems	8
2.1	Notation Conventions and Definitions	9
2.1.1	The optical field	9
2.1.2	Propagation in free space	10
2.1.3	Reflection and Transmission	11
2.2	Radiation Pressure	11
2.3	Two-Mirror Cavity	12
2.4	Optical Springs	15
2.4.1	The Adiabatic Approximation	16
2.4.2	The Frequency-dependent Optical Rigidity	18
2.4.3	Overall behaviour of the optomechanical system	23
2.5	Gaussian Beams and Higher Order Optical Modes	24
2.6	Parametric Instability	27
2.6.1	The Feedback Picture and Parametric Gain	28
2.7	Modelling Radiation Pressure Effects Using FINESSE	31
2.7.1	Optomechanical Couplings	32
2.7.2	Finite Element Modelling of Mechanical Modes	35
2.7.3	Parametric Gain	36
2.7.4	Testing FINESSE	37
3	Modelling Dual-Recycled Parametric Instabilities at Advanced LIGO	41
3.1	Observation and Mitigation of Parametric Instabilities at LIGO	42
3.2	Modelling LIGO Parametric Instabilities	45
3.2.1	Design of the Core Optics in Advanced LIGO	45
3.2.2	Mechanical Modes of an Advanced LIGO Test Mass	46
3.2.3	General Method	48
3.3	Parametric Instability in Increasingly Complex Interferometers	48
3.3.1	‘Common’ and ‘Differential’ Parametric Instabilities	48
3.3.2	A ‘Forest of Modes’	51
3.4	The Signal Recycling Cavity	52
3.4.1	Tuning	52
3.4.2	Gouy Phase	54

3.4.3	Consequences for Advanced LIGO Interferometers	55
3.5	Further Work	57
3.5.1	Is There a Gouy Phase ‘Sweet Spot’?	57
3.5.2	ITM vs. ETM	58
3.5.3	Modelling PIs using parameters from the Livingston detector	60
3.6	Conclusions	61
3.6.1	Suggested areas for further follow-up	62
4	Parametric Instabilities in the Einstein Telescope	63
4.1	Evolution of the Einstein Telescope Design	65
4.1.1	2011 Design Study: The Xylophone Configuration	65
4.1.2	The ET-120K Design	66
4.2	Finite Element Modelling of ET Test Masses	71
4.3	Speeding up PI Modelling for ET using Clipping Effects	74
4.3.1	Apertures as a Source of Loss	74
4.3.2	Clipping of Higher Order Modes	76
4.3.3	Using Clipping Effects to Improve and Speed Up PI Modelling	78
4.4	PI ‘Forests’ in ET-120K arm cavities	79
4.5	Comparison between Advanced LIGO and ET ‘Forests’	83
4.6	PI in a ‘kHz-optimised’ ET-120K	86
4.7	Future work: ‘PI-Aware’ Design for ET	90
4.8	Conclusions	92
5	Bandwidth Manipulation using an Optomechanical Filter	94
5.1	Broadening the Bandwidth of Gravitational Wave Detectors using Negative Dispersion	94
5.2	Idealised Negative Dispersion and Bandwidth Broadening	96
5.3	Generating Negative Dispersion using Optomechanical Systems	98
5.3.1	The Input-Output Relation for a Radiation Pressure Limited Cavity	98
5.3.2	Radiation Pressure Limited Cavity as an Active Filter	103
5.4	Conclusions	107
6	Development of a Tabletop Experiment Towards Optomechanical Bandwidth Manipulation	108
6.1	Thermal Noise	111
6.2	Controlling unstable systems	112
6.3	Frequency Requirements for a Tabletop Optomechanical Filter	113
6.4	Design of the Suspended Mirror	115
6.5	Geometric Design of the Filter Cavity	117
6.6	Optical Layouts	119
6.6.1	Fibre Cavity System	120
6.6.2	Triple Cavity System	122
6.7	Current Experimental Status: Chip Characterisation	125
6.8	Future Work	128
6.9	Conclusions	129
7	Summary and Outlook	131

A	Advanced LIGO Model	134
A.1	Test Mass Geometry and Material Properties	135
A.2	FINESSE Model	135
B	Modelling Complex Interferometers in Finesse	142
B.1	Overview	142
B.2	Example: ET-120K	145
C	Input-Output Relation for a Radiation Pressure Limited Cavity	153
C.1	Linearised Time Domain Equations of Motion	153
C.2	Frequency Domain Input-Output Relation	156
C.3	Simplified form	157
D	Sign conventions for ‘Theorists’ and ‘Experimentalists’	158
E	Controlling MIMO systems: State-Space Approach	160
E.1	Describing Dynamical Systems	160
E.1.1	Frequency Domain	161
E.1.2	State-Space Approach	162
E.2	Controlling Dynamical Systems	163
E.2.1	describing systems with feedback	163
E.2.2	Stability	164
E.2.3	Designing simple controllers	166
	Bibliography	168

List of Figures

1.1	Target sensitivity of the Advanced LIGO design, and contributing noise sources, reproduced from [2].	2
1.2	Schematic optical layout of a current gravitational wave detector	3
2.1	Steady state fields in a two-mirror cavity	12
2.2	Power circulating in a two-mirror cavity	14
2.3	Two-mirror cavity with suspended end mirror	15
2.4	Optical rigidity vs detuning in the adiabatic approximation	17
2.5	Optical rigidity vs detuning using a perturbative model	22
2.6	Spatial distribution of the lowest order Hermite-Gauss modes	25
2.7	Propagation of a Gaussian beam in free space	25
2.8	Feynman diagrams of Stokes and Anti-Stokes modes	27
2.9	A two mirror cavity in which the end mirror is mechanically excited.	28
2.10	Parametric instability depicted as a feedback process	29
2.11	Optical layout modelled in FINESSE, replicating the setup in Corbitt. et. al., 2007	37
2.12	Optical spring test: FINESSE vs. Corbitt et. al., 2007	39
2.13	Parametric instability test: FINESSE vs. Corbitt et. al., 2006	40
3.1	Mechanical modes observed during the first Advanced LIGO observation run.	43
3.2	Optical Layout of a Dual-Recycled Michelson with Fabry-Perot Arms.	45
3.3	Geometry of an Advanced LIGO Test Mass as modelled using COMSOL.	47
3.4	Surface motion maps for mechanical modes used throughout this chapter	47
3.5	Parametric gain of mode 37 with increasing detector complexity	49
3.6	Parametric gain of modes 37, 41 and 257 for changing radius of curvature	50
3.7	A forest of PIs	51
3.8	Parametric gain of modes 37, 41 and 257 for changing SRC tuning	53
3.9	Parametric gain of mode 37 for changing SRC tuning	53
3.10	Parametric gain of mode 37 for changing SRC Gouy phase	54
3.11	Total number and summed gain of PIs for different SRC tuning and Gouy phase	55
3.12	Total number of modes that are unstable within $\pm 2\text{kHz}$ of their COMSOL frequency.	56
3.13	Total number and summed gain of PIs vs combined PRC/SRC Gouy phase	57
3.14	Parametric gain of mode 37 for Input and End Test Masses in both arms	59
3.15	Parametric gain of mode 37 at LLO vs. Advanced LIGO	60
4.1	Proposed optical layout for the Einstein Telescope from the 2011 Design Study	65

LIST OF FIGURES

4.2	Sensitivity curves for ET-HF and ET-LF	67
4.3	Quantum-noise-limited sensitivity curves for a possible ‘ET-120K’ design	68
4.4	Optical layout of the ET-120K interferometer design	69
4.5	Geometries of test masses proposed for ET-HF and ET-120K	72
4.6	Example surface motion maps for ET-HF and ET-120K test masses	73
4.7	Power reflected and lost when a Gaussian beam reflects from a flat, circular mirror	75
4.8	Numerical computation of clipping loss	76
4.9	Transverse intensity distribution of higher order modes with mode order	77
4.10	Loss due to clipping increases with mode order.	77
4.11	Number of PIs and computation time vs number of HOMs included in the model .	79
4.12	Quantum-limited sensitivity curves for ET-120K with clipping	81
4.13	PI ‘forests’ in arm cavities for two test mass radius of curvature choices in ET-120K.	83
4.14	Number of PIs vs number of HOMs modelled for ET-120K	84
4.15	‘Forests’ of PIs in two dual-recycled designs for ET-120K	85
4.16	Quantum-noise-limited sensitivity curve for a kHz-optimised ET-120K	88
4.17	The ‘PI forest’ for the ET-120K kHz-optimised design	89
5.1	Optical layout used to model negative dispersion and bandwidth broadening	97
5.2	Bandwidth manipulation using a ‘magic’ filter cavity	97
5.3	Optical layout used to derive the Input-Output relation	99
5.4	Comparison between direct analytical equations and FINESSE	101
5.5	Lower-to-lower sideband relation for a cavity with moveable end-mirror	103
5.6	The optomechanical filter concept	104
5.7	Key frequency definitions used in the optomechanical filter	104
5.8	$T_1(-\omega)$ at various degrees of approximation.	106
5.9	Bandwidth broadening of a LIGO-like arm cavity using an optomechanical filter. .	106
6.1	Original proposal for implementing an optomechanical filter in a LIGO-like detector	109
6.2	Simplified depiction of a gravitational wave detector with an optomechanical filter	110
6.3	Proposed bandwidth broadening using an optomechanical filter	110
6.4	Example micro-oscillator chip and mounting	116
6.5	4-legged mirror design	116
6.6	Geometry of the Filter Cavity	117
6.7	Beam waist size vs cavity length for various input coupler radii of curvature	118
6.8	Schematic layout of a fibre-based cavity system to demonstrate bandwidth broadening.	120
6.9	Bandwidth of fibre cavities compared to a free space cavity of equivalent length . .	122
6.10	Triple-cavity free-space optical system	123
6.11	Phase accumulation in 2-mirror versus 3-mirror cavity systems.	123
6.12	Optical layout used to characterise mirrors on the AlGaAs chip	126
6.13	Core Optics of the in-air Michelson setup	127
6.14	In-air ‘ringdown measurement’ of micro-mirror ‘A’	128
B.1	Optical layout of ET-120K as modelled in FINESSE for parametric instability studies	146
B.2	The composite beamsplitter	148
B.3	Power in the arms and transmitted from the Michelson vs. CARM and DARM . .	149
B.4	Power in HG modes at the output of the interferometer vs. DARM	151

LIST OF FIGURES

B.5	Quantum-noise-limited sensitivity curve of the ET-120K detector design	152
C.1	Optical layout used to derive the input-output relation	153
E.1	Block diagram of a controlled system	164

List of Tables

2.1	Approximations and definitions used to simplify optical equations	15
2.2	Parameters used to replicate Corbitt <i>et al</i> 's 2007 optical trap experiment	38
2.3	Parameters used to replicate Corbitt <i>et al</i> 's 2006 suspended cavity experiment. . .	40
3.1	Key frequencies derived from the Advanced LIGO design model.	46
4.1	Comparison of optical parameters for ET designs	70
4.2	Mechanical properties of mirror substrates for two ET designs vs. Advanced LIGO	71
4.3	Summary of FEM results, comparing ET-120K and ET-HF mirror designs to LIGO.	74
4.4	Key frequencies and related parameters for the cavities in two ET-120K designs . .	85
4.5	Key frequencies and related parameters for the 'kHz-optimised' ET-120K design .	88
5.1	Parameters used to compare the analytical and numerical input-output relation . .	102
6.1	Mechanical resonant frequencies resulting from various potential mirror geometries.	117
6.2	Geometric and optical parameters of the proposed Filter Cavity.	119
6.3	Parameters proposed for a triple-cavity free-space bandwidth broadening experiment	125
A.1	Parameters used to generate a COMSOL model of a LIGO test mass	135
D.1	Sign conventions used by theorists and experimentalists	159

Glossary

AMD	Acoustic Mode Damper
BBH	Binary Black Hole
BNS	Binary Neutron Star
BS	Beamsplitter
CAD	Computer Aided Design
DRFPMi	Dual-Recycled Michelson interferometer with Fabry-Perot arm cavities
ESD	Electrostatic Driver
ET	Einstein Telescope
ET-120K	120 K cooled design for ET, also sometimes called ‘ET-Voyager’
ET-HF	High Frequency interferometer design for ET
ET-LF	Low Frequency interferometer design for ET
ETM(X/Y)	End Test Mass (in the X- or Y-arm)
FC	Filter Cavity
FDS	Frequency-Dependent Squeezing
FE	Finite Element
FEM	Finite Element Modelling
FINESSE	Frequency domain INterfErometer Simulation SoftwarE
FPMi	Michelson interferometer with Fabry-Perot arm cavities
FSR	Free Spectral Range
HG	Hermite-Gauss
HOM	Higher Order Mode
IMC	Input Mode Cleaner
ITM(X/Y)	Input Test Mass (in the X- or Y-arm)
LG	Laguerre-Gauss
LHO	LIGO Hanford Observatory
LIGO	Laser Interferometer Gravitational wave Observatory
LLO	LIGO Livingston Observatory

LIST OF TABLES

O1/2	Observing run 1/2
OMC	Output Mode Cleaner
PI	Parametric Instability
PR2/3	Power Recycling cavity mirror 2/3
PRC	Power Recycling Cavity
PRFPMi	Power-Recycled Michelson interferometer with Fabry-Perot arm cavities
PRM	Power Recycling Mirror
QN	Quantum Noise
RC	Recycling Cavity
RF	Radio Frequency
RoC	Radius of Curvature
RSE	Resonant Sideband Extraction
SQL	Standard Quantum Limit
SR2/3	Signal Recycling cavity mirror 2/3
SRC	Signal Recycling Cavity
SRM	Signal Recycling Mirror
TM	Test Mass

Finesse syntax

ad	amplitude detector –outputs the complex amplitude of an optical field component at the specified node.
attr	set attribute of component –used to attribute additional properties to optics, such as a mass or curvature.
fsig	signal input –defines an input signal for calculating transfer functions.
Fs0	used in the fsig command –specify a signal input produced due to a force distorting the surface of an optic.
Fz	used in the fsig command –specify a signal input produced due to a force on an optic in the beam axis.
maxtem	maximum mode order –sets the maximum order of higher order optical modes to include in the computation.
smotion	surface motion –defines a generalised surface motion and applies it to an optic.
tf	transfer function –defines a transfer function with respect to frequency, using poles and zeros.
xd	motion detector –outputs displacement in the specified degree of freedom.
zmech	used in the attr command –specify a longitudinal suspension transfer function (i.e. in the beam axis).

Chapter 1

Introduction: Radiation Pressure In Gravitational Wave Detectors

On the 14th September 2015, the Laser Interferometer Gravitational Wave Observatory (LIGO) [2] made the first direct detection of a gravitational wave signal [3]. This marks the start of a new era for astronomy, and is the culmination of over 40 years of experimental effort. The signal, from the coalescence of a binary black hole (BBH) system a billion light-years away, had peak strain amplitude of 1×10^{-21} and increased in amplitude and frequency from 35 Hz to 250 Hz over a period of 0.2 s. Since then, LIGO has successfully detected several such BBH systems [4, 5, 6]. In 2017, Virgo [7] joined the gravitational wave detector network, significantly improving our ability to localise gravitational wave events, as demonstrated by the first three-detector observation of a BBH system on August 14th [8]. The first ‘multi-messenger’ gravitational wave detection was made a short time later on August 17th [9], marking another key milestone for gravitational wave astronomy. This signal, from a coalescing neutron star system, was accompanied by electromagnetic signals across the whole electromagnetic spectrum, and was a landmark moment involving the collaboration of a significant fraction of the global astronomical community, with coordination between gravitational wave, electromagnetic, and particle observatories across the planet.

None of this would have been possible without decades of experimental development. When initially predicted by Albert Einstein in 1916, gravitational waves—a ‘stretching and squashing’ of space-time resulting from the motion of accelerated masses—were believed to be such a small effect that attempting to detect them was considered a practical impossibility. As technologies have

developed this ‘impossible’ task began to look less so, if still extremely difficult; today the results of the gravitational wave community’s work speak for themselves.

Many of the particular challenges of reaching the required sensitivity for gravitational wave detection can be summarised by figure 1.1, which depicts the target sensitivity curve for the Advanced LIGO detectors and the effects of the many contributing sources of noise.




Figure 1.1: Target sensitivity of the Advanced LIGO design, and contributing noise sources, reproduced from [2].

At low frequencies (<10 Hz), the detectors are limited by seismic noise. This includes all kinds of ground motion, from sources including earthquakes, waves crashing on the coast, and human activities such as logging or transportation. Both active and passive techniques are used to minimise the coupling of ground motion into the detector, however these can also introduce noise. The test masses of the detector are suspended to provide passive isolation, but the suspension fibres have associated thermal noise as well as resonant modes. Careful material selection and fibre design ensure that these do not limit the operation of the detector.

In the middle of the detection band, we are limited by coating thermal noise, associated primarily with the Brownian motion of the atoms in the optical coatings applied to the test masses. This is dependent on the material properties and temperature of the coating, the latter of which can be addressed using cryogenic cooling as well as by changing the intensity distribution of the laser beam incident on the coating.

At high frequencies, we are limited only by quantum noise. This is the fundamentally limiting noise of the detector, resulting from quantum fluctuations in the number of photons interacting with the detector optics. At high frequencies, this takes the form of *shot noise*—statistical fluctuations in the rate of arrival of photons, resulting in photon counting noise on readout of the gravitational wave signal. These photon number fluctuations also translate into a varying radiation pressure force on the test masses—*quantum radiation pressure noise*—which becomes the limiting quantum process at lower frequencies. While the signal to noise ratio (SNR) of shot noise improves with higher optical power, the quantum radiation pressure noise SNR worsens. At low frequencies other noises dominate the interferometer, and using high-mass test masses can reduce quantum radiation pressure noise, so a high power is chosen for current detectors: to achieve design sensitivity, LIGO will require 800 kW circulating optical power. The next generation of detectors are expected to significantly increase this.

Optical Configuration of Current Gravitational Wave Detectors




Figure 1.2: Schematic optical layout of a current gravitational wave detector. The configuration is based on a Michelson interferometer with Fabry-Perot arm cavities, formed from the central beamsplitter (BS) and the input- and end test masses (I- and ETMs) in the X- and Y-arms. The power- and signal-recycling mirrors (P- and SRM) form additional cavities which are used to increase the circulating power and shape the frequency response of the interferometer, resulting in the overall dual-recycled configuration.

The overall shape of the quantum-limited sensitivity curve is determined by the optical configuration of the detector. Current gravitational wave detectors are ground-based, kilometre-scale optical systems based on a Michelson interferometer, as depicted in figure 1.2. Gravitational waves exert a strain on space-time, which can be measured by comparing the relative distance between

free masses as the gravitational wave passes. Using a Michelson allows noises that are common to both ‘arms’ of the interferometer, such as laser frequency noise, to be separated from differential signals, including the gravitational wave signal. The interferometer is tuned to a dark fringe on transmission—we perform a null measurement. Fabry-Perot cavities are included in the arms to increase their effective length, amplifying the signal. The mirrors forming these cavities are referred to as the Input- and End Test Masses (I- and ETMs). The shot-noise-limited sensitivity is inversely related to the optical power in the arm cavities. We therefore include a Power Recycling Mirror (PRM), forming the Power Recycling Cavity (PRC) between the PRM and ITMs, to reflect rejected light back into the interferometer and increase the circulating power. Finally, the detector sensitivity is dependent on the storage time of the gravitational wave signal in the detector. This is controlled through the introduction of the Signal Recycling Cavity (SRC), formed by the Signal Recycling Mirror (SRM) and ITMs. The tuning of this cavity can be chosen to amplify or resonantly extract signal sidebands. The latter technique, referred to as resonant sideband extraction (RSE), is used currently by the LIGO detectors, and results in the broad, bucket-shaped sensitivity curve depicted in figure 1.1. The combined use of power- and signal recycling is commonly referred to as Dual Recycling.

Parametric Instabilities

The LIGO detectors are yet to reach the sensitivity targeted by the Advanced LIGO design. After several years commissioning work to implement dual recycling, the detectors are now operated as observatories for extended periods, interspersed with planned commissioning breaks to increase the detector sensitivity and observation time in stages. While parts of this work have been well-anticipated tasks, designing and building a gravitational wave detector is by necessity at the threshold of many aspects of new science, and therefore comes with less-anticipated challenges.

One particular challenge, and the central theme of my thesis, comes as a direct result of the design choices required to minimise seismic and shot noise. By suspending the test masses and using high optical power, radiation pressure becomes a significant effect in gravitational wave detectors. This changes the dynamical behaviour of the interferometer, which has consequences for both the overall sensitivity of the detector and our ability to control it.

In late 2014, a *parametric instability* (‘PI’) was observed for the first time in a gravitational wave detector [10]. These instabilities result from a coupling between the optical field and the internal mechanical modes of the test masses. In cases where the optical power is high, and the resonant

conditions for the optical fields in the interferometer are met, this coupling can result in the optical field driving the motion of the mechanical mode in a positive feedback loop, causing the mechanical mode to ‘ring up’. This results in exponentially growing optical signals that can saturate the sensors used to control the detector, meaning that it is no longer maintained at a useful operating point—a process referred to as ‘losing lock’. As a result, the observation time of the detector is cut, with lock re-acquisition typically taking at least half an hour.

Parametric instabilities were first predicted in 2001 [11], and anticipated to be a phenomenon that could limit the performance of high-powered interferometric gravitational wave detectors. However, it was unclear how significant their effect on LIGO, and detectors like it, would be. Following the first observation at the LIGO Livingston Observatory (LLO), PIs have become a regular occurrence at both LIGO sites, particularly at Hanford (LHO), which chose to upgrade its input power from ~ 25 W to ~ 50 W during the commissioning period between the first and second observing runs in 2016.

Understanding and mitigation of PIs has therefore been a very active area of research in recent years. A variety of mitigation techniques have been explored and are now being used to successfully control PIs at both sites, but work is ongoing to ensure LIGO will remain resilient to PIs as we further increase the optical power.

Future Detectors

While the current generation of gravitational wave detectors must mitigate parametric instabilities retro-actively, the next generation can be designed with parametric instabilities in mind. The Einstein Telescope is a current European proposal for a ‘third generation’ gravitational wave detector, intended to have a broader bandwidth and peak sensitivity an order of magnitude better than current detectors. The facility, which will be underground and use arm lengths of order 10 km, is planned to have a multi-decade lifespan, and it is expected that the detector will evolve significantly over this period.

The interferometer design for ET is still in development. In 2011 an extensive design study was produced [12], proposing a triangular detector layout consisting of three pairs of Michelson interferometers, each pair containing an interferometer optimised independently for low-frequency and high-frequency gravitational wave signals, called ET-LF and ET-HF. While ET-LF would use cryogenic cooling to minimise thermal noise, ET-HF would use high circulating power to re-

duce shot noise. At a workshop in late 2017, it was suggested that the ET site could first use a single-interferometer configuration I refer to as ‘ET-120K’, which uses reduced levels of cryogenic cooling with high optical power, achieving a sensitivity close to that achieved using the ET-HF/-LF ‘xylophone’ configuration.

Parametric instabilities are an example of radiation pressure resulting in an unwanted behaviour that must be minimised in order for gravitational wave detectors to operate as intended. The latter portion of my thesis has instead explored a new concept which would use couplings between the optical field and a mechanical oscillator to our advantage, creating an unstable ‘optomechanical filter’ that could be used to broaden the bandwidth of an interferometric gravitational wave detector without compromising the detector’s peak sensitivity.

The optomechanical filter, proposed in 2015 by H. Miao *et. al.* [13], is an optical cavity in which a mirror is suspended and the cavity dynamics are dominated by radiation pressure. The mechanical resonant frequency of the suspended mirror, input (pump) laser frequency, and cavity’s resonant frequency and bandwidth are carefully selected such that it acts as a source of ‘negative dispersion’: optical fields reflected from it lose, rather than gain, phase. This anomalous behaviour means that the optomechanical filter, when incorporated into an optical system such as a gravitational wave detector, is a type of ‘white light cavity’. The negative dispersion of the filter can be used to counteract the phase accumulation of gravitational wave signals in the arm cavities, meaning that the bandwidth of the detector as a whole is increased.

From concept to implementation, any new technology for improving gravitational wave detector sensitivity typically takes several decades, since it must comply with the detector’s strict requirements regarding noise and reliable operation. Before an idea is tested in a large-scale interferometer, we first determine if it can be practically realised using a combination of numerical modelling and table-top experiments. I have started this process for the optomechanical filter.

1.1 Thesis Overview

In this thesis I present my work studying radiation pressure effects in interferometric gravitational wave detectors. This falls into two major categories: parametric instabilities—a problem for the current generation of detectors, and the optomechanical filter—a proposal to use radiation pressure to improve future detectors.

Chapter 2 provides an overview of key optical and optomechanical concepts that will be used throughout my thesis, and describes how optomechanical couplings are numerically modelled using FINESSE. In chapter 3 I show that parametric instabilities at Advanced LIGO are influenced by the whole core optical configuration, as published in [1]. Chapter 4 demonstrates how design choices for the Einstein Telescope will affect how severely it experiences PIs, and proposes a plan for developing the Einstein Telescope design with parametric instabilities in mind. In chapter 5 I review the concept of the optomechanical filter and compare the results of a numerical model to an analytical description of the filter. Implementation of such a filter is challenging. Chapter 6 first notes the particular challenges of the thermal requirements and controllability of the filter concept, then describes a research programme aimed at addressing the controllability of the filter.

Chapter 2

Optomechanical Systems

In optical systems with high circulating power and suspended optics, the radiation pressure force exerted by the optical field on the optics can contribute significantly to the behaviour of the system. This means that gravitational wave detectors, which incorporate both of these features to reach their sensitivity goals, are susceptible to radiation pressure effects such as *optical springs* and *parametric instabilities* (PIs). While these features can result in additional challenges for the control of the detector, it may also be possible to use such radiation pressure-dependent behaviours to enhance the performance of detectors in the future.

In this chapter I provide a theoretical overview of key optical and optomechanical concepts that will be used throughout my thesis. In section 2.1 I note the sign and definitional conventions used, and then in sections 2.2 to 2.4 I build from describing a simple two-mirror cavity to a description of an optical spring. This derivation is based on a frequency-domain, perturbative description of the optical fields in the sideband picture using plane waves. In section 2.5 I review Gaussian beams and the modal description of optical fields, and in section 2.6 develop the framework and metric used to describe parametric instabilities. Finally in section 2.7 I outline how optomechanical couplings are numerically modelled using FINESSE, in conjunction with finite element modelling (FEM) tools such as COMSOL Multiphysics, and present examples of tests used to verify the numerical tools against experimental results.

2.1 Notation Conventions and Definitions

This chapter describes the behaviour of an optical field in an optomechanical system. Most of these systems are made up of only two types of optical element: propagation in free space, and reflection or transmission through a mirror. There are several different conventions that are commonly used when describing such a system; in this section I outline the conventions used throughout this thesis, which are chosen to agree with those used in [14].

2.1.1 The optical field

Generally, we can describe a monochromatic optical field of angular frequency ω that is propagating in the z direction in terms of its electric field component:

$$E(z, t) = E_0(t) \cos(\omega t - kz + \delta(t)) \quad (2.1)$$

$$= E_1(t) \cos(\omega t - kz) - E_2(t) \sin(\omega t - kz), \quad (2.2)$$

where $E_0(t)$ describes the amplitude of the field at time t , and $\delta(t)$ describes a phase offset. $k = \omega/c$ is the wavenumber of the field, where c is the speed of light.

It is common to write this in exponential form:

$$E(z, t) = E'_0(t) e^{i(\omega t - kz)} + E_0'^*(t) e^{-i(\omega t - kz)}, \quad (2.3)$$

where

$$E'_0(z, t) = E_0(z, t) e^{i\delta(t)}, \quad (2.4)$$

and then simply work with the first term of equation 2.3 in our analysis, remembering that the physically meaningful quantity is the field amplitude $E(t)$.

We will commonly work at a single location, or a single moment in time. In the former case we can fully describe the field by setting $z = 0$, at which the field is given by

$$E(0, t) = \text{Re}\{E'_0(0, t) e^{i\omega t}\}. \quad (2.5)$$

In the latter case, used for example when we assume a *steady state*, we can equivalently set $t = 0$

without loss of generality, so the field is described by

$$E(z, 0) = \text{Re}\{E'_0(z, 0)e^{-ikz}\}. \quad (2.6)$$

2.1.2 Propagation in free space

If it does not interact with anything, a field at a position z_2 can be described in terms of the field at z_1 :

$$E_2 = E_1 e^{-ik(z_2 - z_1)} = E_1 e^{-ikD} = E_1 e^{-i\phi}, \quad (2.7)$$

where $D = z_2 - z_1$ is the distance between the two points and $\phi = kD$ describes the phase accumulated between the two positions, or the *propagation phase*.

We can consider this in the time domain: light arrives at point z_2 at a time $T = D/c$. Therefore

$$E_2(t) = E_1(t - T) = E_0 e^{i\omega(t - T)} = E_1(t) e^{-i\omega T} \equiv E_1(t) e^{-i\phi}, \quad (2.8)$$

since $\omega T = kD = \phi$ by definition.

Detuning

The phase accumulated on propagation is cyclic: $e^{-i\phi} = 1$ whenever $\phi = 2N\pi$, where N is an integer. It can therefore be convenient to think of propagation as being separated into a macroscopic length L , where $kL = 2N\pi$, and a microscopic length x , which describes a small *detuning* in addition to this. The propagation phase is thus described as:

$$\phi = kD = k(L + x) = \omega\left(\frac{L}{c} + \frac{x}{c}\right) = \omega\left(\tau + \frac{x}{c}\right) = \tau(\omega + \Delta), \quad (2.9)$$

where here we have introduced the *propagation time*, $\tau = L/c$. The detuning can be expressed in terms of this frequency, $\Delta = \omega \frac{x}{L}$. In cases where only a single optical frequency is present, we can typically discard the macroscopic term since $kL = 2N\pi$ and describe the behaviour just in terms of the detuning, so $\phi = kx$.

We can also split the frequency into parts: a reference frequency ω_0 which is used to define the macroscopic length, and a small detuning to that frequency, Ω . In its most generic form, the

propagation phase is therefore

$$\begin{aligned}\phi = kD &= (k_0 + k_\Omega)(L + x) = k_0L + k_0x + k_\Omega L + k_\Omega x \\ &= \tau(\omega_0 + \Delta + \Omega + \Delta_\Omega),\end{aligned}\tag{2.10}$$

where here we have redefined $\Delta = \omega_0 \frac{x}{L}$ and introduced $\Delta_\Omega = \Omega \frac{x}{L}$. In most cases of interest, both Ω and x are small compared to ω_0 and L , so the Δ_Ω term can be considered negligible, and by definition $\tau\omega_0 = 2N\pi$ so can be ignored.

2.1.3 Reflection and Transmission

As described in section 2.4 of [14], the magnitude and phase of a propagating electric field at the boundary between two media satisfy relations that are determined by the material properties of those media. This is simple in the case of a single reflective surface, but becomes more complicated for composite optical surfaces, as commonly used, for example, to create highly reflective coatings. Energy conservation, however, imposes the requirement that the phase of the reflected field differs from the transmitted field by a factor of $(N + \frac{1}{2})\pi$ where N is an integer. In cases where knowledge of the relative, rather than absolute, phase accumulated for each field component is sufficient, we can choose to define

$$E_{\text{refl}} = rE_{\text{in}} \quad E_{\text{trans}} = itE_{\text{in}},\tag{2.11}$$

where r and t define the real amplitude reflection and transmission coefficients for the surface. This convention is used throughout this thesis and by FINESSE.

2.2 Radiation Pressure

Electromagnetic waves carry momentum: the momentum density is given by \vec{S}/c^2 , where \vec{S} is the Poynting vector, describing the energy flux density, and c is the speed of light. If a light field travelling in the z direction strikes a surface with incident power P , it will in general impart a force on that surface. In the case of light normally incident upon a perfectly reflecting mirror, the total change in momentum $\Delta\vec{p}$ in time Δt results in a radiation pressure force \vec{F} :

$$\vec{F} = \frac{\Delta\vec{p}}{\Delta t} = \frac{2P}{c}\vec{z}.\tag{2.12}$$

2.3. TWO-MIRROR CAVITY

In all current and future ground-based gravitational wave detectors, the core optics and test masses are suspended, multi-stage pendulums with a low resonant frequency. This reduces the coupling between mirror movements and ground motion (seismic noise), improving the detector sensitivity. The suspended optics can be modelled as Simple Harmonic Oscillators (SHOs), so we can picture each as a single pendulum obeying the equation of motion

$$F_{\text{tot}} = M\ddot{z}(t) = -Kz(t) - b\dot{z}(t) + F_{\text{ext}}(t), \quad (2.13)$$

where the F_{tot} is the total force, M is the mass on the pendulum, z is the direction of motion, K is the spring constant and b is the damping coefficient. F_{ext} is an external force applied to the mass in the z direction, which could also have some time dependence. The resonant frequency of the suspension is $\omega_{m0} = \sqrt{K/M}$. We can describe the damping of the system in terms of its *quality factor*, $Q = \sqrt{MK}/b = M\omega_{m0}/b$, and these equations hold true for any degree of freedom that can be described as a SHO.

In this picture, we can treat a perfectly reflective suspended mirror as a bulk object of mass M , and illuminate it with a laser beam of power P directed at normal incidence at the centre of the mirror. Assuming no clipping losses, this will give a time-independent contribution to F_{ext} given by equation 2.12.

2.3 Two-Mirror Cavity

The simplest optical resonator, often referred to as a Fabry-Perot cavity, is a linear cavity formed by two partially-transmissive parallel mirrors. This configuration is illustrated in figure 2.1, illuminated by a laser at normal incidence to the input mirror.




Figure 2.1: A two-mirror, or *Fabry-Perot*, cavity. The steady-state behaviour of the cavity can be calculated in the frequency domain by propagating the field amplitudes, a , through the system. The resonant condition depends on the cavity length, D and the reflectivities and losses of the two mirrors.

2.3. TWO-MIRROR CAVITY

The optical response of a cavity can be described using a *transfer function*. Transfer functions can characterise any time-independent linear system by taking the complex ratio of the output to input signals, as a function of the signal frequency [15]. Simple optical transfer functions can be calculated by viewing the system in the frequency domain and considering the steady-state response, as described in [14]. The field from the laser that is incident on the input mirror is described by $E_{\text{in}} = a_{\text{in}} \cos(\omega t)$, where ω is the angular frequency of the laser and t is time. In this case of a single frequency field, we can simplify the notation and consider only the complex field amplitude, a . The amplitude of the circulating field incident on the input mirror (see figure 2.1) can be calculated in terms of the input field, resulting in the relation:

$$a_4 = a_{\text{in}} \frac{ir_2 t_1 e^{-i2\phi}}{1 - r_1 r_2 e^{-i2\phi}}, \quad (2.14)$$

where r_1 and r_2 , are the amplitude reflectivities of the two mirrors, t_1 is the transmissivity of the input mirror ($t_1 = \sqrt{1 - r_1^2}$ for a lossless mirror) and ϕ is the phase accumulated as the light propagates the length of the cavity. In the plane-waves approximation, $\phi = kD$ where D is the length of the cavity and $k = \omega/c$ is the wavenumber.

This can then be used to calculate the dependence of the circulating power on the cavity length:

$$P_c(D) = \frac{T_1 P_{\text{in}}}{|1 - r_1 r_2 e^{2i\phi}|^2} = \frac{T_1 P_{\text{in}}}{1 + R_1 R_2 - 2r_1 r_2 \cos(2kD)}, \quad (2.15)$$

where $P_{\text{in}} = |a_{\text{in}}|^2$ is the input power and $T_{1,2} = t_{1,2}^2$, $R_{1,2} = r_{1,2}^2$ are the power transmissivities and reflectivities of the mirrors. This function is plotted in figure 2.2, showing that the circulating power is maximal for $D = N\pi c/\omega = N\lambda/2$ where N is an integer and λ is the wavelength of the laser. When this is true, light entering the cavity is in phase with light already circulating and interferes constructively.

Equation 2.15 can be used to find properties of the optical resonator. The *pole frequency*, f_{pole} is the change in frequency required to reduce the circulating power to half its maximum value, i.e. the half-width-half-maximum. The bandwidth of the resonator, γ is therefore:

$$\begin{aligned} \gamma = 2\pi \times f_{\text{pole}} &= 2\pi \times \frac{c}{2D} \times \frac{1}{\pi} \arcsin\left(\frac{1 - r_1 r_2}{2\sqrt{r_1 r_2}}\right) \\ &= 2\pi \times \text{FSR} \times \frac{1}{2\mathcal{F}}, \end{aligned} \quad (2.16)$$

where we define the $\text{FSR} = c/2D$ as the *free spectral range*, describing the distance between resonant

2.3. TWO-MIRROR CAVITY

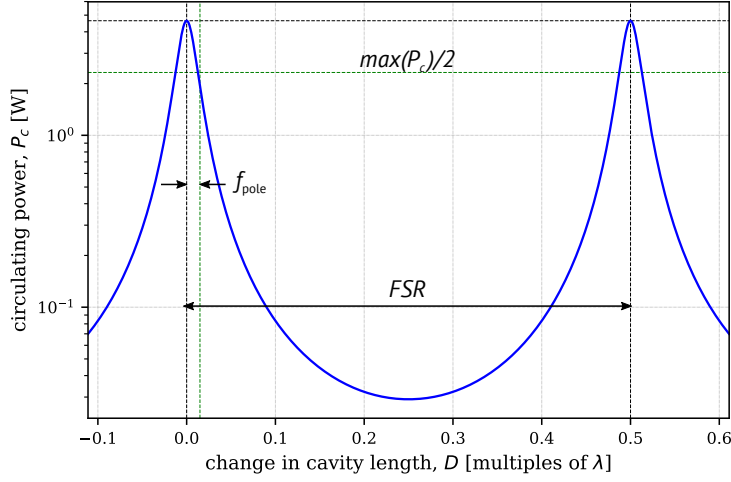


Figure 2.2: Power circulating in a simple 2-mirror cavity as a function of cavity length: the optical power in a cavity is position-dependent. By ‘scanning’ a cavity in this manner, the pole frequency, f_{pole} and free spectral range, FSR can be identified.

peaks, and \mathcal{F} as the *finesse* of the cavity, which is analogous to the Q-factor of a mechanical resonator.

We can also use steady state calculations find the reflection transfer function (sometimes referred to as the *input-output relation*) for the cavity:

$$\frac{a_r}{a_{\text{in}}} = \frac{r_1 - r_2 e^{-i2\phi}}{1 - r_1 r_2 e^{-i2\phi}}. \quad (2.17)$$

In the limit that both mirrors are highly reflective, and the propagation phase accumulated in the cavity is small, we can apply additional approximations and definitions. These are summarised in table 2.1. In these limits equation 2.17 can be re-written in the compact form:

$$\frac{a_r}{a_{\text{in}}} \approx \frac{-(\omega + \Delta) - i(\gamma_1 - \gamma_2)}{-(\omega + \Delta) + i(\gamma_1 + \gamma_2)} \approx \frac{-\omega - i\gamma_1}{-\omega + i\gamma_1} \sim -e^{-2i\frac{\omega}{\gamma}} \quad (2.18)$$

where ω is the laser angular frequency and $\gamma_{1,2} \equiv \frac{T_{1,2}}{4\tau}$ describes the power transmissivities of the mirrors scaled by the propagation time in the cavity, τ . Δ describes the relative detuning between the laser and the cavity as an angular frequency, assuming a fixed external reference. The second simplification is valid when using a perfectly reflective end mirror and zero detuning; the final form comes from a comparison of Taylor expansions.

Assumption/Statement	Simplification
Lossless mirrors	$R + T = 1$
Small propagation phase: $\phi \ll k\lambda$	$e^{-i2\phi} \approx 1 - i2\phi$
Cavity length description	$D = L + x$
Cavity 1-way propagation time	$\tau = L/c$
Detuning	$\Delta = \omega x/L$
Propagation phase using frequencies	$\phi = kD = \frac{\omega}{c}(L + x) = \tau(\omega + \Delta)$
High-reflectivity mirrors: $T \ll 1$	$r_{1,2} = \sqrt{1 - T_{1,2}} \approx 1 - \frac{T_{1,2}}{2}$
‘Bandwidths’ in terms of mirror transmissions	$\gamma_{1,2} \equiv \frac{T_{1,2}}{4\tau}$ so $r_{1,2} \approx 1 - 2\tau\gamma_{1,2}$
$\gamma_{1,2}$ are small since $T_{1,2}$ are small	discard terms quadratic in $\gamma_{1,2}$ and ϕ

Table 2.1: List of approximations and definitions used to simplify optical transfer function equations

2.4 Optical Springs

In cases where one or more of the optics in a cavity is suspended, the radiation pressure force from the light circulating in the cavity will push on the optics, changing the cavity length and therefore whether the light resonates. This changes the optical power in the cavity, and the resulting radiation pressure force. This feedback loop is analogous to a simple harmonic oscillator, and so the behaviour is often referred to as an *optical spring*.

Figure 2.1 and equation 2.15 show that the power circulating in an optical cavity is dependent on the positions of the mirrors. The radiation pressure acting on these mirrors is therefore also position-dependent. If one or more of the mirrors is suspended, the cavity behaviour becomes dynamic and dependent on the power currently circulating in it.




Figure 2.3: A two-mirror cavity with suspended end mirror. The steady-state behaviour of the cavity are again calculated in the frequency domain by propagating the field amplitudes, a , through the system. In this example, we consider a perfectly reflective end mirror, assume both mirrors are (optically) lossless, and ignore clipping effects.

For simplicity, we look at the case of one suspended and one fixed optic, as depicted in figure 2.3, and ignore any clipping or surface imperfections. The radiation pressure force on the end mirror depends directly on the power incident on the mirror, which depends on the position of the end

mirror, x . We can Taylor expand equation 2.12 to understand the behaviour:

$$F_{\text{rad}} = \frac{2P_c(x)}{c} \approx \frac{2}{c} \left(P_c(0) + \frac{dP_c}{dx}(0)x + \dots \right). \quad (2.19)$$

First is a DC term, which is independent of the mirror position. The second term has a linear dependence on x , with a form that matches $F = -K_{\text{opt}}x$. This force is equivalent to having the cavity mirrors connected by a mechanical spring. We therefore define the *optical rigidity* as

$$K_{\text{opt}} = -\frac{2}{c} \frac{dP_c}{dx}. \quad (2.20)$$

K_{opt} can be positive or negative, and is typically complex, as will be described below. In cases where it is positive F_{rad} acts as a restoring force, and so the system is commonly referred to as an *optical spring*.

2.4.1 The Adiabatic Approximation

If the cavity is adiabatically detuned from resonance such that its optical response can be considered instantaneous (i.e. the change in propagation time of the light in the cavity due to detuning is negligible), the power incident on the mirror is as in equation 2.15 above, and we can directly evaluate the adiabatic optical rigidity using equation 2.20 as shown in figure 2.4. We can rewrite this in approximate form using the definitions given in table 2.1 above, assuming that the cavity detuning is small and the transmissions of the mirrors are $T_1 \ll 1$ and $T_2 = 0$. In this case, the denominator in the middle form of equation 2.15 can be rewritten as

$$1 - r_1 e^{2i\phi} \approx 1 - (1 - T_1/2)(1 + 2i\phi) \approx T_1/2 - 2i\phi, \quad (2.21)$$

where we have ignored the very small $T_1 \times \phi$ term. Equation 2.15 thus becomes

$$P_c \approx \frac{T_1 P_{\text{in}}}{(T_1/2 - 2i\phi)(T_1/2 + 2i\phi)} \approx \frac{T_1 P_{\text{in}}}{(T_1^2/4 + 4\phi^2)} \quad (2.22)$$

$$= \frac{4P_{\text{in}}}{T_1} \frac{\gamma^2}{\gamma^2 + \Delta^2}, \quad (2.23)$$

where we use $\phi = kx = \Delta\tau$ and $\gamma \equiv \frac{T_1}{4\tau}$ in the final step.

The optical rigidity in the adiabatic picture is therefore

$$\begin{aligned} K_{\text{opt}} &= -\frac{2}{c} \frac{dP_c}{dx} = -\frac{2}{c} \frac{dP_c}{d\Delta} \frac{d\Delta}{dx} = -\frac{2}{c} \left(-\frac{4P_{\text{in}}}{T_1} \frac{2\gamma^2\Delta}{(\gamma^2 + \Delta^2)^2} \right) \frac{\omega}{L} \\ &= \tilde{K}_{\text{opt}} \frac{\Delta}{\gamma^2 + \Delta^2}, \end{aligned} \quad (2.24)$$

with

$$\tilde{K}_{\text{opt}} = \frac{64P_{\text{in}}k\gamma}{cT_1^2} \frac{\gamma^2}{\gamma^2 + \Delta^2}, \quad (2.25)$$

and varies linearly with detuning for small detunings, as depicted in figure 2.4. It can be useful to describe the detuning as a dimensionless quantity by normalising it using the cavity bandwidth such that $\delta_\gamma = \Delta/\gamma$. This leads to the definition of the adiabatic spring constant K as derived in section 3.1 of [16]¹.




Figure 2.4: Optical rigidity in the adiabatic approximation, as a function of the cavity detuning. The dashed line is a direct evaluation based on the gradient of the circulating power in equation 2.15; the solid line shows the form approximated by equation 2.24. The x-axis here is the detuning, Δ , expressed in degrees: a detuning of 360° is equivalent to a cavity length change of $D \rightarrow D + \lambda$ where λ is the wavelength of the main light field.

We describe the cavity as *blue* ($\Delta > 0$) or *red* ($\Delta < 0$) detuned, depending on the relative resonant frequency of the cavity compared to the frequency of the input laser. Positive detunings increase the cavity length and decrease the resonant frequency of the cavity compared to the laser frequency. In this case, K_{opt} is negative and results in a Hooke’s Law restoring force: moving further above the resonance reduces the power circulating, so the force on the mirror reduces and the mirror moves inwards. This reduces the detuning, so the cavity power rises again, pushing the mirror out

¹note that by choosing to define $F = -K_{\text{opt}}x$, our K_{opt} has the opposite sign to that in [16]. This choice means that a positive value of K_{opt} corresponds to a restoring force, which is consistent with the convention commonly used for mechanical springs.

once more. This is called an optical spring, with an associated resonant frequency.

With red detuning, K_{opt} becomes positive so there is an anti-restoring force. Moving further below resonance, the power drops causing the mirror to move inwards. However here this causes the power to reduce further, and the mirror does not recover its original position: an *optical anti-spring*.

2.4.2 The Frequency-dependent Optical Rigidity

So far the equations above have been valid for mirror motions at frequencies much less than the cavity linewidth: since $1/\gamma$ represents the cavity's response time, motions much slower than this mean that the cavity's response is effectively instantaneous and all motions are quasistatic. However, in general the cavity will take some time to reach the steady state after a change in length, leading to a delay between mirror motion and the change in the applied force. The optical rigidity can then be separated into a position-dependant term and a velocity-dependent term, the latter of which is referred to as the *optical damping*.

This can be taken into account by introducing a frequency dependence to the optical rigidity, to take into account the different phase relations that the carrier and additional frequency components will see in the cavity. This can be achieved using a *perturbative model* of the optical fields, which approximates the behaviour in the linear regime by considering only three field components: the zero-th order carrier at frequency ω , $E^{(0)}$ and two sidebands at frequencies $\omega \pm \Omega$, which are treated as a first order perturbation, $E^{(1)}$. This is valid under the assumption that the amplitude of the oscillation producing the sidebands is small, and is sometimes referred to as the *linear regime* approximation.

The fields in the cavity therefore have form:

$$E \approx E^{(0)} + E^{(1)} = a^{(0)} e^{i\omega t} + a^+ e^{i(\omega+\Omega)t} + a^- e^{i(\omega-\Omega)t} + \text{c.c.} \quad (2.26)$$

where $a^{(0)}$ is the complex field amplitude of the (zero-th order) carrier and a^\pm are the complex field amplitudes of the two (first order) sideband fields, which are introduced through amplitude and phase modulations in the system. We use 'c.c.' to refer to the complex conjugate.

Reflection from a Radiation-Pressure Driven Oscillating Mirror

We first take into account that the suspended mirror is moving, driven by a changing radiation pressure force as the cavity length changes. The displacement of the mirror is described by

$$x(t) = x_{\text{RP}} \cos(\Omega t) \quad (2.27)$$

where x_{RP} here is the amplitude of the mirror motion, which is determined by the power incident on the mirror and the mirror's mechanical susceptibility. Note that we have now separated out the static detuning of the cavity, Δ , from the mirror displacement, $x(t)$.

A mirror oscillating at a frequency Ω and amplitude x_{RP} will phase modulate the optical field reflected from it, since the propagation distance is changing. In the linear regime approximation ($kx_{\text{RP}} \ll 1$) the field reflected from the suspended mirror therefore has form:

$$E_3(\Omega) = r_2 E_2(\Omega) e^{2ikx(t)} \approx r_2 E_2(\Omega) \left(1 + ikx_{\text{RP}} (e^{i\Omega t} + e^{-i\Omega t}) \right), \quad (2.28)$$

as described in [14], where $k = \omega/c$ is the wavenumber of the carrier. To first order this simplifies to

$$E_3 \approx r_2 \left((E_2^{(0)} + E_2^{(1)}) + E_2^{(0)} ikx_{\text{RP}} (e^{i\Omega t} + e^{-i\Omega t}) \right). \quad (2.29)$$

It is useful to collect together the field components with different frequencies (the ‘sideband picture’). In this case we expand E_3 using the notation of equation 2.26 to write

$$E_3 \approx r_2 (E_2^{(0)} + E_2^{(1)}) + r_2 E_2^{(0)} ikx_{\text{RP}} (e^{i\Omega t} + e^{-i\Omega t}) \quad (2.30)$$

$$\begin{aligned} &= r_2 a_2^{(0)} e^{i\omega t} + r_2 (a_2^+ + a_2^{(0)} ikx_{\text{RP}}) e^{i(\omega+\Omega)t} + r_2 (a_2^- + a_2^{(0)} ikx_{\text{RP}}) e^{i(\omega-\Omega)t} \\ &\quad + r_2 a_2^{(0)*} e^{-i\omega t} + r_2 (a_2^{+*} + a_2^{(0)*} ikx_{\text{RP}}) e^{-i(\omega+\Omega)t} + r_2 (a_2^{-*} + a_2^{(0)*} ikx_{\text{RP}}) e^{-i(\omega-\Omega)t} \\ &= a_3^{(0)} e^{i\omega t} + a_3^+ e^{i(\omega+\Omega)t} + a_3^- e^{i(\omega-\Omega)t} + \text{c.c.} \end{aligned} \quad (2.31)$$

Steady-state equations in the cavity

We can now propagate the complex field amplitudes as usual. The zero-th order fields obey the same conditions as those used to derive the simple two-mirror cavity equations in section 2.3.

Assuming that the input field a_{in} is a pure carrier field, the first order field incident on the end mirror is then described by the behaviour of the two sidebands:

$$a_2^\pm = e^{-i\phi_\pm} a_1^\pm = e^{-i\phi_\pm} r_1 a_4^\pm = e^{-2i\phi_\pm} r_1 a_3^\pm = e^{-2i\phi_\pm} r_1 r_2 (a_2^\pm + a_2^{(0)} i k x_{\text{RP}}), \quad (2.32)$$

where ϕ_\pm describes the propagation phase accumulated by the sideband field frequencies: $\phi_\pm = \tau(\omega_0 \pm \Omega + \Delta)$.

Solving equation 2.32 for a_2^\pm then lets us describe the sideband field amplitudes in terms of the carrier amplitude:

$$a_2^\pm = i k x_{\text{RP}} \frac{r_1 r_2 e^{-2i\phi_\pm}}{1 - r_1 r_2 e^{-2i\phi_\pm}} a_2^{(0)} = f^\pm a_2^{(0)}, \quad (2.33)$$

where f^\pm is introduced to compact the notation.

Frequency-dependent Optical Rigidity, $K_{\text{opt}}(\Omega)$

The power incident on the end mirror in this model is given by

$$\begin{aligned} P_2 &= |E_2|^2 \approx |E_2^{(0)} + E_2^{(1)}|^2 \\ &\approx |E_2^{(0)}|^2 + E_2^{(0)*} E_2^{(1)} + E_2^{(0)} E_2^{(1)*} + \mathcal{O}((k x_{\text{RP}})^2), \end{aligned} \quad (2.34)$$

The field amplitudes above can be substituted into this power expansion to give:

$$\begin{aligned} P_2(\Omega) &\approx |E_2^{(0)}|^2 + E_2^{(0)*} E_2^{(1)*} + E_2^{(0)*} E_2^{(1)} \\ &= |a_2^{(0)}|^2 + a_2^{(0)*} (a_2^{+*} e^{-i\Omega t} + a_2^{-*} e^{i\Omega t}) + a_2^{(0)*} (a_2^+ e^{i\Omega t} + a_2^- e^{-i\Omega t}) + \text{c.c.} \\ &= 2|a_2^{(0)}|^2 \left[1 + (f^- + f^{+*}) e^{-i\Omega t} + (f^{-*} + f^+) e^{i\Omega t} \right] \\ &= P_2^{(0)} + P_2^{(1)}(\Omega) e^{-it\Omega} + P_2^{(1)*}(\Omega) e^{it\Omega} \end{aligned} \quad (2.35)$$

Once again, the first term is the DC power circulating in the cavity, while the remaining terms, at frequency $\pm\Omega$, describe the dynamic behaviour of the optomechanical system to first order. The frequency-dependent optical susceptibility, $K_{\text{opt}}(\Omega)$ is therefore given by

$$K_{\text{opt}}(\Omega) = -\frac{2}{c} \frac{dP_2^{(1)}(\Omega)}{dx_{\text{RP}}}, \quad (2.36)$$

which can be evaluated directly as shown in figure 2.5.

We can simplify the analytical form using the *single mode approximation*, which is valid for cases where the frequency of the mirror oscillation is much less than the free spectral range of the cavity, i.e. $\Omega\tau \ll 1$ and so $e^{2i\phi_{\pm}} \approx 1 + 2i\phi_{\pm}$. We then assume $T_2 = 0$ and $T_1 \ll 1$ and use the definitions in table 2.1 as before to write:

$$\begin{aligned}
 P_2^{(1)}(\Omega) &= 2|a_2^{(0)}|^2 kx_{\text{RP}} \left(\frac{ir_1 r_2 e^{-2i\phi_+}}{1 - r_1 r_2 e^{-2i\phi_+}} - \frac{ir_1 r_2 e^{+2i\phi_-}}{1 - r_1 r_2 e^{+2i\phi_-}} \right) \\
 &\approx 2P_c kx_{\text{RP}} \left(\frac{i(1 - T_1/2 - 2i\phi_+)}{T_1/2 + 2i\phi_+} - \frac{i(1 - T_1/2 + 2i\phi_-)}{T_1/2 - 2i\phi_-} \right) \\
 &= 2P_c kx_{\text{RP}} \frac{2(\phi_+ + \phi_-)}{T_1^2/4 + 4\phi_+\phi_- + iT_1(\phi_+ - \phi_-)} \\
 &= 2P_c kx_{\text{RP}} \frac{4\Delta\tau}{T_1^2/4 + 4\tau^2(\Delta^2 - \Omega^2) + 2iT_1\Omega\tau} \\
 &= \frac{2P_c kx_{\text{RP}}\Delta}{\tau(\Delta^2 + (\gamma + i\Omega)^2)}
 \end{aligned} \tag{2.37}$$

where we use that $\phi_{\pm} = \tau(\Delta \pm \Omega)$ as defined above. The frequency-dependent optical susceptibility is thus given by

$$\begin{aligned}
 K_{\text{opt}}(\Omega) &= -\frac{2}{c} \frac{dP_2^{(1)}(\Omega)}{dx_{\text{RP}}} \\
 &= -\frac{4\omega P_c}{cL} \frac{\Delta}{\Delta^2 + (\gamma + i\Omega)^2} \\
 &= \tilde{K}_{\text{opt}} \frac{\Delta}{\Delta^2 + (\gamma + i\Omega)^2},
 \end{aligned} \tag{2.38}$$

where \tilde{K}_{opt} is as defined in equation 2.25, and we see that this reduces to the adiabatic form when $\Omega = 0$.

We see that $K_{\text{opt}}(\Omega)$ can now take complex values; the real term is the spring term (displacement quadrature) while the complex part is the damping (velocity quadrature). When the mirror is quasistatic, with negligible Ω , we return to the simple optical spring described previously. If instead the detuning is very small, $\Delta \ll \gamma$, the behaviour is dominated by the damping. We can see that the damping term is positive for positive detunings and negative for $\Delta/\gamma < 0$: motion is damped above resonance and *anti*-damped below it. This is summarised in figure 2.5. In most systems, then, the system will be unstable: either it is red detuned, with a damped anti-restoring force, or blue detuned, with an anti-damped restoring force.

If, as at Advanced LIGO and other detectors, the arm cavities are *locked* on resonance (i.e. are




Figure 2.5: $K_{opt}(\Omega)$ as a function of the cavity detuning in the perturbative model. The solid line is a direct evaluation based on the gradient of the circulating power; the dashed line shows the form approximated by equation 2.38. The result given by the adiabatic approximation (equation 2.24) is shown for reference. The real part of $K_{opt}(\Omega)$ describes the optical rigidity, while the imaginary part is the optical damping. Both depend on the frequency of the mirror motion, and typically take opposite signs, meaning that the overall system, without additional feedback, is unstable for both positive and negative detunings.

held on resonance by actuators using a feedback loop), the detuning of the mirror is controlled and held at almost zero. This means that the effects of the (anti-)restoring force are negligible, and the optical rigidity is low in comparison to mechanical rigidity of the system. However, the damping term can still be large depending on the frequency of mirror motion. If the system is red detuned ($\Delta < 0$), this can reinforce any mechanical damping present, reducing the quality factor of the mechanical motion—an effect called *cold damping*. However, if the system is blue detuned ($\Delta > 0$) there is antidamping, which can exceed the mechanical damping and cause oscillations to ring up.

2.4.3 Overall behaviour of the optomechanical system

We can take into account the mechanical properties of the suspended mirror to find the overall behaviour of the optomechanical system. In this case, the forces on the suspended mirror can be summarised as:

$$F = F_0 + F_{\text{rad}} + F_{\text{ext}}. \quad (2.39)$$

On comparison with equation 2.13, F_0 is the equation of motion for the free suspended mirror, and we here separate the additional external forces into F_{rad} , the forces due to the radiation pressure of the optical field, and F_{ext} , the forces due to any additional external elements (e.g. gravity). In the frequency domain, the equation of motion for the free mirror is

$$F_0 = m\omega_{m0}^2 x - im\gamma_{m0}\omega_{m0}x \quad (2.40)$$

where we have used γ_{m0} and ω_{m0} as the bare mechanical damping and resonant frequency of the mirror.

As noted above, K_{opt} has real and imaginary components, corresponding to a spring frequency ω_{opt} and a damping rate γ_{opt} . We therefore rewrite the radiation pressure force to match the form of a simple harmonic oscillator as above:

$$\begin{aligned} F_{\text{rad}} = -K_{\text{opt}}x &= -\text{Re}\{K_{\text{opt}}\}x - i\text{Im}\{K_{\text{opt}}\}x \\ &= m\omega_{\text{opt}}^2 x - im\Omega\gamma_{\text{opt}}x. \end{aligned} \quad (2.41)$$

The full equation of motion for this system is therefore described by

$$\begin{aligned} F &\approx (m\omega_{m0}^2 - im\Omega\gamma_{m0} + m\omega_{\text{opt}}^2 - im\Omega\gamma_{\text{opt}})x + F_{\text{ext}} \\ &= (m\omega_m^2 - im\gamma_m\Omega)x + F_{\text{ext}}, \end{aligned} \quad (2.42)$$

Where we have defined a new resonant frequency, ω_m and damping rate, γ_m for the overall optomechanical system (assuming $F_{\text{ext}} = 0$ or is DC):

$$\omega_m = \sqrt{\omega_{m0}^2 + \omega_{\text{opt}}^2} \quad (2.43)$$

$$\gamma_m = \gamma_{m0} + \gamma_{\text{opt}}. \quad (2.44)$$

By direct substitution of equation 2.38 into equation 2.41, we find that the optical spring frequency and damping rate terms are:

$$\omega_{\text{opt}} = \sqrt{\frac{K'(\Delta^2 + \gamma^2 - \Omega^2)}{M}} \quad \text{and} \quad \gamma_{\text{opt}} = \frac{-2\gamma K'}{M}, \quad (2.45)$$

where

$$K' = \tilde{K}_{\text{opt}} \frac{\Delta}{(\Delta^2 + \gamma^2 - \Omega^2)^2 + 4\Omega^2\gamma^2} \quad (2.46)$$

2.5 Gaussian Beams and Higher Order Optical Modes

In the previous sections we have modelled the optical field as a plane wave. This is sufficient to understand one-dimensional effects such as a linear optical spring, however it is not a realistic description of the optical field. To understand more complex behaviours, we move away from the plane wave approximation and account for spatial properties of the light field perpendicular to the propagation direction.

Any 2D spatial distribution can be described as a sum of Gaussian modes [14, 17]. For example, the Hermite-Gauss modes form a complete, orthogonal, infinite, discrete set of spatial distributions, the first few of which are depicted in figure 2.6. The indices n and m determine the number of nodes in the x and y directions (with the beam propagating in the z direction):

$$\begin{aligned} E_{n,m}(x, y, z) &= u_{nm}(x, y, z)e^{ikz} \\ u_{nm}(x, y, z) &= u_n(x, z)u_m(y, z) \\ u_n(x, z) &= \left(\frac{2}{\pi}\right)^{1/4} \left(\frac{\exp(i(n+1)\Psi(z))}{2^n n! w(z)}\right)^{1/2} H_n\left(\frac{\sqrt{2}x}{w(z)}\right) \exp\left(-i\frac{kx^2}{2R_C(z)} - \frac{x^2}{w^2(z)}\right) \end{aligned} \quad (2.47)$$

where $w(z)$ and $R_C(z)$ describe the beam spot size and radius of curvature respectively, as shown in figure 2.7, and $\Psi(z)$ is the Gouy phase, which will be described later. $H_{n,m}$ are the Hermite polynomials. The sum $n + m$ describes the order of the mode; typically any value greater than 0 is considered a *higher order mode* (HOM). This zero-th order mode (HG_{00}) is called the fundamental mode, and depicts the typical laser beam shape with Gaussian intensity profile. However, combinations of HOMs can be used to describe any small perturbations to the beam, including any distortions introduced due to mirror surface roughness, as well as effects such as misalignments, astigmatism, and more [18].




Figure 2.6: Spatial distribution of the lowest order Hermite-Gauss modes. The indices n and m describe the intensity distribution in the x and y directions.




Figure 2.7: Propagation of a Gaussian beam in free space

As shown in figure 2.7, a Gaussian beam has an ‘hourglass’ profile, the narrowest point of which is referred to as the *waist*, w_0 . If the wavelength, waist size, and waist position are known, the spot size, w and wavefront curvature, R_C can be calculated anywhere along the beam axis as it propagates in free space:

$$w(z) = w_0 \sqrt{1 + \left(\frac{z - z_0}{z_R} \right)^2} \quad (2.48)$$

$$R_C(z) = z - z_0 + \frac{z_R^2}{z - z_0}, \quad (2.49)$$

where z_0 is the waist position and $z_R = \frac{\pi w_0^2}{\lambda}$ is referred to as the *Rayleigh range*. These beam properties are determined by the geometry of the optics that the light interacts with. For a beam to resonate in a cavity, the field which has circulated through the cavity must constructively interfere with the field entering. This requires an overlap in the spatial distribution as well as the phase relations discussed in section 2.3, i.e. the cavity must be *geometrically stable*. In the case of a simple two-mirror cavity, the stability is dependent on the cavity length, L and the curvature of

the two mirrors, $R_{1,2}$ according to:

$$g = g_1 g_2 = \left(1 - \frac{L}{R_1}\right) \left(1 - \frac{L}{R_2}\right), \quad (2.50)$$

where the cavity is stable for $0 \leq g \leq 1$ [17]. In more complex optical configurations, the properties of the Gaussian beam at any point, and geometric stability of the system, can be calculated using *ABCD* matrices, as detailed in section 9.13 onwards of [14].

Of particular importance when considering HOMs is the *Gouy phase*, which in equation 2.47 is represented by $\Psi(z)$:

$$\Psi(z) = \tan^{-1} \left(\frac{z - z_0}{z_R} \right). \quad (2.51)$$

This term is introduced since the phase velocity of a Gaussian beam is reduced compared to a plane wave.

The total phase accumulated by a mode HG_{nm} of angular frequency ω when propagating in free space is described by

$$\phi_{nm}(\omega) = \frac{\omega L}{c} + \Psi_{nm}(z) + \delta, \quad (2.52)$$

where, as in section 2.1.2, L describes the macroscopic length of the space, and $\delta = \frac{\omega x}{c}$ here represents any microscopic detuning. $\Psi_{nm}(z)$ refers to the total Gouy phase accumulated by the mode. This is dependent on the mode order:

$$\Psi_{nm}(z) = \left(n + \frac{1}{2}\right) \Psi_x(z) + \left(m + \frac{1}{2}\right) \Psi_y(z) \quad (2.53)$$

$$\stackrel{\Psi_x = \Psi_y}{=} (n + m + 1) \Psi(z) \quad (2.54)$$

meaning that in an optical cavity, each mode order will resonate for a slightly different length, i.e. have a different resonant frequency. Ψ_x and Ψ_y refer to the Gouy phase in the x - and y -directions; in cases with no astigmatism the beam is symmetric and so the Gouy phase becomes degenerate.

The difference in resonant frequency for different mode orders is given by the *higher order mode separation frequency*, f_{msf} [14]:

$$f_{\text{msf}} = \frac{\Psi_{\text{RT}}}{2\pi} \text{FSR}, \quad (2.55)$$

where Ψ_{RT} is the total Gouy phase accumulated on one *round trip* in the cavity, which will depend on the cavity geometry. FSR is the free spectral range of the cavity, as described in section 2.3.

2.6 Parametric Instability

A suspended mirror is not a point mass, as the equations above describe, but rather an extended bulky object, which can oscillate in more than just one dimension. Rather than an ideal light field hitting an ideal mirror exactly centrally and at normal incidence, there may also be an angle of incidence, α , between the field and the mirror normal, or the beam may be off-centre, in which case torques may be induced. As a bulk object, the mirror also has internal resonances: mechanical modes which can distort the mirror surface and therefore change the interaction between the mirror and the beam. These mechanical modes are of particular importance, as their effect cannot be controlled through beam alignment [19].

Any motion of the mirror surface, at angular frequency Ω , will introduce phase modulation to the beam it reflects. As introduced in section 2.4.2, for small modulation amplitudes phase modulation can be considered as introducing two *sideband* fields at frequencies $\omega_0 \pm \Omega$, where ω_0 is the angular frequency of the incident beam. Where these fall within the bandwidth of the cavity, it is called the *unresolved* regime; in the *resolved* sideband regime the sidebands are outside the cavity bandwidth and so see the cavity very differently to the carrier.

The two sidebands represent transfers of energy between the optical field and mechanical motion [11]. Although this is a classical phenomenon, it is common to describe these in terms of quantum mechanics, as depicted in figure 2.8. The lower sideband represents the *Stokes* mode, in which the mirror extracts energy from the light field and is driven at Ω ; the upper sideband is the *anti-Stokes* mode, where the reverse happens—the mirror’s motion is damped as energy is transferred to the light field.




Figure 2.8: Feynman diagrams depicting the modes of coupling between the optical and mechanical modes at mechanical frequency Ω .

In general a cavity will have an asymmetric response to the sidebands, meaning that one will be amplified more than the other. If this is an anti-Stokes mode, this can be useful, since it will damp the mirror’s motion at that frequency. However, if the Stokes mode is favoured, and the overlap between the optical and mechanical modes is sufficient to give a round-trip gain greater

than 1, the resulting antidamping can exceed the mechanical damping so that oscillations ring up: a *parametric instability* will be induced. Typically we use ‘parametric instability’, or ‘PI’, to refer to the unstable coupling of a mechanical mode in a bulk optic to a higher order optical mode.

2.6.1 The Feedback Picture and Parametric Gain

In this section we consider the linear interaction of an optical field with a mechanical motion in the optical set up depicted in figure 2.9. A two-mirror cavity, formed from a fixed input mirror and a suspended end mirror of mass M , is resonant for a high-power Gaussian beam of angular frequency ω_0 . It is controlled so that the DC radiation pressure force on the end mirror can be ignored. A mechanical mode in the end mirror, at mechanical angular frequency Ω_m and quality factor Q_m , phase modulates the incident optical field. This scatters light from the incoming (‘pump’) field into higher order optical modes at sideband frequencies, determined by the shape and frequency of the mirror’s surface motion. These new field components then propagate through the cavity.




Figure 2.9: A two mirror cavity in which the end mirror is mechanically excited.

As described in section 2.5, different HOMs will have different resonant frequencies in the cavity, due to the Gouy phase accumulated according to their mode order. If new field components, introduced by the optomechanical interaction, coincide with these resonances, and have the correct transverse shape, they can be resonantly enhanced by the cavity. If the Stokes mode (lower sideband) is enhanced, this results in a *parametric instability*, or ‘PI’: the mechanical mode in the mirror is driven by radiation pressure at its resonant frequency, due to the beat note between the pump and sideband optical fields. Energy is coupled from the optical field into the mirror, causing the runaway excitation of the mechanical mode. This can saturate control systems used to hold the set up at a useful operating point.

Evans *et al* [20] showed that parametric instabilities can be described as a feedback process, as

depicted in figure 2.10. We thus characterise parametric instabilities using the *parametric gain*, R , which compares the rate of damping induced optically, γ_{opt} to the mechanical damping rate, γ_{m0} [16]:

$$R = -\frac{\gamma_{\text{opt}}}{\gamma_{m0}} \quad (2.56)$$




Figure 2.10: Parametric instability, depicted as a feedback process. A mechanical excitation in the mirror scatters light from the pump field into higher order optical modes at sideband frequencies determined by the frequency of the mechanical motion. These propagate through the optical system, interacting with the pump field and mechanical motion of the mirror via radiation pressure. If the optomechanical interaction is sufficiently strong, and the optical gain of the lower sideband in the optical system is sufficiently high, an instability results, ‘ringing up’ the mirror. This figure is reproduced from [1].

R is defined such that $R = 0$ for no optical damping, $R < 0$ when the optical damping reinforces the mechanical damping ($R = -1$ means the mechanical damping is doubled), and $R > 0$ when the optical damping opposes the mechanical. When $R = 1$, the optical and mechanical damping terms exactly cancel each other, so $R > 1$ implies that the system is unstable.

We see that we can also describe optical springs in terms of a parametric gain, by substituting equation 2.45 into equation 2.56 directly. In this sense, optical springs can be considered the simplest case of a parametric instability: the mechanical excitation is the fundamental motion of the suspended mirror in the beam direction, and it is driven by the fundamental optical mode.

To describe parametric instabilities in their more general sense, we must account for the required resonant conditions for the frequency and shape of the optical sidebands. Braginsky first developed this formalism for a two-mirror cavity in 2001 [11]; here we combine this with the notation shown in [20, 21, 22].

For a higher order mode n driven by a mechanical motion with angular frequency ω_m , the frequency condition is described using the overlap factor, Λ_ω :

$$\Lambda_\omega = \frac{1}{1 + (\Delta\omega/\gamma_n)^2}, \quad (2.57)$$

where $\Delta\omega = \omega_0 - \omega_n - \omega_m$ describes the difference in frequency between the optical beat note and

the mechanical resonant frequency. ω_0 is the frequency of the fundamental HG_{00} mode, which is the pump field. ω_n is the frequency of the higher order mode optical sideband, and γ_n is the optical bandwidth (full width half maximum) of the HOM.

The transverse spatial distribution of the optical field must also match the surface deformation of the mechanical excitation in order to drive it (or suppress it) effectively. This is characterised by a dimensionless spatial overlap factor, Λ_s :

$$\Lambda_s = \frac{V \left(\int f_0(\vec{r}_\perp) f_n(\vec{r}_\perp) u_z d\vec{r}_\perp \right)^2}{\int |f_0|^2 d\vec{r}_\perp \int |f_n|^2 d\vec{r}_\perp \int |\vec{u}|^2 dV} \quad (2.58)$$

Where f_0 and f_n are functions describing the spatial distribution of the fundamental and higher order optical modes respectively, \vec{u} is the spatial vector of displacements in the elastic mode, u_z is its z -component, and $\int d\vec{r}_\perp$ and $\int dV$ correspond to integrating over the mirror's surface and volume, V respectively.

The overall scaling of R is then governed by the energy stored and transferred in the optomechanical system. For this simplest case of a mechanical mode m interacting with a single higher order optical mode n in a 2-mirror linear cavity, the parametric gain is described by

$$R_{n,m} = \frac{4P}{McL} \frac{Q_m}{\omega_m^2} Q_n \Lambda_s \Lambda_\omega, \quad (2.59)$$

where P is the power of the optical field incident on the mirror, M is the mirror mass, c is the speed of light, L is the cavity length, Q_m is the mechanical quality factor associated with the mechanical mode, and Q_n is the optical Q-factor for the higher order optical mode, which can be described in terms of the linewidth or finesse of the cavity for that mode: $Q_n = \omega_n/2\gamma_n = 2L\mathcal{F}_n/\lambda$.

This equation is based on the Lagrangian analysis presented by Braginsky et al [11], which makes several approximations and assumptions. The moving end mirror is treated as perfectly reflective, while the input mirror is assumed to have infinite mass and no optical loss. The mechanical modes of the mirror are also treated as independent degrees of freedom, and the power incident on the mirror is assumed to be dominated by a single frequency, fundamental mode field. The final equation shown for $R_{n,m}$ is also only valid for cases where the mechanical bandwidth is much less than the bandwidth of the higher order optical mode. This is, however, the case for the suspended mirrors in gravitational wave detectors like LIGO, which are chosen to have very high mechanical Q-factors in order to minimise the effect of thermal noise.

The equation above also considers only the effect of a single higher order mode. Typically the motion of the mirror surface will scatter light into multiple higher order modes, each of which will have a parametric gain, depending on their individual behaviour in the optomechanical system. The overall parametric gain of a mechanical mode m in an optical system is then given by the sum of the individual contributions.

As described by Evans et al in [20], we can generalise equation 2.59 further and consider a generic optical system interacting with a mechanical mode. We consider a feedback process, as in figure 2.10, with the optical system acting as a single plant that has a transfer function G_n for each optical mode n propagating through it. The overall parametric gain of a mechanical mode m is then given in its generic form as:

$$\mathbb{R}_m = \frac{8\pi Q_m P}{M\omega_m^2 c\lambda} \sum_{n=0}^{\infty} \text{Re}[G_n] B_{m,n}^2, \quad (2.60)$$

where the spatial overlap of the mechanical and optical mode in this case is typically denoted $B_{m,n}$, where $B_{m,n}^2 = \Lambda_s$ for each higher order mode n , and $0 \leq B_{n,m} \leq 1$. G_n encapsulates all of the properties of the optical system, for example mirror reflectivities, tunings, and curvatures. Such properties will therefore affect which mechanical modes are likely to become unstable in the system. Since \mathbb{R}_m depends linearly on the incident power, mechanical modes that are marginally stable can become unstable if the input power to the optical system is increased.

2.7 Modelling Radiation Pressure Effects Using Finesse

FINESSE [23, 24] is an advanced frequency-domain interferometer simulation tool. It is developed at the University of Birmingham, led by Andreas Freise and Daniel Brown. It aims to quickly and accurately compute the behaviour of light fields in any arbitrary interferometer layout, particularly gravitational wave detectors. In response to the needs of the GW community, FINESSE has been continuously updated to include features such as realistic beams and optics, quantum noise, and scattering effects [25, 26, 27, 28]. It can also be used in combination with a python wrapper, *PyKat* which allows automation of multiple FINESSE simulations as well as providing a stand-alone optical calculation toolkit.

The basic algorithm of FINESSE computes the steady-state field amplitudes in an interferometer, much like the methods outlined throughout this chapter. This is generalised by constructing a

matrix describing the local light field couplings at each constituent component and how these connect together. The interferometer is described as a nodal network: at each component the coupling between all input and output fields is computed and represented as a matrix which will then form a building block for the overall *interferometer matrix*. The steady-state solution is then calculated, such that the field amplitude at any point in the network can be found, and used to evaluate, for example, the circulating power, an error signal, or a transfer function.

The fields can be calculated using either plane waves or paraxial fields; the latter allows the description of any arbitrary transverse beam geometry through the use of Hermite-Gauss modes (as described in section 2.5). Maps can also be integrated onto optics, describing deviations in phase accumulated across a surface on transmission or reflection, a surface motion, or an aperture. Each higher order mode, and similarly any additional optical frequencies present in the detector (for example for sensing and control of mirror positions) must be computed as a separate optical field propagating through the interferometer, which can make the simulation computationally expensive. However, this has recently been improved, for example using reduced order quadrature [29].

Radiation pressure effects, including optical springs and parametric instabilities, were introduced in FINESSE in late 2014 [30]. This was motivated by the first observation of a PI at LIGO Hanford [10] and the following efforts to understand and mitigate the effect at both LIGO sites and other current and future detectors. I have tested the implementation of both these classical radiation pressure effects in FINESSE using simple examples, to ensure that our results are physically meaningful. This is outlined below. I have then used FINESSE to model parametric instabilities and optical spring effects in systems with much greater complexity, as described throughout this thesis.

As described above, longitudinal optical springs can be considered the ‘fundamental’ optomechanical coupling, while parametric instabilities result from ‘higher order’ couplings. As such, all radiation pressure effects can be implemented in FINESSE using the same generalised functions, with some additions to account for different types of surface motion. These are described in greater depth in chapter 3 of Daniel Brown’s thesis [30] and the Finesse manual [31]; here I provide an overview.

2.7.1 Optomechanical Couplings

FINESSE is designed for modelling linear optical systems in the steady state. As such, optomechanical couplings and the propagation of the resulting optical fields through the system are calculated

in a manner similar to that shown in section 2.4.

A signal can be applied to an optic in a FINESSE model by using the `fsig` command to define an oscillation frequency and a type, for example `Fz` describes a force applied in the beam propagation axis, while `Fs0` specifies a force changing the surface of the mirror. Additional properties of the motion are then described by defining a mechanical susceptibility using `tf` and either attributing this to a degree of freedom of the optic (such as `zmech` for the mechanical motion of the mirror in the beam axis) or pairing it with a map describing the surface motion in the beam axis using the `smotion` command.

All fields incident on the oscillating optic will be phase modulated, introducing new frequency components to the model. FINESSE first runs through the file considering only the carrier fields in the static case, building the nodal network and providing information about the carrier amplitude at all nodes that can be used to compute the power (and associated radiation pressure) where needed. FINESSE then re-runs the file, including any control and signal frequencies as perturbations to the carrier, and solves the interferometer matrix until a steady state solution is found.

Radiation pressure effects are inherently non-linear with regards to the field amplitude, due to their dependence on the optical power. This means that some assumptions must be applied to the system to linearise the behaviours to something FINESSE is designed to handle:

- The system is assumed to be controlled such that any DC radiation pressure force, i.e. a DC offset in the position of the optics, may be ignored.
 - This is equivalent to ignoring the DC force terms in, for example, equation 2.13, and assuming the cavity length L in equations such as 2.9 is independent of the optical power.
- The magnitude of an optic's oscillation is small compared to the wavelength of the optical field. Correspondingly, the amplitudes of sidebands generated by the oscillation are also small compared to the optical field incident on the mirror, and so the radiation pressure force on the mirror is dominated by the beat between the carrier field and sidebands at frequency Ω , while the beating of the sidebands with one another at 2Ω is negligible.
 - This is equivalent to requiring $kx_{RP} \ll 1$, as assumed throughout section 2.4.
- The frequency of the optic's oscillation is small compared to the optical field frequency and any control sideband frequencies, and smaller than half the difference in frequency between input optical fields when multiple carrier fields are present. This is so that any beat frequencies between input field components are high enough that the resulting radiation pressure

effects are negligible (since $\chi(\Omega) \propto 1/\Omega^2$), and so that the upper sideband for one carrier is not also the lower sideband of another.

- This is slightly different to the frequency requirements of the analytical derivation above (introduced when taking the single mode approximation), however it will generally be true since mechanical frequencies of interest are typically of order kHz, while control sidebands are usually MHz.

As with the analytical derivations presented above, this means we can only meaningfully model systems for which the assumptions are valid.

Using the assumption that surface motions are small allows FINESSE to linearise its description of phase modulation due to an oscillating surface, treating it as introducing a pair of sidebands around each input field frequency component ω_j . All optical fields are described in terms of their higher order mode composition; as described in [30] and [31], the amplitude of each HOM in the sidebands generated on reflection is generalised to

$$a_{s,jnm}^{\pm} = irkA_s^{\pm} \sum_{n',m'} a_{c,jn'm'} \iint_{-\infty}^{\infty} u_{n'm'}(x,y) z_s(x,y) e^{ikz_0(x,y)} u_{nm}^*(x,y) dx dy, \quad (2.61)$$

where r is the amplitude reflectivity of the surface, k is the wavenumber of the input field component, $a_{c,jn'm'}$ is the amplitude of the incident field component's $n'm'$ -th HOM, A_s^{\pm} describes the amplitude of the surface motion, and $u_{n'm'}$ and u_{nm} describe the spatial distributions of input and reflected higher order modes (see equation 2.47). The indices ‘c’ and ‘s’ are used to distinguish between the original ‘carrier’ field component and the newly generated ‘sideband’ field. The phase of the reflected field is separated into a term describing any static surface distortions, $z_0(x,y)$ and $z_s(x,y)$, which describes surface motions (resulting in phase modulation). For simpler motions such as linear and angular optical springs, calculation of these couplings can be sped up using analytical solutions. More complex cases like parametric instabilities rely on numerical integration routines.

In cases where optics are suspended and have transfer functions attributed to them, the radiation pressure force at signal frequency Ω is calculated taking into account all of the optical fields interacting with the optic, including the incoming and outgoing beams from both sides and all of the field amplitude components which beat to produce a power term at that frequency. The

radiation pressure force on an optic is generalised to:

$$F_{\text{RP}}(\Omega) = \frac{2}{c}(-P_{1i}(\Omega) - P_{1o}(\Omega) + P_{2i}(\Omega) + P_{2o}(\Omega)), \quad (2.62)$$

where the signs depend on which direction is deemed a positive motion relative to the mirror surfaces, and

$$P(\Omega) = \sum_j \sum_{n,m} (a_{s,jnm}^+ a_{c,jnm}^* + a_{s,jnm}^- a_{c,jnm}) \quad (2.63)$$

describes the power fluctuations at frequency Ω in each optical field, the indices referring to the input (i) and output (o) fields at surfaces 1 and 2 of the optic.

The displacement of the mirror is then found using the mechanical susceptibility for the motion of the mirror in the beam axis, so the generalised form of $F_{\text{RP}}(\Omega)$ (summed with any additional forces at this frequency) can be substituted directly into the first form of equation 2.27, and in equation 2.61 $A_s^+ = x(\Omega)$ and $A_s^- = x^*(\Omega)$ in this case.

Finally, the interferometer matrix can be solved to find the steady state field amplitudes in the optical system and other derived quantities.

2.7.2 Finite Element Modelling of Mechanical Modes

In order to model parametric instabilities for current and future detectors, FINESSE requires inputs for the frequency and Q-factor of the mechanical mode, as well as a map of how the front surface seen by the beam is distorted. These depend entirely on the material and geometry of the mirror, and so the mechanical eigenmodes of optics of interest can be computed using finite-element modelling tools, such as the structural mechanics modules of COMSOL Multiphysics or ANSYS. A 3D model of the optic is produced and then discretised to a set of small mesh elements, whose equations of motion, governed by linear elasticity and the optic's material properties, are solved numerically. The resolution of the mesh will affect the accuracy of the model, and should be adjusted until the solution converges. In some cases it may be preferable to use adaptive mesh refinement to increase the mesh density locally and dynamically with the deformation of the optic. The unforced eigenfrequencies and eigenmodes of the mirror can then be computed, from which the distortion of the optic's surface in the direction along the beam path can be extracted to create a mirror surface distortion map.

Particular care must be taken when considering the normalisation of such maps. Several common

scalings are used in finite element modelling tools, since there is no unique combination of mode shapes, ψ_i and modal masses, \tilde{M}_i forming the solution to the eigenvalue equations. A common choice is *modal mass normalisation*, where the mode shapes are rescaled by the modal mass:

$$\hat{\psi}_i = \frac{\psi_i}{\sqrt{\tilde{M}_i}}, \quad (2.64)$$

which allows the equations of motion for each mode to be written in a simple and familiar form:

$$(\Omega_{m,i}^2 + i \frac{\Omega_{m,i}}{Q_{m,i}} - \Omega^2) \vec{p}_i(\Omega) = \tilde{F}_i(\Omega), \quad (2.65)$$

assuming weak damping, where $\vec{p}_i(\Omega)$ describes the amplitude motion spectrum of mechanical mode i , $\Omega_{m,i}$ is the mode's mechanical resonant frequency, and $\tilde{F}_i(\Omega)$ is called the *modal force*.

Modal forces are calculated in FINESSE using the interferometer matrix elements outlined above, and form the basis of how FINESSE computes parametric gains. Therefore surface maps provided to FINESSE must use modal mass normalisation to produce meaningful results. In COMSOL, this is one of the default options available in the output options of the eigenvalue solver.

2.7.3 Parametric Gain

As described in section 2.6.1, parametric instabilities can be viewed as a feedback loop which depends on the transfer functions of higher order optical modes through the optical system. Calculating the transfer function for each HOM individually and summing the results (as in equation 2.60) is possible using FINESSE, but computationally expensive. Instead, FINESSE extracts \mathbb{R} from a single sparse solution of the inverted interferometer matrix.

The open loop transfer function for a motion which results in an optical signal that propagates through the interferometer and returns to drive that same motion, $T(\Omega) = \frac{p(\Omega)}{\Delta p(\Omega)}$, exists as a diagonal element in the inverted interferometer matrix, and so can be easily extracted at the mechanical resonant frequency of interest. From the frequency-dependent equation of state for the mirror we then find that this transfer function can be used to directly evaluate the parametric gain of the i th mode:

$$\mathbb{R}_i = 1 - \text{Re} \left(\frac{1}{T(\Omega_{m,i})} \right) \quad (2.66)$$

2.7.4 Testing Finesse

FINESSE is based on a series of generic relations describing how the optical fields evolve on interaction with different components. In simple cases, we can test that FINESSE produces physically meaningful results by comparing to analytical equations. In other cases, comparison can instead be made to other numerical models, or, as shown here, directly to experimental results.

Optical Spring Test: Comparison to T. Corbitt *et. al.*, 2007

In their 2007 paper [32], Thomas Corbitt *et al* demonstrate an optical spring setup using a 1 g end mirror in a 2-mirror suspended cavity. As shown in section 2.4, all such optomechanical systems are inherently unstable, either due the real part of $K_{\text{opt}}(\Omega)$ producing an anti-spring, or the imaginary part producing an anti-damping effect. Therefore additional control mechanisms must be used to stabilise the overall system. In this paper, the system is stabilised through the addition of a second input optical field (referred to as the subcarrier), which is frequency shifted from the first input (the carrier) by one FSR of the cavity, using an acousto-optic modulator (AOM). Both fields are therefore resonant in the cavity, but can be independently detuned from resonance. Corbitt *et. al.* showed that each optical field results in an optical spring effect, which can be used to stabilise each other. As a result, different choices of relative detunings for the two input fields can be used to explore all four regimes of (in)stability created by the net damping and spring effects.




Figure 2.11: Optical layout modelled in FINESSE to replicate the setup presented in [32]. A suspended cavity with a heavy (250 g) input mirror and light (1 g) end mirror is illuminated by two optical fields: a carrier field with 1064 nm wavelength, and a subcarrier which is frequency shifted from the carrier by the FSR of the cavity. The total input power is 3W, 5% of which is subcarrier. Each optical field is then individually detuned, producing optical springs in the cavity. These are then probed by amplitude modulating (A.M.) the carrier laser and outputting the amplitude and phase of the end mirror’s resulting motion.

In FINESSE we replicate this experiment using the optical layout shown in figure 2.11, which can be compared to the schematic layout in figure 1 of Corbitt’s paper. All parameters (listed in table 2.2)

are taken directly from Corbitt’s thesis [16] where available. We have assumed the heavy input mirror can be treated as having infinite mass, and we use a plane waves model, so there are no mirror curvature or cavity stability requirements. While Corbitt’s experiment uses an AOM to produce the subcarrier field, here we simply added a second **laser** component via a beamsplitter, and define its frequency as offset from the default value ($f_0 = c/[1064\text{nm}]$) by the free spectral range of the cavity². The experiment also requires a Pound-Drever-Hall error signal to lock the cavity and laser frequency. In FINESSE this is unnecessary since cavities with zero detuning are automatically resonant for the default laser frequency, and we do not model additional noise sources.

Parameter	Value
Total input power	3 W
Proportion of power in subcarrier	5 %
Cavity Length	0.9 m
Transmission of m_1	800 ppm
Transmission of m_2	5 ppm
Mass of m_2	1×10^{-3} kg
Resonant frequency of m_2	172 Hz
Q-factor of m_2	3200

Table 2.2: Parameters from [16] used in our model replicating Thomas Corbitt *et al*’s 2007 optical trap experiment

The mechanical transfer function of the system is then extracted from the model by applying an amplitude modulation to the carrier laser and outputting the resulting displacement of the end mirror using an **xd** detector. This is found for four choices of carrier and subcarrier detunings, which are expressed in terms of the cavity bandwidth full-width-half-maximum (so a subcarrier detuning of δ_S gives a subcarrier frequency of $f_S = f_0 + \text{FSR} + \delta_S \times \text{FWHM}$, while the carrier frequency is $f_C = f_0 + \delta_C \times \text{FWHM}$). The resulting transfer functions are plotted in figure 2.12. Solid lines are results modelled by FINESSE, while the crosses are experimental values extracted from Corbitt’s paper using a plot digitiser. FINESSE shows good agreement with the experimental results.

Parametric Instability Test: Comparison to T. Corbitt et. al., 2006

This test, also presented in Daniel Brown’s thesis [30], allows us to test FINESSE’s evaluation of the parametric gain of a mechanical mode, \mathbb{R} independently from its integration of surface motion maps, since this case considers only longitudinal mirror motions. In chapter 3 I detail how FINESSE has been verified further using the first observation of a PI at LIGO.

²This FSR can either be manually calculated, or found by quickly running the FINESSE file with only the carrier and outputting the cavity properties.




Figure 2.12: Comparison to figure 3 of [32], depicting the mechanical transfer function of the end mirror for four choices of carrier (δ_C) and subcarrier (δ_S) detunings. This showed that an optical spring can be modified (cases (a)–(c)) or stabilised (case (d)) using a second optical field. Solid lines are those produced by FINESSE while the crosses are experimental data values extracted from Corbitt’s figure using a plot digitiser. All parameters values were taken from Corbitt’s paper and chapter 6 of his thesis [16].

In 2006, Corbitt *et al* presented experimental results studying the optomechanical dynamics of a 1 m 2-mirror suspended cavity [19]. This included measurements of both the longitudinal optical spring’s transfer function, and the parametric gain of the 28.188 kHz drumhead mode.

Experimentally, the parametric gain of a mechanical resonance is found by exciting the resonance, setting the cavity tuning to the required test value, and measuring the time constant as the mode rings up or down. In FINESSE, we directly extract the parametric gain as a function of cavity tuning as outlined above.

The optical layout of the FINESSE model is simply a two-mirror cavity illuminated by a single laser input, with parameters taken from [16] and [19], as listed in table 2.3. The ‘drumhead mode’ used in the experiment dominantly results in a longitudinal oscillation, with little coupling to higher order optical modes, so in this case we can model the surface motion using `Fz` and `zmech` and consider couplings to the HG_{00} mode only, rather than applying a surface map. The mechanical Q-factor is not stated in the literature, however as described in [30] a fitted value of 0.75×10^6 is found to give good agreement to the experimental results, and is well within the range of expected

values suggested by the paper authors. The resulting plot in figure 2.13 shows the parametric gain as a function of cavity detuning, where δ_γ is the detuning scaled by the cavity bandwidth (full-width-half-maximum). This shows good agreement between the experimental results, the analytical result presented in the paper, and our FINESSE model.

Parameter	Value
Input power	3 W
Cavity length	1 m
Transmission of m_1	0.63%
Curvature of m_1 (derived)	0.5125 m
Transmission of m_2	0
Curvature of m_2 (derived)	0.5125 m
Effective mass of m_2	0.125 kg
Mechanical mode frequency	28.188 kHz
Mechanical Q-factor (fitted)	0.75×10^6

Table 2.3: Parameters from [16, 19] used in our model replicating Thomas Corbitt *et al*’s 2006 suspended 2-mirror cavity experiment. The mirror curvatures are derived: Corbitt’s thesis describes a cavity stability of $g \sim 0.8$ and mirror curvatures at approximately 0.5 m, leading to the curvature values we have used. The mechanical Q-factor of the mirror is not provided, so is fitted as described in [30].




Figure 2.13: Comparison to figure 4 of [19]. The parametric gain of the system is plotted as a function of the cavity detuning, showing that this mechanical mode is unstable for positive detunings. The solid line is produced using FINESSE, while crosses are experimental data values extracted from Corbitt’s figure using a plot digitiser. The dashed line is the analytical curve which is also extracted from the paper figure. Where available, parameter values were taken from Corbitt’s paper and thesis. This result is also presented in [30].

Chapter 3

Modelling Dual-Recycled Parametric Instabilities at Advanced LIGO

This chapter describes research to explore the behaviour of parametric instabilities at LIGO and how this behaviour is influenced by the core optical configuration of the Advanced LIGO design. This work was reported in *The Influence of Dual Recycling on Parametric Instabilities at Advanced LIGO* [1], which was published in *Classical and Quantum Gravity* in September 2017. I am the principal author, having led the modelling work. Daniel Brown implemented radiation pressure effects and the ‘parametric gain detector’ in FINESSE [30], while Miguel Dovale-Álvarez provided the finite element model of the test mass and associated surface shape maps. Figures and text from the paper have been used verbatim; I have modified the content to suit the structure of this thesis.

This chapter is structured as follows. First in section 3.1 I provide an overview of parametric instabilities at the LIGO sites, briefly reviewing the observations and mitigation techniques implemented to date. Section 3.2 presents an overview of how FINESSE is used to model parametric instabilities (‘PIs’) throughout this work, including details of the Advanced LIGO design that are pertinent to modelling PIs. Sections 3.3 and 3.4 are largely reproduced verbatim from [1]. In section 3.3 we study the parametric gain of specific mechanical modes and how this changes when recycling cavities are introduced, including inherent defects such as astigmatism. Section 3.4 focuses on parameters of the signal-recycling cavity and consequences for the parametric gain in a realistic interferometer configuration. We conclude that changes to the tuning or accumulated Gouy phase

of the signal-recycling cavity have a significant impact on parametric gain, and therefore which modes will require suppression. Section 3.5 describes more recent results which tie into this work. Section 3.5.1 extends from section 3.4.2 to explore whether the combined parameter space of Gouy phase in the Power and Signal-Recycling cavities could be used to minimise PIs in the detector. Section 3.5.2 points out that the same surface motion map, when applied to an input rather than end mirror in the arm cavity, will see a different optical plant due to its location, and therefore has a slightly altered optical gain response. Section 3.5.3 compares the general Advanced LIGO design to a model of the Livingston detector, showing that the results do not substantially change. I conclude the chapter by summarising some suggested areas for further exploration on this topic.

3.1 Observation and Mitigation of Parametric Instabilities at LIGO

A parametric instability was first observed in a gravitational wave detector at the LIGO Livingston Observatory (LLO) in 2015 [10]. A 15.54 kHz mechanical mode was observed to grow, saturating the primary gravitational wave detection output electronics.

Since this first observation, both the Hanford and Livingston detectors have been upgraded; in particular during the commissioning period between the first and second observing runs (referred to as ‘O1’ and ‘O2’), the Hanford site increased the laser input power from 20 to 50 W [33]. This increase in power resulted in Hanford observing many more parametric instabilities. During O2, 13 different mechanical modes, associated variously with all four test masses, were observed at Hanford [22, 34]. Examples of measured mechanical modes, and the modelled arm cavity optical response, are shown in figure 3.1. We are currently in a year-long commissioning break between the second and third observing runs, during which both detectors will again undergo power upgrades, aiming to reach 70 W input power. We can therefore expect to observe more PIs in the future.

If no intervention occurs, PIs can saturate the interferometer control systems, causing the detector to move away from a useful operating point. This reduces the observing time of the detector. The mitigation strategy up to the end of O2 has focused on using technologies already built into the interferometers for other purposes. Livingston has so far suffered from only a few PIs due to its lower circulating power. In this case they can mitigate PIs using the ring heaters on the test masses, which form part of LIGO’s thermal compensation system [2]. Adjusting the ring heaters

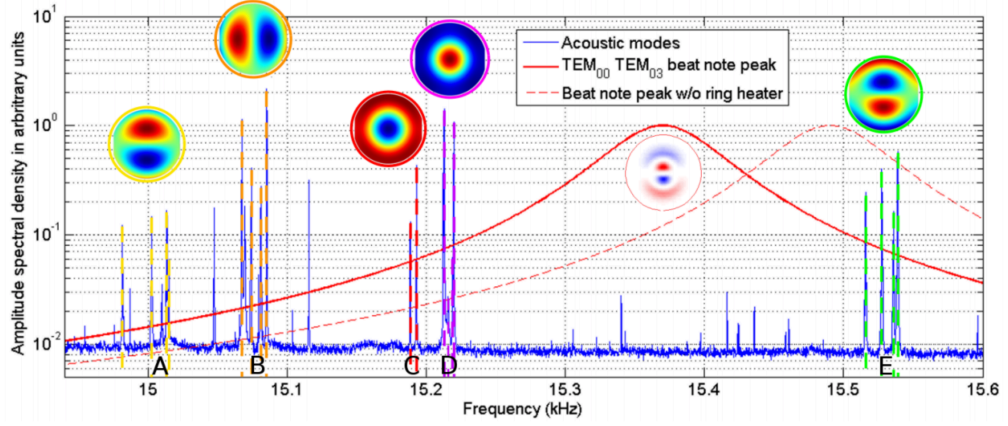


Figure 3.1: Mechanical modes observed during the first Advanced LIGO Observation run. Frequency measurements (in blue) were made by looking at the signal transmitted from the Output Mode Cleaner, and matched to surface shapes by comparison to finite element models. The red traces depict the beat note of the HG00 and HG03 modes in a single arm cavity (i.e. no recycling), when ring heaters are turned on (solid) and off (dashed). The modes are bunched into groups of 4, corresponding to the four test masses. This plot is reproduced from figure 1 of [34].

changes the optical gain of problematic higher order optical modes [35, 36, 37, 38, 39], and move the interferometer into a ‘quiet’ region where no mechanical modes couple strongly enough to the optical field to become unstable. At Hanford, there are too many PIs for this method to work alone. There, electrostatic drivers (ESDs) are used to actively damp the mechanical mode that is unstable [40, 34]. The ESDs are installed on the reaction masses, behind the test masses, to provide low-noise actuation for the control of the detector longitudinal degrees of freedom [41]. The charge across an ESD can also be modulated to provide a force distribution over its surface. This can be designed to counteract the surface motion associated with unstable mechanical modes, actively damping them. Further details of the implementation of ESD damping and PI characterisation methods used at the LIGO sites up to and during O2 are provided in [22].

So far these techniques are proving successful; however as the circulating power is increased towards the design level, the severity and number of PIs will increase resulting in more unstable modes. New mitigation techniques are therefore under development worldwide. The most promising of these currently is a passive technique to directly decrease the Q-factor of a select range of mechanical mode frequencies using ‘acoustic mode dampers’ (AMDs) [42]. A small (mm-scale) reaction mass is attached to the barrel of the test mass via a piezoelectric crystal. The energy of the mechanical mode is then coupled into the crystal and dumped into a resistor, where it is radiated away as heat. This effectively damps the mirror motion and therefore reduces the resulting parametric gain. In March 2018, this was implemented on one end test mass at Livingston [43]. Their experimental

3.1. OBSERVATION AND MITIGATION OF PARAMETRIC INSTABILITIES AT LIGO

results are very positive, and indicate that AMDs could be able suppress all PIs in LIGO up to design power. Although AMDs do increase the thermal noise of the test masses, this was not found to limit the detector sensitivity for LIGO. A proposal to implement AMDs across the LIGO network is therefore currently in development.

Other mitigation techniques have also been explored, including optical techniques such as actively injecting higher order optical modes in anti-phase [44], introducing additional cavities to passively suppress the optical response to higher order optical modes [45], or modulating the higher order optical mode separation frequency [46].

The parametric instabilities observed at LIGO are first generated in the Fabry-Perot arm cavities of the interferometer, where the optical power is highest and directly interacts with the suspended test masses. It is therefore logical to first understand the behaviour of PIs in these cavities. However, the Advanced LIGO design includes dual recycling (see chapter 1), and is therefore a coupled, multi-cavity system. The optical response of the interferometer is thus shaped by the resonant conditions of the recycling cavities as well as the arms. This changes the expected behaviour of PIs in the interferometer, as has been studied both analytically and numerically [20, 47, 48, 49, 50, 51, 52].

The work in this chapter and presented in [1] complements this existing research, aiming to develop a more intuitive understanding for how PIs behave in this multi-cavity system. I use FINESSE to model the full dual-recycled Advanced LIGO design with inherent imperfections. As also shown in [20], introducing dual recycling causes the optical response to change from a single, broad resonant peak, corresponding to resonance in the arm cavity, to a pair of narrower, taller peaks, corresponding to resonance in the full system via the power or signal recycling cavities. Including the astigmatic affects inherent to LIGO's recycling cavities again increases the number of peaks in the optical response. Parameters in the recycling cavities can therefore be used to change which mechanical modes will become unstable, and their corresponding parametric gain. This should be taken into account when considering upgrades to the recycling cavities, and when considering parametric instability mitigation techniques.

3.2 Modelling LIGO Parametric Instabilities

3.2.1 Design of the Core Optics in Advanced LIGO

A complete model of the core optics in Advanced LIGO forms the basis of the simulation, as depicted in figure 3.2. This model uses design parameters given in [2], including using a circulating power of 750 kW. Key frequencies derived from this design are listed in Table 3.1 and full details of the optical parameters can be found in Appendix A.




Figure 3.2: Optical Layout used throughout these simulations, indicating the location of the applied surface motion map. The X- and Y-arm cavities are formed by the Input and End Test Masses (I- and ETMs), and together form a Fabry-Perot Michelson via the beamsplitter (BS). The Power Recycling Cavity (PRC) is formed between the arm cavities and the Power Recycling Mirror (PRM), via Power Recycling cavity mirrors PR2 and PR3. The Signal Recycling Cavity (SRC) is similarly formed between the arm cavities and Signal Recycling Mirror (SRM) via Signal Recycling cavity mirrors SR2 and SR3.

As in the design, the 4 km long X- and Y-arm cavities formed by the Input- and End Test Masses (I- and ETMs) are identical to one another: the ETMs and ITMs have radii of curvature of 2245 m and 1934 m respectively and the cavities are the same length. The path lengths from the X- and Y-arm cavities to the Beam Splitter (BS) differ by 8 cm; this introduces Schnupp Asymmetry which allows the radio frequency control sidebands generated at the input to transmit through to the dark port [53].

Like the general description of ground-based interferometric gravitational wave detectors given in chapter 1, the Advanced LIGO design includes both Power- and Signal-Recycling Cavities (the P-

3.2. MODELLING LIGO PARAMETRIC INSTABILITIES

and SRC'). These are formed by the Power-Recycling Mirror (PRM) and arm cavities, and Signal-Recycling Mirror (SRM) and arm cavities respectively. At LIGO, the recycling cavities are *folded* in order to (a) produce cavities that are mode matched to the arm cavities and geometrically stable, and (b) fit into the space available on site [54]. This comes at the expense of producing a slightly astigmatic beam. Currently, the LIGO detectors operate with the SRM tuned to resonantly extract the signal sidebands (Resonant Sideband Extraction, RSE).

Parameters within this core model, such as mirror positions, angles and curvatures, may then be varied to change the response of the interferometer, and consequently the parametric gain. We have used feedback loops, mimicking those used at the detector sites, to check that these parameter sweeps do not move the model away from the operating point for the interferometer linear degrees of freedom. We refer to in-phase signals that are common to both arms and therefore reflected back towards the laser as ‘common’ signals, while ‘differential’ signals, with a 180° phase difference between arms, are transmitted from the beamsplitter to the detection port.

Cavity	Property	Value
X-arm	f_{msf}	5.155 kHz
	f_{pole}	42.34 Hz
Y-arm	f_{msf}	5.155 kHz
	f_{pole}	42.34 Hz
PRX	x-axis f_{msf}	377.0 kHz
	y-axis f_{msf}	358.3 kHz
	f_{pole}	309.5 kHz
PRY	x-axis f_{msf}	377.5 kHz
	y-axis f_{msf}	359.0 kHz
	f_{pole}	310.0 kHz
SRX	x-axis f_{msf}	288.5 kHz
	y-axis f_{msf}	255.3 kHz
	f_{pole}	420.5 kHz
SRY	x-axis f_{msf}	288.9 kHz
	y-axis f_{msf}	255.6 kHz
	f_{pole}	421.0 kHz

Table 3.1: Key frequencies derived from the Advanced LIGO design model. f_{pole} is the pole frequency, describing the linewidth (half-width-half-maximum) of the cavity. f_{msf} is optical higher order mode separation frequency. In cavities where the beam is astigmatic, the mode separation frequency differs between the x- and y-axes. PRX(Y) and SRX(Y) refer to the cavities formed between the PRM or SRM and ITMX(Y) respectively (see figure 3.2).

3.2.2 Mechanical Modes of an Advanced LIGO Test Mass

The Advanced LIGO test masses, depicted schematically in figure 3.3, are 40 kg fused silica cylinders of radius 17 cm and thickness 20 cm with flattened sides and ‘ears’ for attaching the monolithic




Figure 3.3: Geometry of an Advanced LIGO Test Mass as modelled using COMSOL. The 40 kg mirrors are fused silica cylinders with flattened sides and ‘ears’ which allow monolithic fibre suspensions to be attached.

suspension fibres. In our model we use a Young’s Modulus value of 73.1 MPa and assume the mirror is at room temperature (293 K). Full details of the mirror geometry and material properties used for this modelling are provided in Appendix A.

Surface motion maps for these mirrors were produced using COMSOL as described in section 2.7.2, using the ‘extra fine’ mesh option. A symmetry boundary condition is applied in the xz -plane to speed up the computation time. This resulted in a set of 800 test mechanical modes in the 0-60 kHz range. Example mirror surface shapes are shown in figure 3.4, where listed mode numbers are those generated by COMSOL. Mode 37 has strong spatial overlaps with HG03 and HG21 optical modes and is associated with the first observation of a PI in a LIGO detector [10].

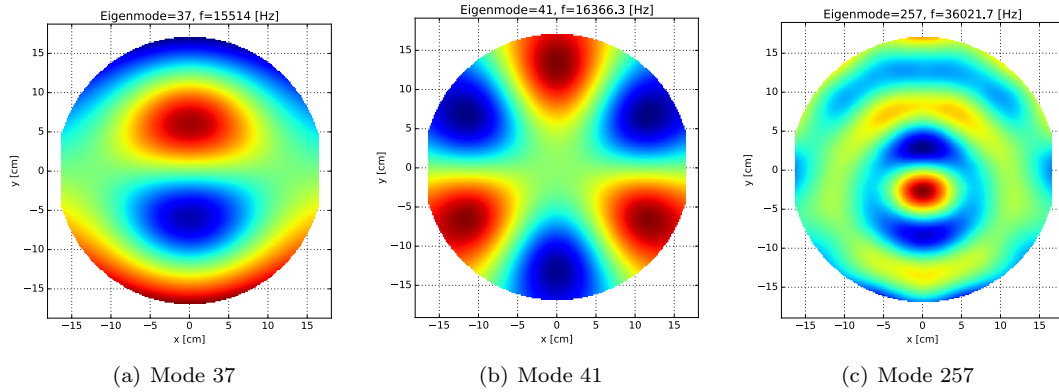


Figure 3.4: Surface motion maps for mechanical modes used throughout this chapter. Mode 37 is associated with the first observed parametric instability at Advanced LIGO. Modes 41 and 257 are examples of mechanical modes that interact with the dual recycled interferometer differently to mode 37, as described in section 3.4.

3.2.3 General Method

As described in section 2.7, FINESSE computes the parametric gain of the m th mechanical mode of a mirror in an interferometer by extracting the open loop transfer function of the mirror motion back onto itself, $T(\omega_m)$, from the inverted interferometer matrix and then using the real part according to equation 2.66. Since FINESSE computes the light field amplitudes of all Hermite-Gauss (HG) modes to a specified order, $T(\omega_m)$ contains components from all these optical fields directly. Unless stated, all plots in this chapter are produced by evaluating higher order optical mode couplings up to 10th order. By default, FINESSE treats all optics as having infinite transverse extent – aperture effects such as loss due to clipping are not considered in this chapter, but will be explored extensively in chapter 4.

Each simulation of a parametric instability in this chapter applies one surface motion map to the End Test Mass of the X-arm (ETMX), as shown in figure 3.2. The simulation also takes the resonant mechanical mode frequency and Q-factor as inputs; by default we use a Q-factor of 10^7 and the resonant frequency computed by COMSOL. This means that we can explore the combined parameter space of mechanical mode frequency and interferometer parameters. Note that since we only study effects due to ETMX, modes from different test masses and any cross-coupling between these are not considered in this study.

3.3 Parametric Instability in Increasingly Complex Interferometers

3.3.1 ‘Common’ and ‘Differential’ Parametric Instabilities

Figure 3.5 depicts the parametric gain of mode 37 (see figure 3.4) as a function of the resonant frequency of the mechanical mode, using the method described in section 3.2.3. We compare a single X-arm cavity to Michelson Interferometers with just Fabry-Perot arms (FPMi), Power-Recycling (PRFPMi), and Dual-Recycling (DRFPMi). We find that the presence of the power- and signal-recycling cavities significantly shapes the optical response and resulting parametric gain, in agreement with [20]. To allow direct comparison, the input power was adjusted to maintain a constant power circulating in the arm(s) in all cases.

In both the single arm cavity and FPMi cases we see a typical single broad peak. This corresponds

3.3. PARAMETRIC INSTABILITY IN INCREASINGLY COMPLEX INTERFEROMETERS

to an overlap between the mechanical mode frequency and the 5.16kHz mode separation frequency of the arm cavity, which allows a 3rd order optical mode to resonate.

Introducing the power-recycling mirror results in a cavity coupling between the X-arm and both the power-recycling cavity (PRC) and Y-arm. The condition for resonance is therefore complicated. We see the introduction of *two* new peaks, since the PRC includes spherically curved mirrors at non-normal incidence, producing an astigmatic beam. This results in the HG03 and HG21 modes picking up different amounts of Gouy phase in the cavity. We describe these peaks as *common* peaks due to their association with the reflection port of the Michelson. The frequency separation between these common peaks and the original single cavity peak is 182 Hz, while the separation between the two common peaks is 36 Hz.

Similarly, adding the signal-recycling mirror produces an additional set of couplings via the signal-recycling cavity (SRC). This time the new resonance condition results in two *differential* peaks, offset from the single cavity resonance by 30 Hz. These two peaks are unresolved due to the low finesse of the SRC, appearing as a broadening of the peak when compared to the non-astigmatic case. We also see that the original broad peak is suppressed.

Our model allows us to treat the resonant frequency of each mechanical mode as a tuneable

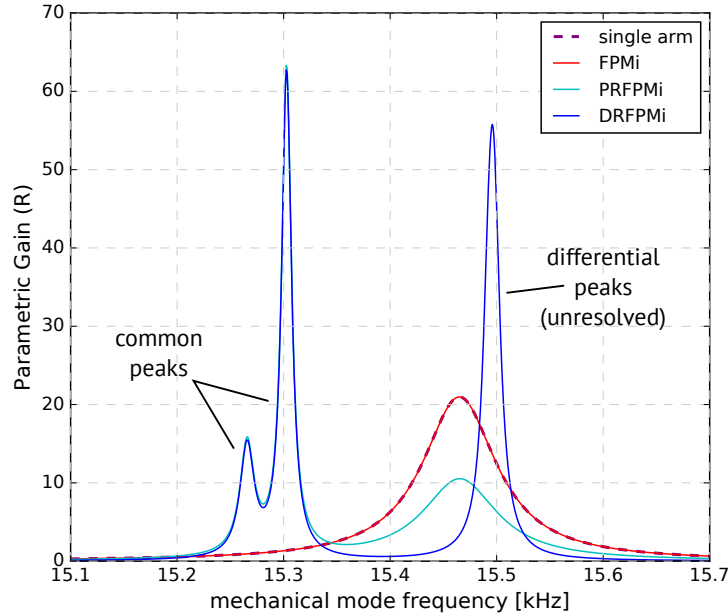


Figure 3.5: Parametric gain of mode 37 as a function of mechanical mode frequency in the case of a single arm cavity, Michelson interferometer with Fabry-Perot arm cavities, Power-Recycled-, and Dual-Recycled Fabry-Perot Michelson. The arm circulating power is 750 kW in all traces.

3.3. PARAMETRIC INSTABILITY IN INCREASINGLY COMPLEX INTERFEROMETERS

parameter, as discussed in section 3.2.3. Changes of frequency on the scales explored here are not something we expect to see in reality, however plots of this kind are useful diagnostic tools. They allow us to explore the response of the interferometer to a mechanical mode, whose resonant frequency may shift and is unknown prior to measurement, without changing the interferometer state.

Experimental work at both the Hanford and Livingston detectors has attributed a mirror motion with the shape of mode 37 to observed parametric instabilities at 15.53kHz. We see that in our model this falls within a differential peak in parametric gain, indicating that properties of the signal-recycling cavity could also be used to influence the gain of this mode in the interferometer, as will be shown in section 3.4. However, other modes will match different resonant conditions in the interferometer, for example resonating via the PRC. Improving the behaviour for one mechanical mode may worsen the situation for another.

The internal properties of the arm cavities can be used to suppress parametric instabilities. Unlike the SRC properties discussed below, changes to the radii of curvature (RoCs) of the test masses are known to influence parametric gain by altering the optical response. They are therefore one focus of efforts to suppress PIs at the detector sites. Figure 3.6 plots the parametric gain of modes 37, 41 and 257 on ETMX when the RoCs of all four test masses are increased simultaneously by the same amount. Once again we find four peaks in the trace for mode 37. We can also see that a





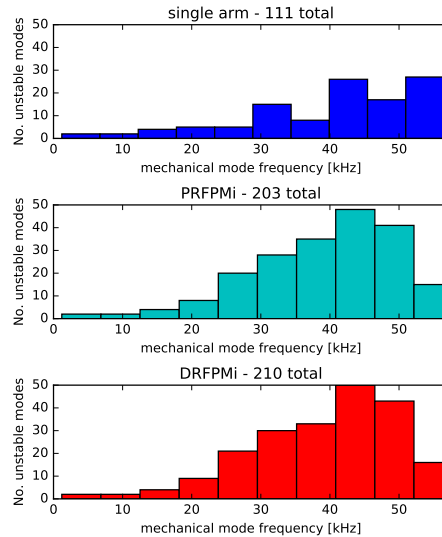
Figure 3.6: The parametric gain of modes 37, 41 and 257 depend on the radii of curvature (RoC) of the test masses. In this plot, the RoC of all four test masses are changed by the same amount simultaneously. For context, the reference curvature of ETMX is 2248m.

significant change in RoC is expected to stabilise mode 37, but overcompensation could result in instability through mode 41.

3.3.2 A ‘Forest of Modes’



(a)



(b)

Figure 3.7: A forest of PIs: each data point relates to a specific mechanical mode that has positive parametric gain within ± 2 kHz of its COMSOL computed resonant frequency. 3 cases are depicted: a single arm cavity, PRFPMi, and DRFPMi interferometers. We find that the DRFPMi configuration could produce twice the number of unstable modes when compared to the single cavity.

Figure 3.7 illustrates the importance of including the full DRFPMi interferometer in PI studies for LIGO. Building the interferometer in stages as described in section 3.3.1, we plot all mechanical modes up to 60 kHz that are found to be unstable within 2 kHz of their COMSOL frequency. Each point then marks the peak value of \mathbb{R} found for each eigenmode. This allows for inaccuracies in our simple mechanical model and a range of interferometer parameters, creating a ‘worst case scenario’ for PIs at LIGO. Critically, we find that the dual-recycled interferometer could suffer from twice the number of PIs when compared to the single cavity case. Note also that this plot refers to PIs exclusively due to mechanical modes in ETMX. Modes from different test masses, and any cross-coupling between these, are not considered in this study. Therefore the total number of PIs could at worst be quadruple that depicted.

3.4 The Signal Recycling Cavity

3.4.1 Tuning

In addition to Advanced LIGO’s current broadband operation using resonant sideband extraction (RSE), the tuning of the signal-recycling cavity can be adjusted to produce an operational mode that is optimised for a particular gravitational wave source [55]. In particular, an SRC detuning of $\phi = 16^\circ$ is proposed for optimal binary neutron star detection [2], where a tuning of 360° corresponds to a mirror displacement of one wavelength. We find that the parametric gain of some mechanical modes has a strong dependence on the tuning of the signal-recycling cavity length.

Figure 3.8 depicts the parametric gain on ETMX for modes 37, 41 and 257 as a function of position of the signal-recycling mirror (SRM), expressed as tuning relative to RSE. Each mechanical mode is modelled at its COMSOL determined frequency (see figure 3.4). Detuning the SRC causes a minor alteration to the operating point of the interferometer (see Section 3.2.3); however, actively tuning the interferometer linear degrees of freedom to maintain the operating point did not significantly change our results.

For mode 37, we find a broad peak in parametric gain, resulting in instability for the nominal tuning and an increase in parametric gain for negative detunings. The SRM and Input Test Mass in the X-arm (ITMX) can be viewed as a compound mirror with a frequency-dependent reflectivity determined by the phase accumulated in the SRC. Changes to the position of the SRM therefore alter the effective reflectivity of ITMX as seen by the higher order optical modes.




Figure 3.8: The parametric gain of modes 37, 41 and 257 (see figure 3.4) depend on the tuning of the signal-recycling cavity relative to RSE.

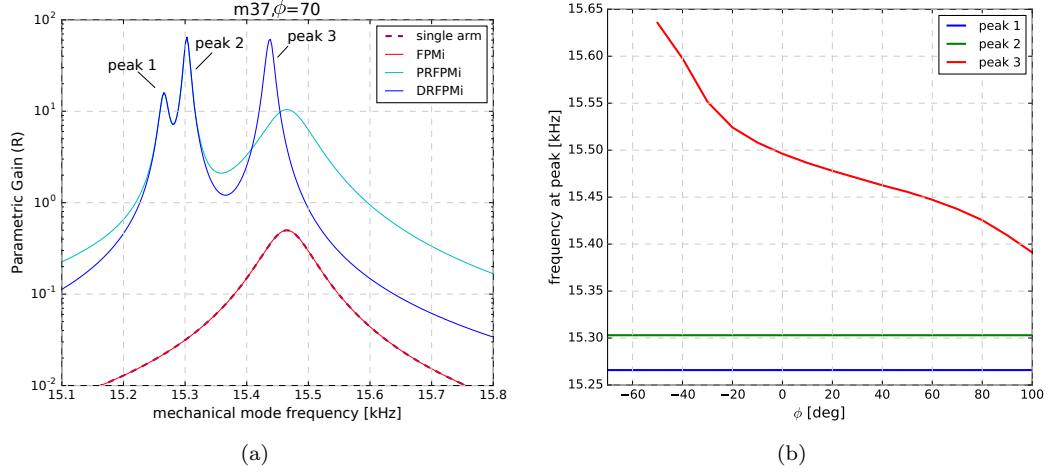


Figure 3.9: Parametric gain of mode 37 as a function of mechanical mode frequency, plotted for different choices of SRM position, expressed as SRC tuning relative to RSE. In 3.9(a) we show the example $\phi = +70^\circ$. On comparison to figure 3.5 we see that the differential peak has shifted in frequency and amplitude, while the common peaks introduced by the PRC are not affected. In 3.9(b) we plot the mechanical mode frequencies resulting in peaks of parametric gain as the SRM position changes. This confirms that only the peak identified as ‘differential’ is affected by SRC tuning.

Figure 3.9 directly tracks the interferometer response to the tuning of the SRC. Figure 3.9(a) plots parametric gain as a function of mechanical mode frequency for the example of $\phi = 70^\circ$, including depicting the power-recycled and single cavity responses for reference. Figure 3.9(b) extracts the mechanical frequency corresponding to each of the three found peaks. We see that the common mode peaks are left unaltered, while the frequency of the differential peak doublet strongly depends on the position of the signal-recycling mirror.

3.4. THE SIGNAL RECYCLING CAVITY

The differential mode peak frequency is not depicted for tunings of -70° and -60° since at these values the peak coincides with the common mode peaks and cannot be resolved. In this sense, the differential mode peak appears suppressed. We note that although mode 37 falls within the differential peak, and is influenced by the SRC, other mechanical modes will fall within the common mode peaks, and thus be independent of parameters in the SRC, as shown in figure 3.9(b).

Figure 3.8 also shows the behaviour of the parametric gain as a function of SRC tuning for modes 41 and 257, both at full design power. For mode 41, we find that the mode is stable for all planned SRC tunings. However, we find that mode 257 could become unstable if the operational mode were to be switched away from RSE to positive detunings. The set of mechanical modes that could result in parametric instability can change depending on the SRC tuning.

3.4.2 Gouy Phase

Proposed upgrade plans for Advanced LIGO include replacing the SRM with another mirror of a different curvature. This would alter the Gouy phase accumulated in the SRC and therefore the optical gain of higher order optical modes in this cavity. We find that changes in the Gouy phase have the same effect as SRC tuning for mode 37.

In figure 3.10, we directly set the value of the Gouy phase accumulated in a single pass through the




Figure 3.10: Parametric gain of mode 37 for changing SRC Gouy phase. 3.10(a) shows the parametric gain of mode 37 as a function of mechanical mode frequency for the case of $\Psi = +10^\circ$. 3.10(b) extracts the peak frequencies at each value of Ψ . As in figure 3.9, we see that only the differential peak is affected by changes to the accumulated Gouy phase in the SRC. The missing points in the trace for $\Psi \sim 50 - 60^\circ$ correspond to values for which the third peak cannot be resolved.

space between the SRM and telescope mirror SR2, Ψ (see figure 3.2). This allows us to mimic the effect of changing the radius of curvature of one of these mirrors without additional design time to re-mode-match the model. As in figure 3.9, we see that only the differential mode doublet is affected by the change, and that the frequency range for which $\mathbb{R} > 1$ changes with Gouy phase, to the point where we can suppress the differential peak of mode 37. Note that the periodicity of this behaviour is double that of the SRC tuning case; this is due to setting the *one-way* rather than round-trip phase.

3.4.3 Consequences for Advanced LIGO Interferometers

We have calculated the parametric gain of 800 mechanical eigenmodes for discrete tunings in $-100^\circ \geq \phi \geq +100^\circ$ and Gouy phases in $0^\circ \geq \Psi \geq +100^\circ$. This is summarised in figure 3.11. In all cases the mechanical modes are modelled at their calculated COMSOL frequency and the power circulating in the X-arm cavity is 750kW. In the case of SRC tuning, we find a minimum of 1 and maximum of 6 unstable modes in our model. The summed \mathbb{R} plots are dominated by the gain of just one or two modes, as can be seen by comparing the shape of the lower plot in figure 3.11(a) to the trace for mode 37 in figure 3.8. We also find that the current tuning of the SRC sits at a local minimum in terms of number of modes, however the gain of this mode is relatively high when compared to other minima in the upper trace: at $\phi \simeq -40^\circ$ we are closer to suppressing all modes. As expected from comparing figures 3.9 and 3.10, this behaviour is also shown for the case of Gouy phase changes.

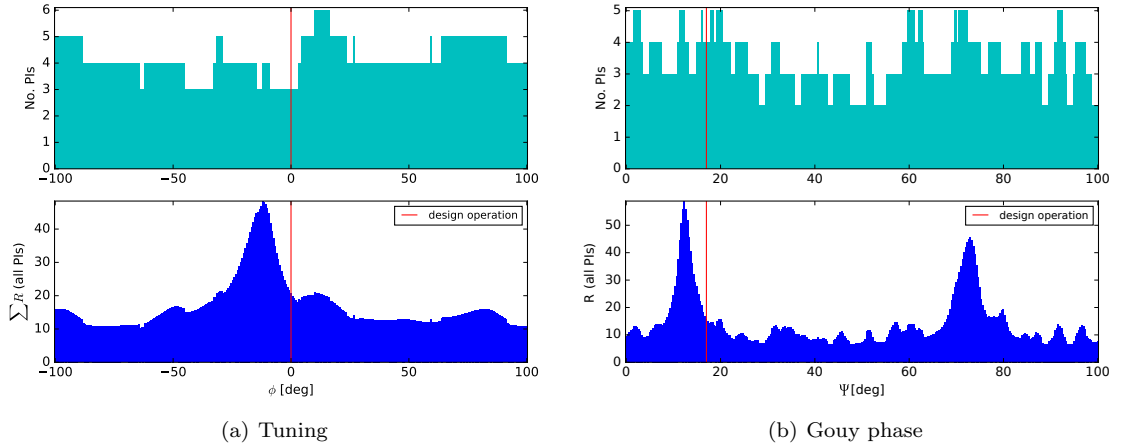


Figure 3.11: Total number of unstable modes and summed parametric gain for different choices of detuning, ϕ (3.11(a)) and Gouy phase, Ψ (3.11(b)). The design operating point of the interferometer is marked for reference.

Relative to the design operating point as show in figure 3.11, the number of unstable modes increases for small positive SRC detuning but the gain of a single PI increases for the equivalent negative detuning. The position and radii of curvature of the mirrors in the signal-recycling cavity alter the parametric gain of mechanical modes and have a strong influence on which PIs will appear. This may have consequences for the current mitigation scheme, such as targeting new mode shapes or requiring stronger actuation to damp higher gain modes. However as seen in our results, the number of additional PIs that could arise in the current configuration are limited.

The SRC ‘Forest’




Figure 3.12: Total number, and mean parametric gain, of modes that are unstable within $\pm 2\text{kHz}$ of their COMSOL frequency. The influence of the SRC is diluted here since this range of mechanical frequencies also allows most modes to resonate via a common rather than differential mode.

The approach used in section 3.3.2 can also be used to study the influence of SRC tuning, as summarised in figure 3.12. However, we find that in this ‘worst case scenario’, whereby all interferometer parameters combine to maximise the number of PIs, the influence of SRC tuning on the total number of PIs is diluted. By allowing the mechanical frequency to sweep over a 4kHz range, peaks in parametric gain due to both power- and signal-recycling are included, and a value of $R > 1$ anywhere in this range is treated as a count of 1 unstable mode. Since the majority of modes are able to resonate in the PRC (given an appropriate choice of mechanical frequency), changing the tuning of the SRC just influences the minority of modes that are *only ever* resonant via the SRC.

Therefore in this case the method used in figure 3.11—using the exact mechanical mode frequency computed by COMSOL—is more informative, and represents a single example interferometer configuration which could be tuned using parameters of the SRC. Conveniently, this method is also faster to compute, since only one parameter is varied to generate each data point. It can therefore be run using a much higher resolution x-axis, as shown.

3.5 Further Work

3.5.1 Is There a Gouy Phase ‘Sweet Spot’?

As shown in sections 3.4.2 and 3.4.3, the number of unstable modes, and the parametric gain of these, depends on the phase accumulated in the Signal-Recycling Cavity. While there is no ideal region with no PIs, there are regions where both of these qualifiers are relatively low. Adjusting the tuning (i.e. propagation phase) of the SRC in order to minimise PI is ill-advised, since this will change the operational mode of the interferometer and therefore the overall sensitivity curve of the detector. However, the Gouy phase accumulated in the recycling cavities could be adjusted without changing this operating point, by changing the radius of curvature of the mirrors. It should be noted that this will change the geometric stability of the cavity, which will have consequences for the control the detector, as explored in [56, 57].




Figure 3.13: Total number of unstable modes and summed parametric gain for different choices of one-way Gouy phase, ψ in the Power- and Signal-Recycling Cavities.

Figure 3.13 explores whether there is a ‘sweet spot’ for Gouy phase in the recycling cavities. This

plot uses the same technique as that used to create figure 3.11, but now explores the 2-dimensional parameter space formed by the Gouy phase accumulated in both the SRC, ψ_S and PRC, ψ_P . Each horizontal slice of the 2D plots forms another version of figure 3.11(b) with a different value of PRC 1-way Gouy phase, counting the number of modes that are unstable at their COMSOL-computed mechanical resonant frequency, and summing the total gain of these unstable modes. As usual, we include calculation of all couplings to higher order Hermite-Gauss optical modes up to 10th order. The ψ_S and ψ_P labels on the axes correspond to the values of Gouy phase applied directly to the space between the S(P)RM and S(P)R2 (see figure 3.2). This is a single, symmetric value applied to Gouy phase in both the x- and y-axes, meaning that the astigmatic effects created in the recycling cavities are removed in this simulation.

Rather than a unique ‘sweet spot’, figure 3.13 shows that there are many small regions with few, or zero, PIs. The Gouy phase of the recycling cavities for the Advanced LIGO design was determined based on the geometric stability of the cavities and the need for good alignment signals from these cavities. Compromising between these two led to chosen values of $\psi_S \sim 19^\circ, \psi_P \sim 25^\circ$ [2], as indicated in figure 3.11(b). Figure 3.13 shows that the effect of recycling cavity Gouy phase on PIs should also be used when the same decision is being made for future detectors. For the case of our Advanced LIGO model, which is dependent on the exact set of optical parameters used and set of mechanical modes tried, this plot indicates that values of $\psi_S \sim 50^\circ$ and $\psi_P \sim 40^\circ$ may experience fewer PIs than the current design.

3.5.2 ITM vs. ETM

As noted in section 3.2.3, the simulations above only explore the behaviour of a single mechanical mode of the End Test Mass in the X-arm of the Advanced LIGO interferometer model. All four test masses are suspended in the same manner, have the same general geometry, and interact directly with the same optical field in the arm cavities. Therefore in the Advanced LIGO model we expect them to produce PIs for the same surface motion maps. In reality, the four test masses are not completely identical, but are extremely similar, meaning that the same type of surface motion has a slightly different resonant frequency for each test mass (typically a few Hz, see figure 3.1). Based on this, we can expect that the estimated number of stable modes is four times the numbers predicted by plots such as figure 3.11.

Figure 3.14 introduces one caveat to this, however: the optical response for a surface motion map

3.5. FURTHER WORK

applied to an ITM is not the same as that for an ETM. In the figure, we plot the parametric gain of mode 37 as a function of mechanical mode frequency as usual. This time, however, we apply this same map to all four arm cavity mirrors in turn.

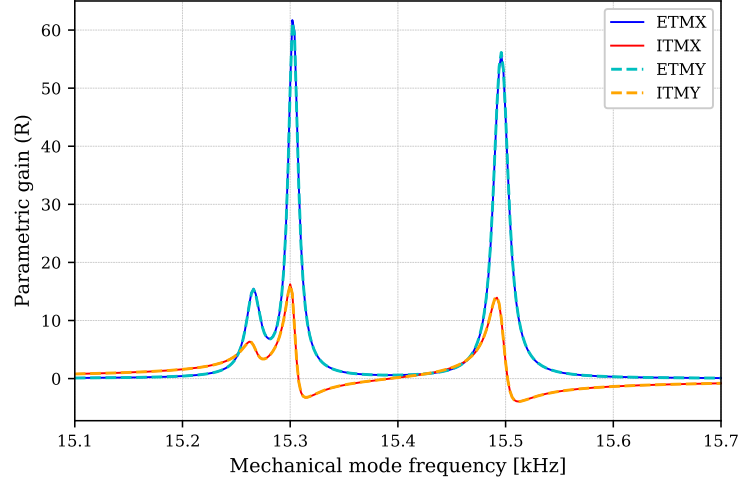


Figure 3.14: Parametric gain of mode 37 when the surface motion occurs at each of the Input- and End Test Masses in the X- and Y-arms.

The dominant force driving the motion of the optic is the optical field in the arm cavity, since this is enhanced by the cavity finesse and therefore the power is a factor of several hundred higher than the field in the PRC. However, while an ETM only experiences radiation pressure from the arm cavity directly, ITMs interact with light from both inside and outside the cavity, and the combined radiation pressure forces from both sides will therefore affect the total parametric gain. The change in optical response for input versus end mirrors comes from the location of the mirrors and the optical fields that are therefore incident on them.

In reality this interaction would be more complicated still, since the test masses are bulky, composite optics, and so the optical field can interact with multiple surfaces of the mirror. In our FINESSE model, only the highly reflective surface of the optic has a defined surface motion which interacts with the optical field on both sides. The bulk material and anti-reflective surfaces act only as a simple low-reflectivity plane and a change in the optical path length for existing fields, without introducing additional radiation-pressure driven interactions.

Further investigation is needed to verify this explanation for the difference between ITM and ETM behaviour, for example by building a minimal FINESSE model of a power-recycled arm cavity using thin optics and exploring how the parametric gain of a mode applied to the ITM or ETM varies with power recycling gain for a given arm power.

3.5.3 Modelling PIs using parameters from the Livingston detector

The study above explores the influence of recycling on parametric instabilities in an optical model of the Advanced LIGO design. The optical configurations of the two sites differ slightly from this design (and each other) in small but important ways—for example, the radii of curvatures of the test masses are not all identical, meaning that the arm cavities do not have identical beam parameters and there is a small amount of mode mismatch at the beamsplitter.

Studies [48, 20] have shown that arm mismatch will influence parametric gain—since, as illustrated by figure 3.6, changing the radii of curvature of the test masses will change the optical response—but that this effect is not expected to significantly change the number of PIs that the detector experiences.




Figure 3.15: Parametric gain of mode 37 in a model of the LIGO Livingston detector (LLO) compared to the Advanced LIGO (aLIGO) design, shown for the case of a single arm cavity (‘1 arm’) and the full DRFPMi configuration (DR).

Figure 3.15 agrees with this statement: here, we plot the parametric gain of mode 37 as a function of mechanical mode frequency, comparing the results for the Advanced LIGO design file (as seen in figure 3.5) to a model of the LIGO Livingston detector, ‘LLO’. This model [58] uses measured values for the optics [59] and setup.

The traces for the Advanced LIGO and LLO are not significantly different in form. The response of the single arm cavity is shifted up by ~ 0.2 kHz. In the dual-recycled (DR) configuration, this means that the common mode peak is further away from the arm cavity envelope and becomes smaller, while the differential peak is amplified. However, while this shows that mismatched arms will change the absolute behaviour of parametric instabilities in the interferometer, it does not

fundamentally change the results presented above.

3.6 Conclusions

Parametric instabilities have been observed at the LIGO sites since 2015, when the detectors came back online after a major period of upgrades towards the Advanced LIGO design. Initially, it was feared that these could become a show-stopping phenomenon, since PIs can saturate the detector control scheme. An intense period of research followed both at the sites and across the globe, building on theoretical and experimental studies developed since 2001, to better understand this behaviour and develop mitigation schemes for PIs. This has proved largely successful so far, meaning that the Hanford site has been able to double its input laser power without detrimentally degrading their duty-cycle. Further mitigation schemes are in development which, it is hoped, will allow the detectors to be run at full design power without experiencing PIs. For the most part, these schemes aim to damp the mirror motion directly, or change the optical response of the arm cavities.

We have used FINESSE to study parametric instabilities in the context of the full Advanced LIGO design. This work provides a route towards the characterisation of PIs in the LIGO detectors as they reach design power, in a way that is not possible with more generalised models. By matching the parameters of our model to measured values at the detector sites, the model could then be used to find the optimal operating conditions.

The dual-recycled configuration of the interferometer greatly expands the parameter space that determines the resulting optical gain of the system. In particular, we have shown that parameters outside the Fabry-Perot arm cavities can also affect the parametric gain of a mechanical mode, to the extent that the number and gain of unstable modes may change.

By contrasting figures 3.6 and 3.8, the complexity of the picture is clear: while the instability of mode 41 is minimally affected by SRC tuning and strongly affected by changes in RoC, the reverse appears true of mode 257, and mode 37 is affected by both. The list of important optical parameters is therefore extensive, and all will influence the likely number of PIs that will affect gravitational wave detectors as the operating power increases.

For the parameters in the Advanced LIGO design, we find that the tuning and Gouy phase accumulated in the signal-recycling cavity will influence the total number of PIs, and the gain of these

modes. For differential modes parametric instability depends on properties in the SRC, while for common modes instability depends on the PRC. Therefore if parameters in the SRC are to be changed, a PI mitigation scheme based on per-mode damping is expected to remain effective for common mode PIs, but may require changes for differential modes. These results are not significantly altered when the parameters are adjusted to match those of the LIGO sites or when the surface motion occurs at a different test mass.

In the combined parameter space of Gouy phase accumulated in both the power- and signal-recycling cavities shown in figure 3.13, we see that there are regions that are expected to result in more or fewer parametric instabilities than the Advanced LIGO design. Choosing such a low-PI region could help ease the requirements for mitigation schemes. We therefore recommend that parametric instability be included as a consideration when designing the recycling cavities of future gravitational wave detectors, such as the Einstein Telescope [12], LIGO Voyager or Cosmic Explorer [60], which will use higher power and larger optics than LIGO.

3.6.1 Suggested areas for further follow-up

The research topics considered in this chapter show the importance of considering the whole detector when attempting to understand parametric instability. There are many possible avenues for further exploration. In chapter 4, I build on the work presented here to study how the next generation of detectors can be designed with PIs in mind. In addition to this there are several opportunities for useful follow-up work directly resulting from the work presented here, of which I highlight two examples.

First, we can experimentally verify the results of our simulations that dual recycling changes the optical response for a PI in a significant way. This could be done, for example, by producing a version of figure 3.6 or figure 3.8 using up-to-date parameters from one of the sites, and comparing the results to measured ring-down values such as those shown in figure 5.28 of [22].

Secondly, those dealing with PI mitigation at the sites regularly experience behaviours that are yet to be fully explained. For example, the error signals sent to the electrostatic drivers at Hanford have sometimes required a 180° phase flip to be arbitrarily introduced when the mechanical mode is associated with an ITM. This naturally follows on from the investigation in section 3.5.2. In this case, modelling the system could help understand the mechanism for the phase flip and allow it to be anticipated or stabilised.

Chapter 4

Parametric Instabilities in the Einstein Telescope

This chapter outlines my study of parametric instabilities in a proposed third-generation gravitational wave detector called the Einstein Telescope (ET). This is a European proposal, which aims to build a gravitational wave observatory with a peak sensitivity that is 10 times better than the current (second) generation of detectors. This is expected to increase event detection rates significantly so that gravitational wave astronomy can become as routine as astronomy using electromagnetic observatories, meaning that events such as the multi-messenger detection of a binary neutron star system in 2017 [9, 61] may become commonplace and used for astronomical surveys and more.

In order to achieve this goal, ET requires many technologies to be pushed forward or newly developed. In 2011, an extensive design study [12] was produced, based on the current knowledge and predictions about the state of the field. This proposes a ‘xylophone configuration’ of three pairs of dual-recycled Michelson interferometers, referred to as ET-HF and ET-LF, with frequency-dependent squeezing, nested together to form an equilateral triangle of detectors underground. Since then, current detector facilities have continued to develop towards their target sensitivities, and encountered unforeseen challenges. Plans for other third generation detectors like Cosmic Explorer have also advanced. Key decisions for the Einstein Telescope will be made in the near future, therefore the ET design is now undergoing revisions and updates, taking into account the growth in technical knowledge, various advances in technologies, and the expectation that ET will

be part of a wider network. The ET facility is expected to have a lifespan of many decades, and therefore the detector design is expected to evolve with time. Recently, a proposal has been made to begin by implementing a single broadband interferometer rather than a pair. This simpler configuration, dubbed ‘ET-120K’, would result in a sensitive bandwidth close to the xylophone configuration, which could then be adapted over time to reach the ET facility’s sensitivity limit. Later configurations could then take the form of an updated xylophone design, or more diverse topologies such as a Sagnac interferometer.

At the time of the ET design study, parametric instability was known of, but yet to be observed. While PIs had been anticipated as a likely phenomenon for interferometric gravitational wave detectors since they were first predicted by Vladimir Braginsky in 2001 [11], this did not attract significant attention until the first observation of a three-mode optomechanical instability was reported in March 2015 [62], and first observation of a PI in a gravitational wave detector announced shortly afterwards in April the same year [10]. As such, the ET design study’s consideration of PIs is preliminary, and restricted to considering PIs in a model of an ET Fabry-Perot arm cavity. It concludes that due to ET-LF’s low power, PIs are not anticipated to be a problem there. However, for ET-HF the significant circulating optical power and size of the mirrors is expected to result in around 7 times as many unstable modes as seen in LIGO.

Since then, several aspects of PI at ET have been explored. In [63] and [64] Sergey Strigin compares the use of cryogenically cooled silicon or sapphire test masses at ET to the room-temperature fused silica test masses at LIGO, showing that these material choices can be expected to increase the number of PIs. In [65], he shows that tuning the radius of curvature of the test masses can be used to reduce the number of PIs in an ET-like Fabry-Perot cavity. In [66], Zhang Jue *et. al.* study PI effects using cryogenically cooled silicon test masses, as proposed for several third-generation detector designs, with results that are consistent with [64].

In this chapter I numerically model PIs in several potential ET design configurations. As in chapter 3, I concentrate on the optical response of the interferometer to mechanical vibrational modes applied to a single end test mass, and study how this can be shaped by the optical parameters of the detector. In section 4.1, I describe the development of the Einstein Telescope design, providing overviews of the optical configuration proposed in the 2011 design study (ET-HF and ET-LF) and the recent ‘ET-120K’ design suggestion. In section 4.2 I compare the test mass designs proposed for ET-HF and ET-120K, using finite element modelling to explore the mechanical modes expected. In section 4.3, I explain how clipping of the beam limits the coupling of higher order optical modes

one can expect to see, and use this to speed up modelling of PIs in FINESSE. In sections 4.4 to 4.6 I consider 3 design options for ET-120K. First, I compare the number of PIs expected in a single 10km arm cavity in the case that clipping loss is small to the case where clipping loss is higher in order to reduce thermal noise. Second, I study these two cases in the full dual-recycled interferometer and compare both to LIGO. I then consider an alternate set of parameters designed to optimise the detector sensitivity in the 1-4 kHz band. Lastly in section 4.7 I propose a plan for designing third generation gravitational wave detectors with parametric instabilities in mind, learning from our experiences at LIGO.

4.1 Evolution of the Einstein Telescope Design

4.1.1 2011 Design Study: The Xylophone Configuration




Figure 4.1: Proposed optical layout for the Einstein Telescope, reproduced from the 2011 Design Study [12]. The observatory is formed from a set of 3 gravitational wave detectors, which are nested together to form a 10 km equilateral triangle underground. Each detector consists of 2 interferometers, tuned to be sensitive to high (Hz-kHz) and low (mHz-Hz) frequency gravitational waves. The outputs are combined to produce a broadband, high-sensitivity observatory – the so-called ‘xylophone configuration’.

The 2011 Einstein gravitational wave Telescope Conceptual Design Study outlines the science case,

site selection and suspension requirements, and optical design plans for ET. The optical design, like the current generation of detectors, would be based on dual-recycled Michelson interferometers with Fabry-Perot arm cavities, however the observatory would be formed from 3 identical, nested detectors forming an equilateral triangle, as shown in figure 4.1. The 60° angle between the arms means that each detector would be equally sensitive to both polarisations of gravitational waves, while co-locating three detectors in this manner would allow the Einstein Telescope to localise events. Each detector would consist of *two* interferometers, designed to be sensitive to different signal frequencies (referred to as ET-HF and ET-LF for the high- and low-frequency interferometers respectively). To reduce seismic noise, the observatory would be 100-200 m underground. To reduce radiation pressure noise, the mirror masses would be increased by a factor of 50. At high frequencies, sensitivity is shot noise limited, so a high circulating power of 3 MW was proposed for ET-HF. A higher order LG_{33} optical mode for the carrier field was also proposed, to reduce the resulting coating thermal noise. In ET-LF, thermal noise would be reduced by cryogenically cooling the optics, necessitating a change in the material of the optics and therefore a change in laser wavelength. Both the ET-HF and ET-LF interferometers would employ frequency-dependent squeezing to overcome the Standard Quantum Limit set by the combined effects of quantum radiation pressure and shot noises.

Research has been underway worldwide to develop and test all these technologies, both for ET and other future gravitational wave detectors like LIGO Voyager and Cosmic Explorer. The sensitivity curves shown in figure 4.2 are based on the parameters from the design study, as listed in table 4.1, predicting that the xylophone configuration for ET would achieve an order of magnitude improvement in sensitivity when compared to LIGO, and widen the peak detector bandwidth by $\sim 50\%$. ET-LF would be expected to be limited by Newtonian noise below around 10 Hz, and seismic noise below 2-3 Hz, while ET-HF would be limited by suspension thermal noise below ~ 30 Hz and mirror thermal noise below ~ 250 Hz. At high frequencies both detectors would be quantum noise limited.

4.1.2 The ET-120K Design

At the 2017 Einstein Telescope Design Update Workshop, plans for ET underwent review, with the aim of both renewing interest in the project and updating the design itself to reflect progress in the field post-detection. There, it was suggested that the xylophone configuration may not be necessary to achieve ET's initial science goals. The xylophone configuration is a means of avoiding the conflicting design requirements necessary to achieve good sensitivity at both high and low




Figure 4.2: Sensitivity curves for the High-Frequency and Low-Frequency interferometers proposed in the 2011 Einstein Telescope Design Study, reproduced from Figure 7 of [12]. These include projected noise contributions from all known sources. The ET-HF (red) and ET-LF (blue) curves together give the final curve for ET, which has an order of magnitude overall sensitivity improvement when compared to Advanced LIGO. The total broadband sensitivity of ET is also significantly wider than that of the Advanced LIGO design. The quantum-noise-limited sensitivity curve of the Advanced LIGO design is shown for reference.

gravitational wave frequencies: for example high laser power is required to improve the detector sensitivity at high frequencies, but this will increase the thermal noise of the test masses. Using a xylophone configuration means that a lower optical power and cryogenic cooling can be used at lower frequencies where thermal noise is limiting, while at high frequencies, where the detector is quantum-limited, a high optical power can be used.

The development of the LIGO Voyager design [67], a proposal to upgrade the current LIGO facilities including cryogenically cooling the test masses to 123 K, has shown that some of these compromises are less constraining than initially believed. As such, the idea emerged that the Voyager design could be adapted to an ET-scale (10 km) facility, resulting in a interferometer with sensitivity similar to the xylophone design. The resulting design is therefore dubbed ‘ET-Voyager’, or ‘ET-120K’.

The ET-120K design is still very much in development. Table 4.1 lists suggested parameters for the core interferometer in this new design, as well as the corresponding values for both Voyager, and ET-HF and -LF. Figure 4.3 depicts the resulting quantum-limited sensitivity curve, alongside the equivalent curves for the xylophone configuration. ET-120K can be viewed as a ‘warmer ET-LF’ or ‘long Voyager’. The optical layout is thus largely identical to ET-LF. The circulating power of 3 MW is higher than that proposed for ET-LF, and is designed to improve ET-120K’s sensitivity

at higher frequencies.

Like LIGO, ET-120K will be based on the dual-recycled Michelson with Fabry-Perot arms (DRF-PMi) described in chapter 1. Figure 4.4 shows a schematic of the proposed layout of the core optics. Assuming that ET-120K retains the triangular configuration, the two arms of each interferometer will be angled at 60° to each other rather than the 90° depicted here, and additional folding may be required to fit the input and output optics into the site (see figure 4.1 for how this may work). However, for the purpose of this study we consider the simplified layout depicted in figure 4.4.

ET uses a similar naming convention to LIGO. The Fabry-Perot X- and Y-arm cavities are identical and 10 km long, formed by the Input- and End Mirrors (I- and ETMs). The radius of curvature for the I- and ETMs is designed so that the beam spot size on the mirrors is as large as possible, which helps reduce coating thermal noise. The upper limit to the spot size is set primarily by the maximum mirror substrate sizes anticipated to be available, and then by the geometric stability of the arm cavities. Due to the 10 km arm length, I have taken the recycling cavity design in my model from the original ET study, while the mirror transmissions, losses and materials are those proposed for Voyager. Test mass curvatures are based on the ET study, however with adjustments due to the projected clipping losses in the original design (see section 4.4). Since ET-120K is proposed to make use of technologies developed for LIGO Voyager, I assume a laser wavelength of $2\mu\text{m}$ and identical test mass properties.

Unlike LIGO, where curved folding mirrors are used in the recycling cavities to ensure that they




Figure 4.3: Quantum-noise-limited sensitivity curves for a possible ‘ET-120K’ design (black), compared to the quantum-limited curves of ET-HF and ET-LF. All curves are the result of plane-waves models and computed using FINESSE; models of ET-HF and -LF were provided by V. Adya.

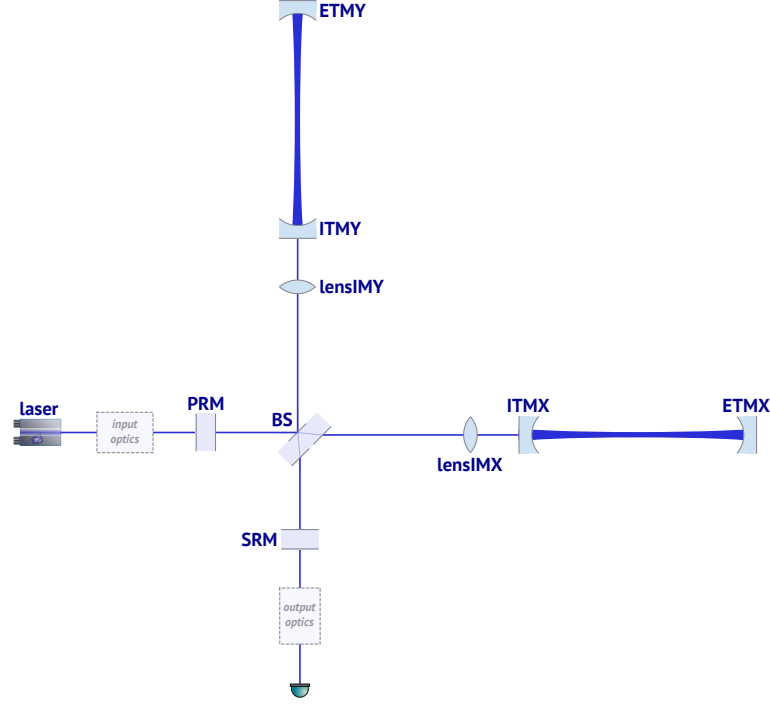


Figure 4.4: Optical layout of the ET-120K interferometer design, depicting the core optics only. While the core design is again a dual-recycled Michelson with Fabry-Perot arms, ET plans to incorporate focusing elements (‘lensIMX/Y’) directly into the input test masses rather than including telescopes in the recycling cavities, as used at LIGO.

are geometrically stable, the arm and recycling cavity lengths in ET can be designed to ensure geometric stability. ET will still require focusing elements, since without these the beamsplitter substrate would need to be significantly larger than would be available. The design study proposes geometrically symmetric recycling cavities, and incorporating focussing lenses into the ITMs, which are then further from the beamsplitter compared to the aLIGO design. This would allow the beamsplitter and recycling mirrors to be significantly smaller than the test masses. This means that the power and signal recycling cavities, formed by the Power- and Signal-Recycling Mirrors (P- and SRMs) and the arm cavities respectively, are linear except for the central beamsplitter, and do not introduce astigmatism to the beam. However, this does reduce the number of pick-off points that could be used for monitoring control signals from the interferometer.

ET-120K is also expected to use frequency-dependant squeezing (‘FDS’) [68, 69, 70, 71]. Squeezers are currently being installed at LIGO; full deployment of FDS is expected to increase the BNS range by 30% [72]. It may also employ balanced homodyne readout [73, 74, 75]. This is currently under investigation for the LIGO A+ design [60] and LIGO Voyager due to the reduction in readout noise compared to DC readout, and tuneable readout angle. In the long-term, it is foreseen that

4.1. EVOLUTION OF THE EINSTEIN TELESCOPE DESIGN

ET could be further upgraded to incorporate elements such as optomechanical filters (see chapter 5 and [13]) and other technologies designed to optimise the quantum-limited detector response. In this study of parametric instabilities we focus on the core optics of the DRFPMi. The behaviour here should be understood, and the design modified to minimise PIs if necessary, before considering the effects of further input and output optics on PIs.

	LIGO Voyager	ET-120K	ET-LF	ET-HF
Arm Power	2.82 MW	3 MW	18 kW	3 MW
Laser wavelength	2 μm	2 μm	1550 nm	1064 nm
Beam Shape	HG_{00}	HG_{00}	HG_{00}	LG_{33}
Readout Scheme	Homodyne	Homodyne	DC	Homodyne
Additional QN elements	FDS: 1 x 300 m FC	FDS -	FDS: 2 x 10 km FCs	FDS: 1 x 300 m FC
Arm Cavities				
Arm Length	4 km	10 km	10 km	10 km
Temperature	123 K	120 K	10 K	290 K
TM RoC	-	6690 m*	5580 m	5690 m
- ITM	1798 m	-	-	-
- ETM	2492 m	-	-	-
ITM Transmission	3000 ppm	3000 ppm*	7000 ppm	7000 ppm
ETM Transmission	5 ppm	5 ppm	6 ppm	6 ppm
TM Scatter loss per surface	10 ppm	10 ppm	37.5 ppm	37.5 ppm
TM Material	Silicon	Silicon	Silicon	Fused Silica
TM Mass	200 kg	200 kg	211 kg	200 kg
TM diameter	~ 45 cm	45 cm	> 45 cm	60 cm
TM thickness	~ 55 cm	55 cm	~ 50 cm	30 cm
Central Interferometer				
Optic material	Fused Silica	Fused Silica	Fused Silica	Fused Silica
Recycling Cavity Length		310 m*	310 m	310 m
- Length ITM-BS		300 m*	300 m	300 m
- Focusing element	-	in/near ITM; f = 303 m	in/near ITM; f = 303 m	in/near ITM; f = 303 m
Temperature	room temp.	room temp.	room temp.	room temp.
PRM Transmission	3.3%	3.3%	4.6 %	4.6 %
SRM Transmission	2.7%	2.7%*	20 %	10 %
SRC Tuning	RSE	RSE	detuned	RSE

Table 4.1: Comparison: parameters of the core optics in my ET-120K design files versus the earlier ET xylophone configuration [12] and proposed LIGO Voyager [67] designs. All acronyms are as listed in the glossary. Starred (*) values in the ET-120K design indicate parameters whose values I explore in this study. While the eventual design for ET-120K is likely to use Homodyne readout and Frequency-Dependent Squeezing, my model only considers the core optics of the interferometer. Sensitivity curves are therefore shown with DC readout and no squeezing.

4.2 Finite Element Modelling of ET Test Masses

The parameters of the test masses proposed for both ET-HF and ET-120K are listed in table 4.2. Compared to LIGO test masses (also listed), they will be significantly heavier, to reduce the effect of quantum radiation pressure noise, and in the case of ET-120K made from crystalline silicon rather than fused silica, so that they can be cryogenically cooled to reduce thermal noise¹. At around 18 K and 125 K, the coefficient of thermal expansion for crystalline silicon has zero-crossings, resulting in a null in thermo-elastic noise [12]. However, silicon is opaque to 1064 nm light, becoming transparent at longer wavelengths such as 1550 nm or 2 μm , thus the use of silicon test masses also results in a shift in wavelength for the interferometer.

	aLIGO	ET-120K	ET-HF [12]
Bulk Material	Fused Silica	Crystalline Silicon [67]	Fused Silica
Temperature, T	room temp.	123 K [67]	room temp.
Density, ρ	2203 kg/m ³	2330 kg/m ³ [12]	2203 kg/m ³
Poisson's ratio, ν	0.17	0.22 [12]	0.17
Young's modulus, E	72.6 GPa	156 GPa [67]	72.6 GPa
Radius, r_1	17.00 cm	22.5 cm [67]	31 cm
Thickness, t	20.00 cm	55 cm [67]	30 cm
Half-width, r_2	16.325 cm	-	-
Height of flats, h	9.49 cm	-	-

Table 4.2: Mechanical properties of mirror substrates for two ET designs, compared to Advanced LIGO.

The test masses for ET are much larger than those at LIGO. Due to expected limitations in fabrication techniques, the diameter of a silicon test mass (TM) is restricted to about 50 cm. To achieve the desired ~ 200 kg total mass, the ET-120K TMs are therefore 55 cm thick – significantly longer than the ratio used for LIGO mirrors. The resulting geometry is shown in figure 4.5. Since ET-HF instead uses fused silica, which is not anticipated to have the same fabrication restrictions, the ET-HF TMs use the same radius-to-thickness ratio as LIGO.

The geometry of the ET test masses is currently modelled as a simple cylinder. In practice, flats will need to be added to the sides to act as attachment points for the suspension fibres. FE models should be updated to include these, and any other geometrical features, as these are decided.

For ET-120K (and, indeed, ET-LF) *crystalline* silicon will be used for the bulk substrate of the mirror. As a result, the Young's Modulus of the material depends on the orientation of the crystal, with different values for each plane, varying from 130 to 188 GPa [76]. In our COMSOL model we simplify our analysis by treating the material as *amorphous* silicon with an estimated value of

¹The mechanical loss of fused silica increases at low temperatures, making it unsuitable for cryogenic operation.

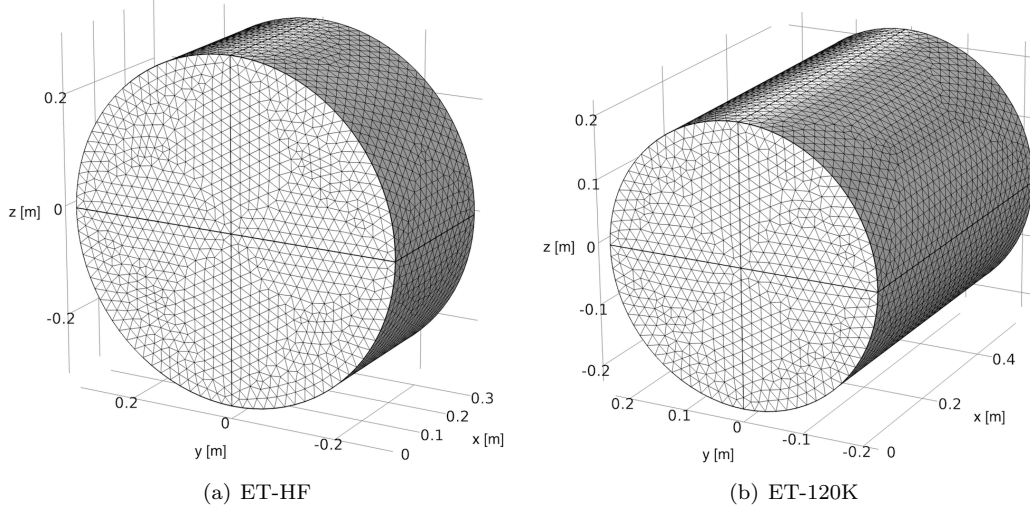


Figure 4.5: Geometries of test masses proposed in the ET-HF and ET-120K (and LIGO Voyager) designs, as modelled in COMSOL by M. Dovale-Álvarez. Dimensions are listed in table 4.2. The mirrors are modelled as simple cylinders currently; details such as suspension points should be added to future models once these designs have been developed.

156 GPa, as is currently used in models of the Voyager design [67]. Future work should update this model to use an accurate model of the mirror crystal structure and corresponding values of Young’s Moduli – as shown in [22], this will affect the frequencies of the resulting modes.

Example acoustic modes of the ET test masses are depicted in figure 4.6. To allow comparison between numbers of parametric instabilities in LIGO and ET, we limit our mode set to the same frequency range of 0-60 kHz. We find that for ET-HF there are 3838 eigenmodes of the mirror in this frequency range—5 times more than for LIGO test masses. This is due to the change in radius-to-thickness ratio and overall increase in mass relative to LIGO. The mass increase naturally leads to a decrease in the frequencies of all modes, so that more modes fall within our chosen window. This increase in mode density naturally means that we may expect more mechanical modes to overlap with an optical resonance, resulting in more PIs.

One important consequence of this is that the frequency of the ‘drumhead’ mode (see figure 4.6(a)) is much lower, at 4.2 kHz. This mode is important because, unlike ‘bulk’ motions such as the fundamental DC motion (the ‘piston’ mode) or twisting modes in pitch or yaw, this mode cannot be controlled using linear or angular alignment control loops, and is therefore the lowest frequency potential PI. We thus use this as a low-frequency cut-off for mechanical modes to include in PI modelling. At 4.2 kHz, the ET-HF drumhead mode currently falls within the 1-5 kHz region of the gravitational wave detection band that is of interest for probing the neutron star equation of state

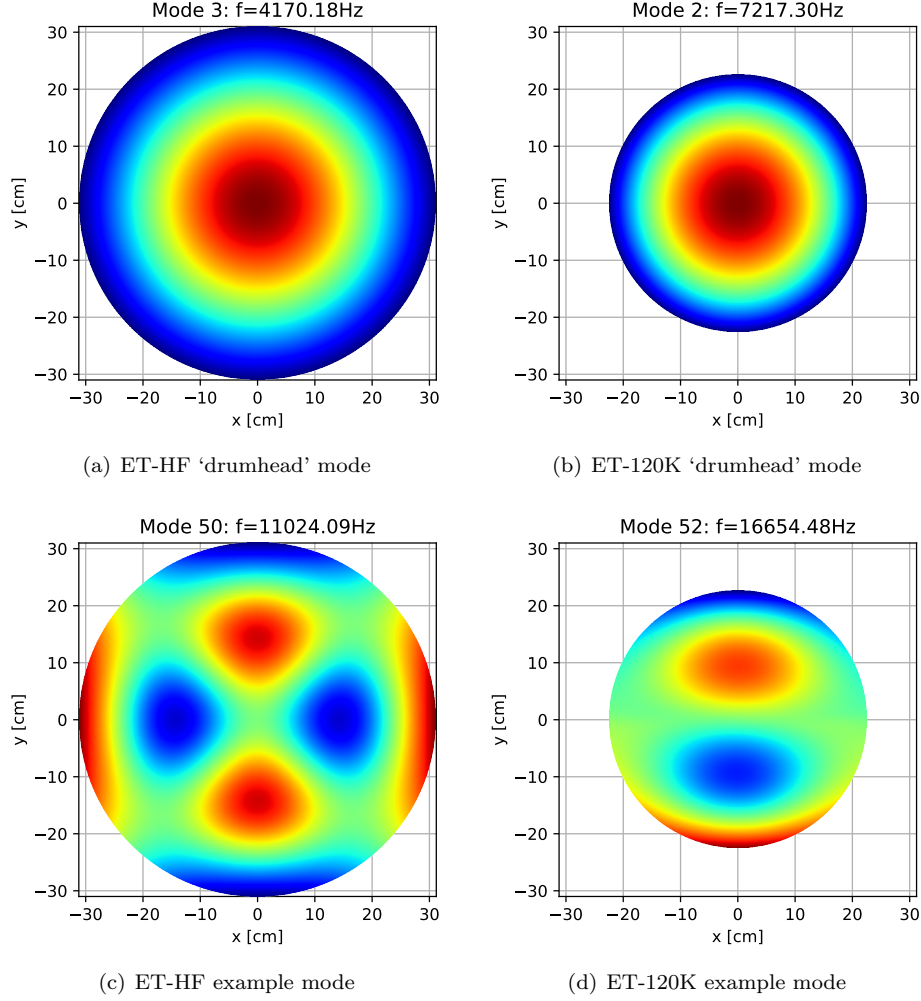


Figure 4.6: Example surface motion maps of mechanical eigenmodes of ET-HF and ET-120K test masses. Mode numbers are identifying tags generated by COMSOL. The colours represent the relative surface displacement: red is positive, blue negative, and green null.

[77]. Care should therefore be taken to design the interferometer such that this mode is unable to ring up in any reasonable operational mode.

By contrast, over the same frequency range there are 1467 eigenmodes for ET-120K. This number is slightly lower than that found for by Zhang *et. al.* in [66], who model the extremely similar LIGO Voyager test mass. This is most likely due to their higher cut-off frequency of 74 kHz. Although still higher than LIGO, this number is much lower than ET-HF. We can understand this by considering the eigenfrequencies f_n of a free string

$$f_n = \frac{n}{2L} \sqrt{\frac{T}{\mu}} = \frac{n}{2L} \sqrt{\frac{Y\sigma}{\rho}}, \quad (4.1)$$

where T is the tension and μ is the mass per unit length; Y is the Young's modulus, σ is the strain in the string and ρ is the density. Although the mass is similar between the two mirror designs, the Young's modulus of silicon is more than double that of fused silica, meaning that the decrease in mode spacing with mass is somewhat compensated. A second consequence of the change in material is that the drumhead mode in ET-120K is now shifted up to 7.2 kHz, thereby avoiding any conflict with the gravitational wave detection band, including the region of interest for the neutron star equation of state.

Table 4.3 summarises the key results of Finite Element Modelling potential mirror designs for ET, compared to the results found for Advanced LIGO mirrors. Based on this alone, it seems that using silicon test masses, even with their elongated geometry, may be advantageous for reducing the PI problem for ET. However, the optomechanical coupling of these modes in the interferometer will differ, and therefore affect the strength of this statement.

	aLIGO	ET-120K	ET-HF
No. eigenmodes in 0-60 kHz range	800	1497	3838
Frequency of drumhead mode	8.1 kHz	7.2 kHz	4.2 kHz

Table 4.3: Summary of FEM results, comparing ET-120K and ET-HF mirror designs to LIGO.

4.3 Speeding up PI Modelling for ET using Clipping Effects

Since we have a significantly higher mechanical mode density for ET than we did for Advanced LIGO, the calculations will be slower. We therefore make use of clipping effects to reduce the maximum number of higher order optical modes which must be included in the calculation to get representative results, as described in [31]. The method is summarised below.

4.3.1 Apertures as a Source of Loss

By default, all FINESSE simulations assume that optics have infinite extent in directions transverse to the beam path. This means that a mirror is treated as a reflective plane, perhaps with a defined radius of curvature. Any light incident on the mirror will be reflected or transmitted, and continue to propagate through the system. However in reality mirrors are of finite size, and therefore act as apertures. Light incident on such a mirror can therefore escape from the system, depending on the transverse distribution of the beam and the dimensions of the aperture. We describe this loss

as *clipping*.

For a fundamental (HG₀₀) gaussian beam centred on a circular mirror, the fractional loss due to clipping is given by:

$$\begin{aligned}\theta_{\text{clip}} &= \frac{P_0 - P_{\text{mirror}}}{P_0} = \frac{P_0 - \int_0^{D/2} I_0 e^{-2r^2/w^2} 2\pi r dr}{P_0} \\ &= e^{-D^2/2w^2},\end{aligned}\tag{4.2}$$

where w is the radius of the beam as defined in equation 2.48, D is the diameter of the mirror, and I_0 is the intensity at the centre of the incident beam. P_0 is the total power in the incident beam, while P_{mirror} is the power inside the locus of the mirror. Typically, we design optical systems such that clipping losses are less than 1 ppm per optic. This corresponds to $D \gtrsim 5 \times w$.

Figure 4.7 shows how the clipping loss and power reflected from a perfectly reflecting flat mirror changes with the mirror diameter, expressed as a multiple of the beam spot size. The effects of circular and square apertures on curved mirrors in optical resonators are explored in detail in [17].




Figure 4.7: Fractional power reflected and lost when a gaussian beam of radius w is reflected off a flat, circular mirror of diameter D , calculated analytically from equation 4.2. A ratio of $D/w \sim 5.3$ corresponds to clipping losses of less than 1 ppm.

FINESSE models apertures by using higher order optical modes to compute the coupling. This modal method is not ideal, since in principal an infinite number of higher order modes is required to recreate a sharp edge. Figure 4.8 illustrates how the power reflected from the mirror is approximated by FINESSE when higher order modes up to different orders are included in the calculation. We see that even including 15th order modes is insufficient to recreate the analytical result when

there is significant clipping. Including more higher order modes also increases the number of calculations FINESSE must do; the total number of calculations scales as $\frac{1}{2}(m+1)(m+2)$ where m is the maximum mode order included (the ‘maxtem’). Therefore a compromise may need to be made between computation time and simulation accuracy: the maxtem included in the model should be sufficient for the numerical result to match the analytical solution to a chosen level. However for cases with minimal clipping loss ($D/w > 5$), a low maxtem is sufficient to accurately model clipping effects.




Figure 4.8: Numerical computation of power reflected from a mirror of diameter D when illuminated by a 1 W gaussian beam of diameter w , using FINESSE. For low D/w ratios with significant clipping losses, the numerical calculation is unable to completely match the analytical result. Increasing the maximum higher order mode included in the calculation (‘maxtem’) improves the accuracy of the model, but also increases the computation time.

4.3.2 Clipping of Higher Order Modes

As described in section 2.5, higher order optical modes describe a range of different intensity distributions transverse to the beam propagation direction. Figure 4.9 shows three example Hermite-Gauss modes, using the same colour map scaling in all three cases. The dashed line marks a locus of $D = 6 \times w$, which we know (from figure 4.7) has clipping losses of $\sim 10^{-8}$ for HG_{00} . Since the higher order modes have a broader intensity distribution, clipping losses from the same mirror when illuminated by these will increase with mode order, as illustrated in figure 4.10.

This allows a limit to be placed on the number of higher order optical modes that should reasonably be included in optical simulations. This limit differs from the compromise discussed in section 4.3.1: here, the motivation to use a higher maxtem comes from effects in the simulation other than the aperture map, and we assume that we have designed our interferometer such that clipping losses for




Figure 4.9: The transverse intensity distribution of higher order modes spreads out with mode order. The three figures show different Hermite-Gauss modes with the same fundamental beam parameter and input power. The dashed circle marks a locus of diameter $D = 6 \times w$. Since the higher order modes have more intensity further from the centre of the mirror, they will experience more significant clipping.




Figure 4.10: Loss due to clipping increases with mode order. Here, a pure Hermite-Gauss mode is shone on a flat mirror of diameter $D = 6 \times w$. The solid line shows the loss due to clipping for modes of type HG_{x0} . These are the most asymmetric modes of their order, with a narrow intensity distribution in one axis and broad in the other. This means that these modes have the greatest clipping loss of their order. For comparison, all of the HG modes at 3rd, 5th and 9th order are also shown—the difference between modes of the same order is far smaller than the difference in loss between mode orders. This plot was produced numerically using FINESSE with maxtem 15.

the fundamental mode are small. Instead, some additional effect, such as a parametric instability, results in some higher order mode content also scattering into the optical system. Sufficient maxtem must be used to represent these new scattered fields. However, if no apertures are included in the model, no light will be lost from the system, and so the power in the new fields is artificially inflated. Including apertures introduces clipping losses, which will have a negligible effect on the fundamental mode but an increasingly significant effect for higher order modes, as shown in figure 4.10. Above a certain threshold order, the optical mode will experience clipping losses that are sufficient to reduce its power to below the level at which it makes a significant contribution to the physical effect being studied. Therefore this threshold becomes the limiting order for the

simulation.

4.3.3 Using Clipping Effects to Improve and Speed Up PI Modelling

The study of PIs at Advanced LIGO described in chapter 3 does not include aperture effects. Mirrors without maps are treated as perfect reflective planes, and the maps generated for the mechanical modes are defined with surface distortions out to the diameter of the mirror, and contain zeros beyond this, i.e. beyond the mirror diameter, they are again treated as a perfect reflective plane.

This means that the numbers of unstable modes given by plots such as figure 3.7 will experience the artificial inflation described above: mechanical modes that interact primarily with high optical mode orders will appear to have a higher gain than they would in reality, since parametric gain depends on the amount of power in the optical field and how much this field is amplified as it cycles through the interferometer.

Figure 4.11 demonstrates this effect for the case of a single LIGO arm. First, the file is run as in figure 3.7, using no apertures and including higher order modes up to 10th order. This allows a list of unstable modes to be collected and used as a subset for faster modelling. Then the same procedure is followed for four cases: no apertures, a single aperture on the ETM (i.e. the mirror which experiences the PI), a single aperture on the ITM, and apertures on both arm cavity mirrors. The radius of the aperture, 17 cm, is 5 times the spotsize on the ETM, and so clipping of the fundamental mode is of order 1 ppm and can be considered negligible. The code is run with different maximum optical mode orders, and the total number of unstable modes, as well as the time take, recorded.

The upper plot shows the total number of unstable modes in each case. While the original version, with no apertures, shows no sign of converging, we see that applying an aperture to the ETM significantly reduces the number of modes we expect to be unstable, and that this number does not substantially change when optical modes above 6th order are included. We also see that including an aperture on the ITM does not have a significant effect on this number. The lower plot tracks the total computation time of each run. Increasing the maxtem has the biggest effect on the time taken, but we also see that applying more apertures increases the integration time. Therefore since the number of unstable mechanical modes does not appear to be affected much by applying a map to the ITM, we can optimise the running time of this type of simulation, while




Figure 4.11: Number of unstable mechanical modes in a LIGO-like Fabry-Perot cavity, and corresponding computation time, resulting from changing the maximum order of HOM coupling (‘maxtem’) included in the model. Including aperture effects in the model limits the number of unstable modes found, resulting in a convergence that does not occur when the mirrors are treated as infinite reflective surfaces. Four cases are compared: no apertures (blue), apertures on one mirror (red: ETM, black: ITM), and apertures on both mirrors forming the cavity. In this example specifying an aperture for the ETM has the most significant effect. Since the number of unstable modes has converged once 6th order optical modes are included, the required computation time can be significantly reduced.

producing more realistic results, by choosing to model PIs with apertures in the arm cavities, and, in the case of LIGO mirrors, running numerical calculations including higher order optical modes up to 6th order.

4.4 PI ‘Forests’ in ET-120K arm cavities

Before considering the full interferometer, it is prudent to understand how design decisions for the arm cavities will affect the number of PIs. I therefore repeat the ‘worst-case-scenario’ modelling shown in section 3.3.2 for ET, to create a list of mechanical modes that could become unstable within ± 2 kHz of their COMSOL-computed eigenfrequency. Here, I consider two arm cavity designs for ET-120K, using test mass radii of curvature $R_{c1} = 5580$ m and $R_{c2} = 6690$ m. In both cases, the cavity is symmetric (I- and ETMs have the same radius of curvature), however the former is chosen to match the ET-LF design while the latter is designed to reduce clipping loss to 1 ppm.

The choice of radius of curvature for the arm cavities is driven by three main considerations: stability, loss, and noise. The geometric stability of a cavity, measured by its g -factor (see section 2.5), affects how easily the cavity can be controlled [78]. The closer a cavity is to geometric instability, the smaller the mode separation frequency becomes, and the more divergent the beam becomes.

Small misalignments and mismatches thus couple more strongly into higher order modes. It is therefore desirable to design the arm cavities to have g -factors to be close to 0.5.

As explained in section 4.3.1, the finite radius of the test masses mean that some light will be lost from the cavity, referred to as clipping loss. This loss will reduce the amount of power circulating in the arm cavity, and ultimately the amount of power reaching the output port of the interferometer. This means the signal to shot noise ratio will decrease, and the overall detector sensitivity at high frequencies is reduced. It is therefore desirable to minimise clipping losses, typically to 1 ppm or less per optic. Since the diameter of the mirror is constrained by mass requirements and manufacturing processes, and the cavity length is constrained by the target sensitivity of the detector, the spot-size to mirror-diameter ratio is also largely controlled by the radius of curvature of the test masses.

The geometric and clipping loss requirements both point to a relatively small spot size on the test masses, and therefore a relatively large radius of curvature. However, this will have consequences for thermal noise. The power spectral density (PSD) of displacement due to Brownian thermal noise for a fundamental Gaussian beam illuminating a block of material is given by [12, 79]²:

$$S_x(f, T) = \frac{4k_B T}{\pi f} \phi_m U = 2k_B T \frac{1 - \nu^2}{\pi^{3/2} f Y w} \phi_m, \quad (4.3)$$

where k_B is the Boltzmann constant, T is the temperature of the environment, f is the frequency, ϕ_m the mechanical loss, and U describes the strain energy of the material. This is determined by the material properties—the Poisson’s ratio ν and Young’s modulus Y , and the radius of the spot on the mirror, w .

The equivalent PSD for thermo-elastic noise is

$$S_{TE}(f, T) = \frac{4k_B T^2 \alpha^2 (1 + \nu)^2 \kappa}{\pi^{5/2} \rho^2 C^2 f^2 w^3} \quad (4.4)$$

where we now introduce the coefficient of thermal expansion α , thermal conductivity κ , mass density ρ , and heat capacity C . Typically the Brownian noise is the dominant noise source to consider.

We see that both of these have inverse relationships to the spot radius, w : $S_x \propto 1/w$ and $S_{TE} \propto 1/w^3$, so we can reduce the thermal noise by increasing the area of the spot interacting with the mirror. For example, one can achieve a larger spotsize by changing the cavity length or mirror

²NB: reference [12] differs by a factor of ν in the equation for S_x ; the two references agree for S_{TE} . Here I consider only scaling with w so the resulting statements are not affected.

radii of curvature (thus changing the cavity stability) [78], or by using a higher order optical mode [80, 81, 82]. In both cases, this comes at the expense of increasing the clipping loss. There is therefore a trade-off between shot noise and thermal noise when designing the arm cavities.

The two arm cavity designs considered here prioritise the two effects differently. Choosing $R_{c2} = 6690$ m corresponds to a spot size of $w_2 = 8.6$ cm. This results in a clipping loss of 1 ppm per optic. If instead we choose $R_{c1} = 5580$ m the clipping loss increases to 61 ppm per optic. The spot size in this case is $w_1 = 10.2$ cm, which means that the brownian thermal noise decreases by a factor of $S_{x1}/S_{x2} = w_2/w_1 = 0.84$ and thermoelectric noise decreases by $S_{TE1}/S_{TE2} = (w_2/w_1)^3 = 0.59$. However, as shown in figure 4.12, this comes at a significant cost for the quantum limited sensitivity of the full detector: In both cases clipping reduces the detector quantum limited sensitivity, however in case 1 the arm power decreases by $P_1/P_2 = 0.85$, leading to large drop in sensitivity.




Figure 4.12: Quantum-noise-limited sensitivity curves for ET-120K with two choices of test mass radius of curvature. When the test masses are modelled including apertures at their radius, clipping can limit the detector sensitivity. While the case of $R_c=6690$ m has 1 ppm clipping, and therefore negligible loss in sensitivity, choosing $R_c=5580$ m causes significant clipping and almost an order of magnitude loss in the quantum-limited detector sensitivity.

Finally we can consider how parametric instabilities behave in the two cases. Test Mass Radius of Curvature affects parametric instability in two ways:

1. Change in Gouy Phase

As demonstrated by section 3.3.1 figure 3.6, this will change which higher order optical modes are able to resonate in the interferometer.

2. Change in Clipping Loss

As highlighted by section 4.3, this reduces the parametric gain. The effect is twofold: firstly, the overall decrease in circulating power reduces the radiation pressure driving all modes. In addition, the higher order optical modes generated by interaction with the mechanical mode are more significantly clipped than the fundamental mode, again reducing the gain. This effect increases with mode order.

Figure 4.13 shows the resulting ‘forests’ of mechanical modes. In each case, the full set of 1467 ET-120K mechanical modes was tested in the arm cavity. The circulating power is set to match that used in the full ET-120K design, i.e. 3 MW. Based on the results of section 4.3, all calculations include higher order modes up to 6th order, and for this test apertures are specified for both the ITM and ETM. Each data point plotted marks a mode that is unstable within ± 2 kHz of its COMSOL eigenfrequency. The coordinates attributed to it in the plot then correspond to its COMSOL frequency and the maximum parametric gain found within the searched window. The 71 unstable modes expected in the case of $R_c = 5580$ m is significantly lower than the 182 modes seen in the $R_c = 6690$ m case, as expected due to the significant change in clipping losses. The slight clustering of unstable modes every ~ 15 kHz corresponds to the FSR of the cavity in both cases. Table 4.4 lists this and other key parameters resulting from each choice of R_c . In these cases we also clearly see the inverse proportionality of \mathbb{R} to the mechanical mode frequency.

We can also see this reduction effect due to clipping in figure 4.14. These follow the same procedure used to produce figure 4.11, now for the two ET-120K arm designs considered. With $R_c = 6690$ m, the clipping level is similar to that in the LIGO case, and so we see a similar behaviour: once an aperture is applied to the ETM, the number of PIs converges and does not change significantly when computed with more than 6th order optical mode couplings. By contrast, the increased clipping of $R_c = 5580$ m means that the number of PIs converges more quickly and is already stable by `maxtem` 4. Note that, compared to LIGO, here the inclusion of apertures on the *input* test mass also has a significant effect on the number of unstable modes. This is because unlike LIGO, the ET arm cavity designs here are symmetric in radius of curvature. As a result, the spot size on LIGO ETMs is larger than on the ITMs, and therefore the ETM is the dominant source of clipping. In ET, the clipping effect on PIs is more balanced.

We therefore have a trade-off between maximising the quantum-limited sensitivity, and reducing thermal noise and PI. A compromise must be made. Many techniques have been developed to reduce PIs (see section 3.1), and as figure 4.13 shows, there will likely be PIs even in the case of significant clipping. Therefore the decision on test mass RoC should, as in the past, prioritise




Figure 4.13: PI ‘forests’ for two test mass radius of curvature choices for ET-120K. Here, we consider a single 10 km arm cavity, and test masses of 45 cm diameter. Each data point marks the peak parametric gain found for a mechanical mode that is unstable within ± 2 kHz of its COMSOL-computed eigenfrequency. This therefore represents a ‘worst case scenario’ for the arm cavities in the two cases. The choice of $R_c = 5580$ m (blue) results in less than half the number of PIs compared to $R_c = 6690$ m (red), due to significant clipping of the beam at each test mass.

thermal and shot noise arguments to first determine the detector sensitivity. In addition to the large-scale RoC choice discussed here, it may also be possible to tune the radius of curvature on a smaller length scale, as demonstrated in [66] and section 3.3.1. However, this length scale is likely to be accessible using ring heaters, while the RoC choices demonstrated here most likely would be too large for that.

4.5 Comparison between Advanced LIGO and ET ‘Forests’

Figure 4.15 compares the ‘worst case scenario’ for parametric instabilities in Advanced LIGO versus the ET-120K designs described in sections 4.1.2 and 4.4. As with figures 3.7 and 4.13, it plots the mechanical modes on ETMX which are unstable within ± 2 kHz of their COMSOL eigenfrequency, marked at their eigenfrequency and the maximum parametric gain found in the frequency range considered. Compared to figure 3.7, we now include clipping effects in the LIGO model, reducing the number of expected PIs from 210 to just 73 modes. For consistency, all detector designs are

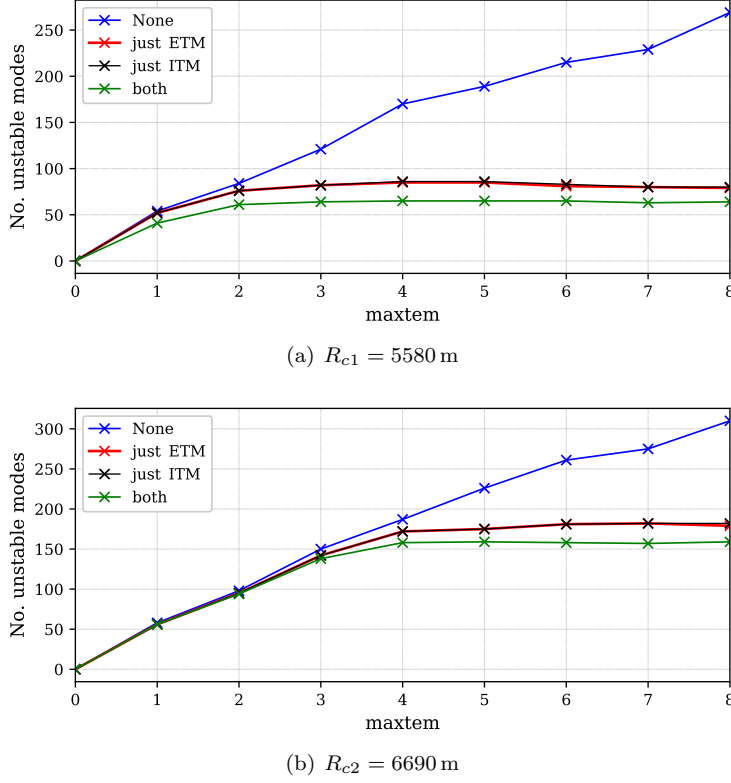


Figure 4.14: Number of unstable mechanical modes resulting from changing the maximum order of HOM coupling (‘maxtem’) included in models of two designs of ET-120K arm cavity. The four cases depicted model the cavity including apertures on the test masses as stated. Unlike LIGO, the ET arm cavities are geometrically symmetric, therefore ITM and ETM apertures have a similar influence on the behaviour.

plotted with aperture maps applied to both the ITM and ETM in the arm containing the surface motion map, and optical couplings are then computed up to 6th order.

The ET designs considered have not been optimised to minimise PIs in the arms or core interferometer as a whole. This gives a baseline understanding of how PIs behave in the designs, which can then be improved upon. Each design includes only the core optics, i.e. it is a dual-recycled Michelson interferometer with Fabry-Perot arms, but does not include any input or output optics such as mode cleaners or frequency-dependent squeezing. The model for each detector design is tuned and mode matched as described in appendix B, with quantum-limited sensitivity curves (assuming perfect mode matching) as shown in figure 4.12. The key frequencies and cavity parameters resulting in the two ET-120K designs are listed in table 4.4; equivalent parameters for the Advanced LIGO design are provided in table 3.1.

In figure 4.15, we first compare the ‘forests’ for LIGO (black) to ET-120K with $R_c=6690$ m (red). These two cases are both designed to have low (<1 ppm) clipping losses, and therefore behave




Figure 4.15: ‘Forests’ of parametric instabilities in two designs of ET-120K compared to the Advanced LIGO design. In all cases, the optical configuration is a dual-recycled Michelson with Fabry-Perot arm cavities. The surface motion map is applied to ETMX, and apertures are included on ETMX and ITMX at the radius of the mirror in each case. For LIGO (black) and ET-120K with test mass $R_c = 6690$ m (red), there is negligible clipping. This ET design has more than double the number of PIs expected compared to LIGO. By contrast, the heavily-clipped case of $R_c = 5580$ m experiences very few PI.

Design		X-arm	Y-arm	PRX	PRY	SRX	SRY
$R_c = 6690$ m	\mathcal{F}	2073.97	2073.97	4.23	4.22	4.24	4.25
	FSR	14.99kHz	14.99kHz	480.6kHz	480.7kHz	480.5kHz	480.6kHz
	f_{pole}	3.6Hz	3.6Hz	57kHz	57kHz	56.6kHz	56.6kHz
	f_{msf}	5kHz	5kHz	29.9kHz ^(x)	29.9kHz	29.9kHz ^(x)	29.9kHz ^(x)
	Ψ_{RT}	-120.7 deg	-120.7 deg	22.4 deg ^(x)	22.4 deg	22.3 deg ^(x)	22.4 deg ^(x)
$R_c = 5580$ m	\mathcal{F}	2073.97	2073.97	4.22	4.22	4.24	4.25
	FSR	14.99kHz	14.99kHz	480.1kHz	480.1kHz	480.0kHz	480.0kHz
	f_{pole}	3.6Hz	3.6Hz	56.8kHz	56.8kHz	56.5kHz	56.5kHz
	f_{msf}	3.1kHz	3.1kHz	17.2kHz ^(x)	17.3kHz	17.1kHz ^(x)	17.2kHz ^(x)
	Ψ_{RT}	-75.2 deg	-75.2 deg	12.9 deg ^(x)	13.0 deg	12.9deg ^(x)	12.9 deg ^(x)

Table 4.4: Key frequencies and related parameters for the cavities in the two ET-120K designs considered. ‘(x)’ indicates that the values in the x- and y-axes differ due to astigmatism introduced by the beamsplitter. In all cases, the difference is of the order of the last significant figure given. Property names are: \mathcal{F} = finesse, FSR = free spectral range; f_{pole} = pole frequency, f_{msf} = mode separation frequency, Ψ_{RT} = round-trip Gouy phase.

somewhat similarly – for example, the quantum-limited sensitivity curve is not significantly affected by the inclusion of apertures in the model, and clipping results in the number of PIs converging

once 6th order HOMs are included. We see that in this case, the maximum number of PIs expected in the dual-recycled ET-120K is roughly triple that seen in LIGO; this is a larger ratio than one may expect based purely on the change in number of mechanical modes.

Comparing figure 4.13 to figure 4.15, we see that the number of unstable modes in the $R_c=6690$ m case does increase with dual recycling, as found for LIGO, however the effect is much weaker than in LIGO’s case due to the change in geometric properties – and therefore resonant conditions – in the recycling cavities. This can be seen in the order of magnitude change in pole and mode-separation frequencies in the ET-120K recycling cavities compared to LIGO.

I have also included the ‘forest’ produced in the ET-120K case of $R_c=5580$ m. As noted above, this case matches that proposed in [12], however it introduces significant clipping losses. Unlike the previous cases considered, including the full dual-recycled interferometer seems to dramatically *reduce* the number of parametric instabilities expected in this model, compared to the single-arm case shown in figure 4.13. This is a result of the small changes in recycling cavity geometries that are required to mode-match to the new arm cavity eigenmode. As with all ET-120K models considered in this chapter, this test file is very idealised, since it is very well mode-matched, tuned and aligned, and is a minimal model of the core optics. Further work should explore whether this reduction in PIs when using dual-recycling is a real feature that could be used advantageously, or whether it relies on perfect interferometer conditions that will be hard to achieve in a physical gravitational wave detector.

4.6 PI in a ‘kHz-optimised’ ET-120K

The current generation of gravitational wave detectors are designed to be maximally sensitive in the 30-400Hz frequency range. They were optimised primarily for binary neutron star (BNS) signals, with enough sensitivity to detect BNS signals from a distance of up to 200 Mpc and stellar-mass binary black hole (BBH) systems to a redshift of $z=0.5$. The lower frequency cut-off limits the maximum total mass of an observed system to $\lesssim 200M_\odot$. The typical goal when designing the next generation of detectors is to deepen and widen this sensitivity range in order to sense signals from further away, and from heavier systems. For example, the ET design study [12] intends to improve on LIGO/Virgo design sensitivity by an order of magnitude, and extend the frequency range down to 1 Hz, meaning that BNS systems could be detected up to redshift $z\sim 2$, stellar mass BBH systems detected to the edge of the observable Universe, and the maximum total mass of a

system would reach $10^2 - 10^4 M_\odot$, meaning that astronomers could start to probe the population of intermediate mass BBHs.

The first ‘multi-messenger’ detection of a binary neutron star system in 2017, named GW170817 [9, 61], has sparked a renewed interest in increasing the sensitivity of gravitational wave detectors at *high* frequencies, i.e. moving into the kHz frequency band. The signal observed included information from the in-spiral of the system, and was sufficient to allow an unprecedented amount of follow-up observation from electromagnetic and particle observatories. However, the post-merger oscillations of the system, at kHz-frequencies, were too weak to be detected by LIGO. This portion of the signal is particularly crucial for probing the neutron star equation of state. The presence of electromagnetic and/or particle counterpart signals means this type of system is a useful probe of many aspects of fundamental physics in an extreme environment.

Both current and future detectors could be adapted to prioritise different science cases such as this. In [83], D. Martynov *et. al.* demonstrate that the quantum limited sensitivity curve of a detector such as LIGO or ET can be optimised to reach high sensitivity in the kHz range, by adjusting the optical resonance of the coupled SRC and arm cavities. In the case of a 10km arm detector like ET, the parameters changed are [84]:

- increase in ITM transmissivity, to $T_{\text{ITM}} = 2.9\%$
- reduction in SRM transmissivity to $T_{\text{SRM}} = 2.4\%$
- reduction in SRC length, to 91.1 m

These result in the quantum-limited sensitivity curve shown in figure 4.16 (black). The curves for the original ET-120K design (red) and Advanced LIGO (blue) designs are also shown for reference. No apertures are used. These parameters are chosen to maximise the sensitivity in the range from 1-4 kHz with an arm circulating power of 3 MW. Squeezing could then be used to further improve the sensitivity. While the sensitivity at high frequencies is improved compared to LIGO and the original ET-120K design, the sensitivity in the middle of the detection band is less improved. At low frequencies ($< 10\text{Hz}$), the quantum-limited sensitivity is also improved, however in practice the detector will be limited by many other noise sources in this band (seismic, controls, and so on).

Here I have chosen to assume that the use of cryogenics and developments in bulk and coating material research mean that thermal noise will not limit ET-120K, so that the adapted design uses test mass curvatures of $R_c = 6690\text{m}$. Since the PRC and SRC lengths are now significantly different, the distance between the ITMs and beamsplitter, which is common to both recycling




Figure 4.16: Quantum-noise-limited sensitivity curve for ET-120K when the optical design is optimised for gravitational wave signals in the 1-4kHz frequency band. The optimised curve (black) has an order of magnitude improvement in sensitivity in the target frequency range when compared to the non-optimised design (black), but is less sensitive in the central 10-1000Hz range. Both designs are more sensitive than the advanced LIGO design sensitivity (blue dashed).

cavities, must be reduced and the PRC-BS length becomes large as a result. While this does not change the resulting sensitivity curve, it will have practical consequences since, for example, the spot size on the beamsplitter will be different.

Design		X-arm	Y-arm	PRX	PRY	SRX	SRY
$R_c = 6690$ m;	\mathcal{F}	213.32	213.27	4.14	4.14	4.17	4.17
SRC & ITMs	FSR	14.99kHz	14.99kHz	480.6kHz	480.7kHz	435.8kHz	435.8kHz
optimised for	f_{pole}	35.1Hz	35.1Hz	58.0kHz	58.0kHz	52.3kHz	52.3kHz
kHz signals	f_{msf}	5.0kHz	5.0kHz	29.9kHz ^(x)	29.9kHz	78.6kHz ^(x)	78.6kHz ^(x)
	Ψ_{RT}	-120.7 deg	-120.7 deg	22.4 deg ^(x)	22.4 deg	64.9 deg ^(x)	64.9 deg ^(x)

Table 4.5: Key frequencies and related parameters for the ‘kHz-optimised’ ET-120K design. ‘(x)’ indicates that the values in the x- and y-axes differ, on the order of the last significant figure given, due to astigmatism introduced by the beamsplitter.

This design adaption is interesting for PI studies since it breaks the geometric symmetry of the SRC/PRC designs used so far in ET proposals. This can be seen by comparing the key parameters of the new design, listed in table 4.5, with those in table 4.4. In the geometrically symmetric case, the PRC and SRC recycling effects on higher order optical modes are expected to be similar to one-another. In LIGO, where there is no such symmetry, a mechanical mode could become unstable through resonance via either the power- or signal-recycling cavity, with different resonant conditions. For many mechanical modes, such as the example ‘mode 37’ used throughout chapter 3, this results in a pair of sharp peaks in the optical response, with further splitting associated with recycling cavity astigmatism. This mode would therefore be included in plots like figure 3.7 in cases

with either just power- or just signal-recycling. However, some mechanical modes may generate HOMs that are only able to resonate via one of the two recycling cavities, so we find more unstable modes in the dual-recycled than power-recycled cases. Overall this means that there are more ways for a mechanical mode to be unstable in cases where the PRC and SRC have differing resonant conditions, as the degeneracy is lifted.




Figure 4.17: The ‘PI forest’ for the ET-120K kHz-optimised design (red) is not significantly different to that of the ‘original’ ET-120K design.

Figure 4.17 shows the resulting ‘forest’ of PIs in the kHz-optimised design, with its significantly shorter SRC. We see that, as expected, the number of potentially unstable modes per test mass increases. The gains and frequencies of unstable modes also change a little. However, in this case with negligible clipping loss the number of unstable modes is already high, so that 19 more potential PIs is not a particularly dramatic change in behaviour. Design proposals such as this are not immediately negated due to their effect on PIs.

4.7 Future work: ‘PI-Aware’ Design for ET

Through the many examples in this chapter and chapter 3, I show that the entire design of an interferometric gravitational wave detector influences parametric instability in it, and the magnitude of the problem PIs will present in keeping the detector controlled and able to observe gravitational wave signals. Therefore it is important to consider PIs while we are still in the design and development stages of third-generation detectors. The experience we have developed during the commissioning and running of the advanced LIGO detectors can be used to our advantage.

‘PI-Aware’ detector design should make use of the mitigation techniques that have proven successful for LIGO, such as those described in section 3.1, but can also directly adapt the core detector design in ways that are not possible to try in sites that already physically exist. The strategies for minimising PIs in future detectors fall into four broad categories, which can be ranked by how complicated they could be to implement compared to a design that does not consider PIs:

1. Advantageous choices of detector parameters

Can the detector design itself be modified to naturally reduce the number of PIs we might expect without adding additional components or control systems, and without impacting the detector sensitivity or controllability?

Examples: geometric and material properties of the test masses, cavity designs

2. Passive mitigation

Additional elements introduced to deliberately reduce PIs in a passive way without significantly changing the detector design. This introduces additional components to the detector, with behaviours that may impact the detector sensitivity.

Examples: acoustic mode dampers (AMDs)

3. Active mitigation

Active control strategies to mitigate PIs if and when they ring up. This can introduce new control loops, and may introduce new components depending on the actuators and sensors used.

Examples: ring heaters, electrostatic drivers (ESDs)

4. ‘Exotic’ additional measures

More complex, or as-yet untested, mitigation strategies that require more significant adjustments to the detector design.

So far at LIGO the strategies have fallen into the third category. The ring heater tunings used at Livingston were not set using an active control loop, but were a manual correction set in response to the observed PIs. The ESDs used at Hanford are used in active control loops for a list of known mechanical modes. It is hoped that the use of AMDs will, in the future, mean that these active techniques are no longer necessary. Based on this, there are a few key steps in designing detectors, from the perspective of parametric instabilities:

1. Determine the key parameters of the detector for it to achieve its science case. For example, the quantum noise requirements will determine the power circulating in the arm cavities, the mass of the mirrors and the general topology of the detector, while thermal noise requirements will affect the material choices for the test masses.
2. Make a model of the core interferometer, as detailed as is available, and produce a ‘forest’ plot (e.g. figure 4.15) to assess the initial situation of how many parametric instabilities can be expected, and with what gains.
3. Determine whether AMDs can be used to suppress all the expected PIs without increasing the test mass thermal noise to the extent that it limits the detector sensitivity.
4. If thermal noise constraints, or other issues, mean that AMDs cannot be used (or used exclusively), determine whether the core detector parameters can be used to suppress the remaining instabilities without compromising the detector sensitivity. This should first consider the test mass and arm cavity design, and then move to include the full core interferometer design and beyond.
5. If necessary, explore other more exotic passive mitigation strategies, and strategies that affect the sensitivity curve of the detector.
6. Include sufficient actuators and sensors in the planned control scheme so that, should a PI still arise, active mitigation can be used. As with LIGO, it is likely that such actuators will already be present in the detector control scheme for other purposes, however this additional role should not be overlooked.

The work demonstrated throughout this chapter falls largely into steps 2 and 4. Examples of the latter include consideration of the test mass radii of curvature and resulting clipping (see section 4.4 and [66]), exploring material and geometry choices for the test masses (as in section 4.2 and [63, 64]), and then considering parameters in the wider interferometer such as the length and Gouy phase accumulated in the recycling cavities (as explored in both this chapter and chapter 3).

If acoustic mode dampers prove as successful as they so far appear to be, these may largely solve the issue of parametric instability for LIGO and should therefore be given early consideration when looking at PIs in future detectors.

4.8 Conclusions

The Einstein Telescope is a proposed third-generation gravitational wave detector in Europe. It is intended to have a peak sensitivity that is ten times that of current detectors like LIGO, over a wider frequency range. The technology improvements required to achieve this include using significantly higher circulating optical power, (~ 3 MW) and heavier test masses (~ 200 kg). As a consequence, it is expected that ET will experience more parametric instabilities than LIGO. This could significantly limit the observation time of the detector.

The ET design is still in development, and is expected to evolve over the facility's multi-decade lifetime. I have studied how various design decisions for ET will affect PIs.

The designs referred to as 'ET-HF' and 'ET-120K' use the same circulating power and mirror mass. However, the ET-HF mirrors are made from fused silica, while ET-120K mirrors will be silicon, since this design uses cryogenic cooling to reduce thermal noise. This material difference means that the mirror geometry is also expected to be different due to fabrication limitations. While both ET designs have a significantly higher mechanical mode density than LIGO test masses, the mode density in the ET-120K test mass design is less than half that of the ET-HF test mass. This is an indication that ET-120K is likely to experience fewer PIs than ET-HF. In addition, the drumhead mode—the mechanical mode which could produce the lowest frequency PI—has a lower frequency for ET-HF, at 4.2 kHz, than ET-120K at 7.2 kHz. The sensitivity band of interest for probing neutron star physics is 1-4 kHz, therefore the ET-HF design should be optimised, or controlled, to ensure that this mode is not unstable.

Thermal noise can be reduced by increasing the spot size of the laser beam on the mirror. However, since the mirror has finite dimensions, this increases clipping loss, reducing the quantum-limited sensitivity of the detector. Clipping, however, has a positive effect for PIs since higher order optical modes, which have a larger spatial distribution than the fundamental mode, are clipped more. I compared the sensitivity curves and parametric instabilities resulting from designing the ET-120K arm cavities to (a) keep clipping loss below 1 ppm per test mass, and (b) reduce the test mass Brownian thermal noise by 20% per optic. In the latter case, clipping effects resulted

4.8. CONCLUSIONS

in a significant decrease in the quantum-limited sensitivity of the detector, but did also reduce the number of parametric instabilities in a model of the arm cavity from 182 to just 71 modes. While dual recycling increased the number of unstable modes in the low-loss case, it *decreased* the number of PI in the highly-clipped case to just 17. This should be explored further.

The detector design can also be optimised to specifically study neutron star physics, by adjusting the signal-recycling cavity length and the SRM and ITM transmissivities. In this case I found that more mechanical modes are able to become unstable in the system, since the SRC and PRC properties are no longer identical and have different resonant conditions. In the low-clipping-loss detector design considered, the absolute number of instabilities added is insignificant when compared to total number.

Finally, I developed an action plan to incorporate PI considerations directly into the design process for future detectors, ensuring that PIs do not limit detector performance. This can be used to ensure that PIs are minimised and controllable in detectors like ET. This broad study includes considering the detailed design of the optics, using the mitigation strategies that have proven successful for LIGO, and considering the effect of the optical design using the techniques I have demonstrated.

Chapter 5

Bandwidth Manipulation using an Optomechanical Filter

This chapter describes the concept of an ‘optomechanical filter’—a radiation-pressure-dominated, suspended cavity which is operated in an unstable regime such that the field reflected from it exhibits negative dispersion. As described by Haixing Miao *et. al.* in [13], such a filter could be used to increase the sensitivity of gravitational wave detectors at high frequencies, where they are limited by quantum shot noise.

Section 5.1 describes the motivation for a negative dispersion element and its application to gravitational wave detection. Section 5.2 then illustrates the effect of an idealised negative dispersion element on an arm cavity. In section 5.3 I provide the input-output relation of the radiation-pressure-driven cavity, then compare the approximated analytical result to that generated using FINESSE. I then show how this can be used as a filter for another cavity, resulting in overall increase in the bandwidth of the system.

5.1 Broadening the Bandwidth of Gravitational Wave Detectors using Negative Dispersion

The quantum-limited sensitivity curve of interferometric ground-based gravitational wave detectors is shaped by the configuration of the detector. Considering a dual-recycled Michelson interferome-

ter with Fabry-Perot arm cavities, the tuning of the signal recycling mirror can be used to operate the detector in a high-sensitivity, narrow-bandwidth mode (‘signal recycling’), or a lower-sensitivity, broadband mode (‘resonant sideband extraction’). The frequency response of the detector is therefore a compromise between peak sensitivity and bandwidth [85].

There are many approaches to improving the quantum-limited sensitivity of such detectors, particularly at high frequencies, where current gravitational wave detectors are shot-noise limited. Increasing the optical power in the detector reduces shot noise, boosting the high-frequency sensitivity at the expense of increasing quantum radiation pressure noise at low frequencies. Frequency-dependent squeezing [68, 69, 70, 71] can be used to improve detector sensitivity across the detection bandwidth. Other schemes have also been suggested, such as adjusting the tuning of the signal-recycling mirror to track particular signals of interest.

Another option, proposed to broaden the bandwidth without compromising the peak sensitivity, is to introduce an element that compensates for the dispersion of the optical gravitational wave signal as it propagates in the arm cavities. This component would be required to exhibit the opposite phase behaviour to that of free space: light must dissipate, rather than accumulate, phase. We refer to this as *negative dispersion*. In an ideal case, this could be used to create a ‘white light cavity’, i.e. a cavity that is resonant for all wavelengths of light. In practice, it would broaden the bandwidth of the detector, extending the peak sensitivity to higher gravitational wave frequencies.

There have been several proposals for producing a ‘white light cavity’. Examples include using an atomic medium which is pumped to generate negative dispersion at a point of zero absorption [86, 87, 88, 89], introducing pairs of diffraction gratings into the arm cavities to create a frequency-dependent optical path length [90, 91], and using stimulated Brillouin scattering to generate negative dispersion in optical fibres [92]. Such passive negative dispersion elements have now been ruled out [89].

In this chapter I study the use of an active optomechanical system to generate negative dispersion. The idea, proposed by H. Miao *et. al.* in 2015 [13], is to use an optical cavity containing a suspended mirror as an active *optomechanical filter* which can be introduced into the signal-recycling cavity of gravitational wave detectors. A similar idea, using optomechanical couplings within a fibre, is also in development [93]. As described in chapter 2, cavities containing suspended optics exhibit dynamic behaviour, resulting from a coupling between the optical field and the mechanical motion of the optics. In the regime where the cavity is blue detuned ($\Delta > 0$), the optomechanical system is anti-damped, with the optical damping term Γ_{opt} taking the opposite sign to the mechanical

damping of the suspended mirror, γ_{m0} . In cases where radiation pressure dominates, $|\Gamma_{\text{opt}}| > |\gamma_{m0}|$, the system is dynamically unstable. In this regime, the lower sideband generated on interaction with the mirror exhibits negative dispersion, as demonstrated in section 5.3 below. Through careful choices of relative frequencies, we can then use this as an active optomechanical filter.

5.2 Idealised Negative Dispersion and Bandwidth Broadening

The idealised behaviour of a negative dispersion component is to exactly compensate for the phase accumulated by a gravitational wave signal as it propagates through a cavity, without introducing any additional effects, at all frequencies. The accumulated phase, or ‘dispersion’, of a gravitational wave sideband at frequency Ω , as it propagates the length of the arm cavity, is described by

$$\phi_A = \Omega L_A / c \quad (5.1)$$

where L_A is the arm cavity length and c is the speed of light. For it to have infinite bandwidth, we therefore require $\phi_A = N\pi$ (see section 2.3) for all Ω , i.e. the cavity length should be $L_A = 0$, and the arm cavity in this case is just a mirror. This is obviously not useful for gravitational wave detection: the measured signal relies on a differential change in path length, and the signal amplitude scales with the length over which it is measured. We therefore want a long arm cavity. The cavity also acts to amplify the signals by the cavity finesse. Instead, we wish to achieve an *effectively* zero-length cavity by introducing an additional optical element somewhere in the cavity with a phase response

$$\phi_f = -\Omega L_A / c \quad (5.2)$$

so that the total phase on 1-way propagation through the cavity is $\phi_A + \phi_f = 0$. Phase like ϕ_f is referred to as *negative dispersion*, for obvious reasons.

Using only simple optical components, we can think of introducing a ‘magic’ filter cavity with a negative length, so that it exhibits negative dispersion, and coupling this to the end of the arm cavity as shown in figure 5.1. Equation 2.18 noted that the effective reflectivity of a cavity is approximated by $-e^{-2i\Omega/\gamma_f}$, i.e. it primarily acts to change the phase of the input field. The total phase is scaled by the new cavity’s finesse, \mathcal{F}_f , so that by setting the new cavity length as $L_f = -\frac{\pi}{2\mathcal{F}_f} L_A$, we can largely compensate for the phase accumulated by the signal in the arm cavity.




Figure 5.1: Optical layout used to model ideal negative dispersion and resulting bandwidth broadening. A dispersive ‘arm’ cavity of length L_A is coupled directly to a ‘magic’ filter cavity which exhibits negative dispersion due to having a negative length L_f . Signal sidebands at frequency Ω are injected as phase modulation sidebands by shaking the input mirror, and are then propagated through the system. The overall bandwidth of the system is measured via the complex field amplitude $E_{\text{out}}(\Omega)$.




Figure 5.2: Bandwidth manipulation using a ‘magic’ filter cavity. The bandwidth of the arm cavity (blue) can be altered by coupling it to a second cavity of appropriate length. When this cavity has positive dispersion (red), the bandwidth of the overall system is reduced. If instead the filter has negative dispersion (black), the bandwidth can be increased.

Figure 5.2 shows the effect of adding this additional, idealised cavity. The end mirror transmissivity is zero, and therefore when on resonance the peak amplitude of the reflected field is unchanged. When the filter cavity length is positive, the overall bandwidth of the system is reduced when compared to the arm cavity with no additional optics. Setting the length as above, the bandwidth is significantly broadened and we can see that the phase accumulated in the arm has been compensated by the filter. At high frequencies, the bandwidth of the filter cavity itself becomes the limiting feature, therefore using any cavity-based negative dispersion element is not expected to produce an infinite-bandwidth gravitational wave detector, but can significantly broaden it. In this model the coupled cavity responses are also 180° out of phase with that of the single cavity response. This is not a physical change in behaviour, but results from the choice of reflection and

transmission phase convention used to build the model. Choosing to shift the phase of transmitted beams by $\pi/2$ means that there is an addition factor of -1 for every extra mirror added in series in a coupled cavity system.¹

The goal of any ‘white light cavity’, or negative dispersion filter, is to replicate this behaviour as closely as possible using physical components.

5.3 Generating Negative Dispersion using Optomechanical Systems

In an optical cavity with suspended mirrors and significant circulating power, an optomechanical feedback loop can be produced due to the radiation pressure from the optical field pushing on the mirrors. Depending on the relative frequency of the input laser compared to the cavity resonance, this can result in an optical spring with anti-damping, or a damped optical anti-spring. An overview of optical springs is provided in chapter 2. In this section, I outline how an optical spring with anti-damping, in the correct regime, can act as an optical element with negative dispersion.

5.3.1 The Input-Output Relation for a Radiation Pressure Limited Cavity

As in section 2.4, we consider the case of a two-mirror cavity in which the end mirror is suspended, and the dynamics of the system are dominated by radiation pressure effects. Such a system also alters the phase relationship between the various frequency components of the optical field. This behaviour can be explored by measuring the transfer function of the system, also referred to as the *input-output relation* of the system. Here we consider the behaviour of a pair of sidebands which are generated by amplitude modulating the input field at frequency ω , as depicted in figure 5.3.

The input-output relation of interest compares the complex amplitude of the sidebands on input and on reflection from the cavity. This is typically first derived using a ‘quadrature picture’ approach, where the optical field is described in terms of its amplitude and phase (see for example section 2.5 of [94]), but can also be derived directly using the ‘sideband picture’, i.e. in terms of the frequency components of the field, in a similar manner to section 2.4. In both cases, the optical fields are treated as plane waves. Appendix C provides the sideband picture derivation in full.

¹We also see this behaviour in FINESSE, since it uses the same convention. It means that a cavity which is modelled as a 3-mirror series system in which the central mirror has $T = 1$ will have a 180° offset in the phase response compared to a model of the equivalent two-mirror system for all frequencies.




Figure 5.3: Optical layout used to derive the Input-Output relation

For a sideband with angular frequency ω , generated around an off-resonant pump field with detuning Δ relative to the cavity resonant frequency ω_0 , the input-output relation on reflection is

$$E_{\text{out}}(\omega) = \left(1 - \frac{2i\gamma [-i\gamma - \Delta + \omega - \hbar|g|^2\chi_{xx}(\omega)]}{(\gamma + i\omega)^2 + \Delta [\Delta + 2\hbar|g|^2\chi_{xx}(\omega)]} \right) E_{\text{in}}(\omega) + \frac{2i\hbar g^2 \gamma \chi_{xx}(\omega)}{(\gamma + i\omega)^2 + \Delta [\Delta + 2\hbar|g|^2\chi_{xx}(\omega)]} E_{\text{in}}^*(-\omega), \quad (5.3)$$

where γ is the bandwidth of the cavity, $\chi_{xx}(\omega)$ is the force to displacement mechanical transfer function of the suspended mirror, and $g = i\sqrt{2P_c\omega_p/(\hbar Lc)} e^{i\phi}$ where P_c is the power incident on the suspended mirror, ω_p is the pump laser frequency, L is the cavity length and ϕ describes the phase difference between the input carrier and the carrier incident on the suspended mirror. This equation assumes that the single mode approximation is valid, that the phase accumulated in the cavity and input mirror transmission are small enough to be approximated by Taylor expansions, and that the end mirror is perfectly reflective.

We can write equation 5.3 more compactly as

$$E_{\text{out}}(\omega) = T_1(\omega)E_{\text{in}}(\omega) + T_2(\omega)E_{\text{in}}^*(-\omega), \quad (5.4)$$

where we introduce $T_1(\omega)$ as the ‘upper-to-upper sideband transfer function’, describing the component of the output field at frequency ω which is contributed by the input field at that frequency. $T_2(\omega)$ is then the ‘lower-to-upper sideband transfer function’, which describes the field cross-coupled from the input lower sideband to the output upper sideband. In cases where radiation pressure does not affect the system, i.e. $g\chi_{xx}(\omega)$ is negligibly small, there is no such cross-coupling and the input-output relation reduces from equation 5.3 to equation 2.18.

The radiation pressure force on the end mirror at frequency ω is generated by the beating of the carrier field at frequency ω_p and the sideband fields at $\omega_p \pm \omega$. The total power incident on the

mirror is therefore

$$P_2 = |E_2(0) + E_2(+\omega) + E_2(-\omega)|^2, \quad (5.5)$$

where $E_2(\pm\omega)$ describes the sideband fields incident on the end mirror and $E_2(0)$ is the carrier field at angular frequency ω_p . The driving radiation pressure force at frequency ω on the suspended mirror therefore has form

$$F_{\text{RP}}(\omega) \sim \frac{2}{c}(E_2(\omega)E_2^*(0) + E_2^*(-\omega)E_2(0)), \quad (5.6)$$

and the displacement of the mirror in the frequency domain is correspondingly given by $x_{\text{RP}}(\omega) = \chi_{xx}(\omega)F_{\text{RP}}(\omega)$. This mirror displacement results in phase modulation sidebands at $\pm\omega$ whose complex amplitudes, for both the upper and lower sideband, contain contributions from both $E_2(\omega)$ and $E_2(-\omega)$.

Since ω describes any generic sideband frequency, the lower sideband is also described by equation 5.3, which we write explicitly as

$$E_{\text{out}}(-\omega) = T_1(-\omega)E_{\text{in}}(-\omega) + T_2(-\omega)E_{\text{in}}^*(\omega). \quad (5.7)$$

$T_1(-\omega)$ therefore describes the lower-to-lower sideband transfer function, while $T_2(-\omega)$ is the upper-to-lower sideband transfer function.

The approximated analytical result given by equation 5.3 is compared to the output of a FINESSE model of the same system in figure 5.4. The input sidebands are generated by amplitude modulating the pump laser input, and measured using appropriate `ad` detectors². The parameters of the model in both cases are as listed in table 5.1. FINESSE linearises the optical system by pre-computing the power incident on the suspended mirror and using this to describe the optomechanical coupling from the mirror motion into the sideband fields, much as used above. By default, all modulations are modelled as producing a single pair of sidebands. FINESSE then propagates the three fields through the optical system without applying further approximations. As shown, the analytical and numerical models show good agreement in this parameter regime.

² There is one nuance required to be able to model the four transfer function components in FINESSE. By design, FINESSE will compute the total complex field amplitude of a frequency component at each node in the network. Therefore using an `ad` detector to measure the field reflected from the cavity at frequency ω will include contributions from both $T_1(\pm\omega)$ and $T_2(\pm\omega)$ inherently. In order to separate out the contributions from the input upper and lower sidebands, an (unphysical) frequency-dependent beamsplitter is used to selectively transmit only the upper or lower input sideband, and the carrier, into the cavity. The complex amplitude of the reflected field is then only the component contributed by the transmitted sideband.

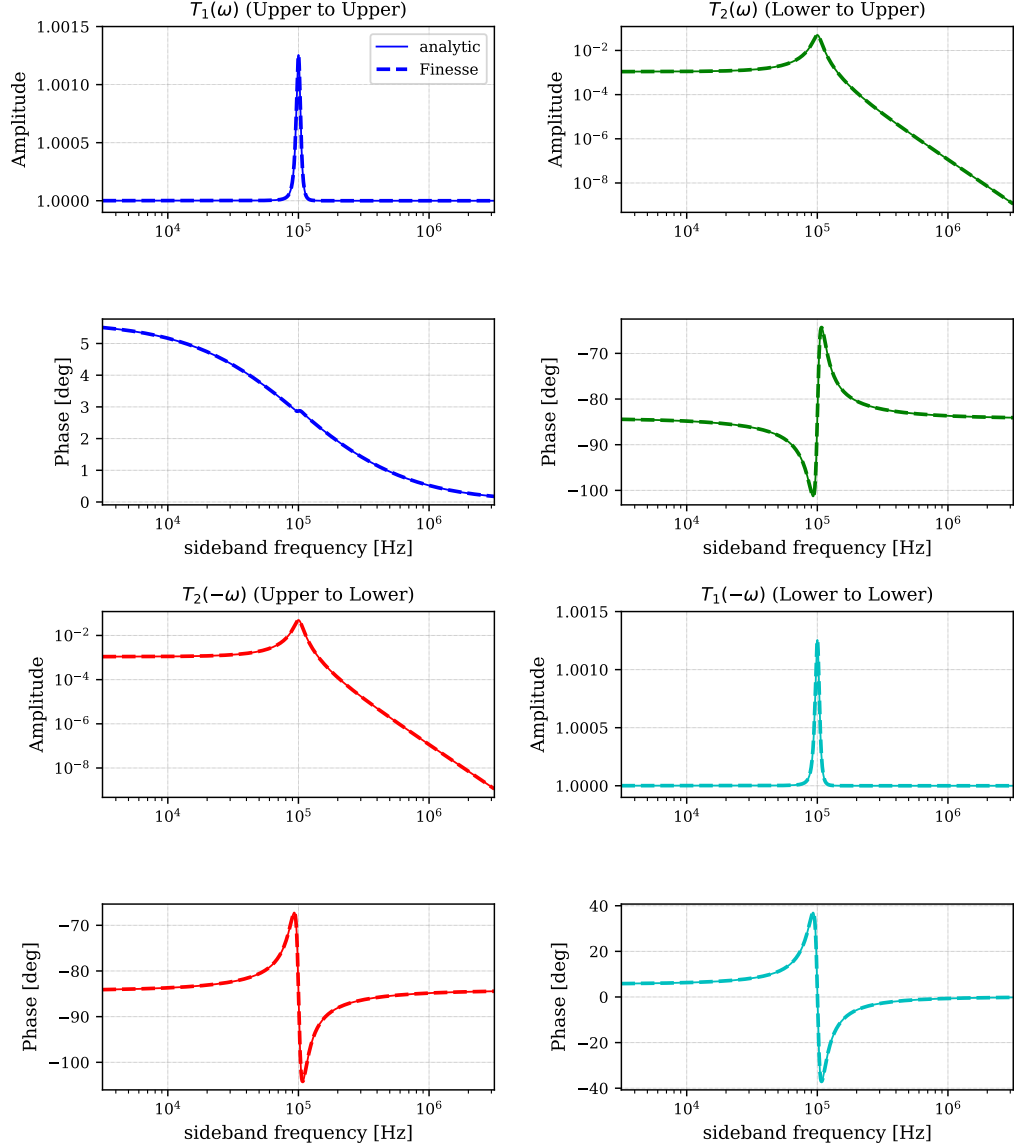


Figure 5.4: Comparison between direct analytical equations and FINESSE calculation of transfer functions for a 2-mirror cavity containing a suspended mirror. In this case the pump laser is significantly positively detuned from cavity resonance such that $\Delta = \omega_m$ and the lower sideband $-\omega$ is well resolved, and resonant in the cavity. The analytical equations (dashed lines) are those presented in equation 5.3, which are valid in the single mode approximation, and assuming that the input mirror transmissivity and accumulated phase are both small. Parameters for the comparison model were therefore chosen to match this regime. FINESSE (solid lines) does not apply all of these approximations, instead solving steady state equations directly. We see that the FINESSE result matches the approximation in this regime. Note that the x-axis in all cases is the frequency *magnitude*.

Parameter	Value
cavity resonant frequency $\omega_0/2\pi$	$c/1064\text{ nm}$
cavity length	1.19 m
input mirror transmission	500 ppm
end mirror transmission	0
cavity bandwidth $\gamma/2\pi$	5 kHz
laser detuning $\Delta/2\pi$	+100 kHz
mechanical frequency $\omega_m/2\pi$	100 kHz
mechanical Q-factor	10^{11}
mirror mass	5 mg
input power	70 W
circulating power P_c	1.4 kW

Table 5.1: Parameters used to compare the input-output relations generated using equation 5.3 to an equivalent finesse model. Parameters directly used in the analytical equation are noted by their symbol, others are derived, or as used in FINESSE.

Figure 5.4 shows Bode plots of the transfer functions for the upper-to-upper and lower-to-lower sidebands $T_1(\pm\omega)$ and cross-coupled terms $T_2(\pm\omega)$, for the case of positive detuning, chosen such that $\Delta = \omega_m \gg \gamma$, i.e. the resolved sideband regime. This choice is particularly relevant for the optomechanical filter. In this case, the upper sideband is far from resonance in the system, and therefore the upper-to-upper sideband transfer function $T_1(\omega)$ looks largely as would be expected for a cavity without radiation pressure effects and a perfectly reflective end mirror. The cross coupled terms $T_2(\pm\omega)$ behave identically aside from a sign change in their phase response. This is not a fundamental change in their behaviour relative to one-another, but instead comes from their phase response being defined relative to the carrier field at ω_p in this reference frame. The lower-to-lower sideband transfer function $T_1(-\omega)$ exhibits the key behaviour of interest. As noted above, the setup is configured such that the lower sideband will be resonant in the cavity when its frequency magnitude matches the mechanical resonant frequency. We therefore see a resonant peak at $|\omega| = \omega_m$ and corresponding phase change, which is also observed in the cross-coupled terms. A small peak is also generated in the amplitude of all four transfer functions as a result of this, due to the cross-coupling effects generated in the mirror interaction.

Figure 5.5 takes a closer look at the lower-to-lower sideband coupling, and compares the behaviour of $T_1(-\omega)$ when the end mirror is fixed in place to when it is suspended with significant incident radiation pressure. When suspended with significant radiation pressure, the primary effect of interest for this work is the sign change observed in the phase response: this is the anomalous, or *negative*, dispersion we require to use the optomechanical system as a filter for bandwidth broadening. In addition, we confirm that power from the pump field, ω_p , is coupled into the sideband field on interaction with the moving mirror, causing the small increase in the amplitude

of the reflected field relative to the input at $|\omega| = \omega_m$.

In the process of comparing FINESSE to analytical derivations of equation 5.3, we have identified several differences between the sign conventions used typically by theorists and experimentalists in the gravitational wave community. These are noted in appendix D.

5.3.2 Radiation Pressure Limited Cavity as an Active Filter

Equation 5.3 described the input-output relation for a sideband of a pump laser input, entering an optical cavity driven by radiation pressure. From the perspective of an input field at the frequency of the lower sideband, the blue-detuned unstable optomechanical cavity produces negative dispersion, as shown in figure 5.5. We can therefore use this as a bandwidth-broadening filter when coupled to another optical cavity, by designing it such that the lower sideband of the pump laser corresponds to the resonant frequency of the test cavity, such that it maximally resonates. This corresponds to setting the detuning of the pump laser to $\Delta = \omega_m$. Figures 5.6 and 5.7 provide a summary of key frequencies used and how these interrelate.

The stable, linear ‘arm’ cavity is held on resonance at the probe frequency ω_0 . Audio-frequency




Figure 5.5: Lower-to-lower sideband input-output relation for a cavity containing a moveable end-mirror. The amplitude modulated sidebands are generated around a pump laser input which is far blue-detuned from the cavity resonance. In this regime, the cavity dynamics are dominated by optical anti-damping. Both the upper and lower reflected sidebands are shown in two cases: (black) fixed end mirror, and (cyan) suspended mirror with high circulating power. This shows that in cases where radiation pressure dominates the dynamics of the system, the phase change takes the opposite sign to usual: negative dispersion. The detuning here is chosen such that $\Delta = \omega_m = 100$ kHz, meaning that the lower sideband is maximally resonant in the cavity.




Figure 5.6: The optomechanical filter concept: a cavity containing a suspended mirror is illuminated by a high-power, off-resonant pump field, with frequency ω_p chosen such that the lower sideband produced in the optomechanical interaction is resonant in the cavity at probe frequency ω_0 . This is coupled to a second, stable cavity which also has resonant frequency ω_0 . Since the lower sideband $\omega_p - \omega_m = \omega_0$ is resonant, fields at that frequency experience negative dispersion on reflection from the unstable cavity. Signal sidebands, generated around the probe at $\omega_0 \pm \Omega$, where $\Omega < \gamma_A$, therefore experience this anomalous dispersion and the overall bandwidth of the system is broadened. Figure 5.7 shows the relative frequencies graphically.




Figure 5.7: Key frequency definitions used in the optomechanical filter, based on figure 4(b) of [13]

signal sidebands, Ω , are generated around the probe, and propagate in the cavity. The effective bandwidth of the arm cavity is described by the behaviour of the signal sidebands in the system, and how these interact with the optomechanical filter. The frequency reference must therefore be shifted. We first describe the frequency in equation 5.3 as $\omega = \Omega - \omega_m$ (see figure 5.7), and then shift the reference frame from the pump frequency, ω_p , to ω_0 . In the blue-detuned case that is required to achieve negative dispersion, the input-output relation for the upper signal sideband $+\Omega$ thus becomes

$$E_{\text{out}}(\Omega) = T_1(\Omega)E_{\text{in}}(\Omega) + T_2(\Omega)E_{\text{in}}^*(2\omega_m - \Omega) \quad (5.8)$$

The detuning of the pump, and therefore the mechanical frequency of the suspended mirror ω_m , should be sufficiently large that the pump is not resonant in the optical system. The signal sidebands, meanwhile, should be able to propagate through the entire system in order to interact

with it. The relevant frequency regime for the optomechanical filter is therefore

$$\omega_m \gg \gamma \gg \Omega. \quad (5.9)$$

In this case, the input-output relation can be further simplified from equation 5.3 to

$$E_{\text{out}}(\Omega) \approx \frac{-\Omega + i\gamma_{\text{opt}}}{\Omega + i\gamma_{\text{opt}}} E_{\text{in}}(\Omega) + \left(\frac{\gamma}{\omega_m} \right) \frac{\gamma_{\text{opt}}}{\Omega + i\gamma_{\text{opt}}} E_{\text{in}}^*(2\omega_m - \Omega) \quad (5.10)$$

where γ is the bandwidth of the cavity and we define

$$\gamma_{\text{opt}} \equiv \frac{\hbar |g|^2}{2m\omega_m\gamma} \quad (5.11)$$

and assume that the contribution of the mechanical damping γ_m is negligibly small compared to γ_{opt} .

In the resolved sideband regime, $\gamma \ll \omega_m$, the contribution of the second term to the input-output relation is negligible since γ/ω_m is small. In this regime, the optomechanical filter can be viewed as a frequency-dependent mirror with effective reflectivity

$$r_{\text{eff}}(\Omega) = T_1(\Omega) \sim \frac{-\Omega + i\gamma_{\text{opt}}}{\Omega + i\gamma_{\text{opt}}} \quad (5.12)$$

$$\sim e^{2i\frac{\Omega}{\gamma_{\text{opt}}}}. \quad (5.13)$$

Figure 5.8 shows how the differing degrees of approximation change the expected optical response of the filter from the perspective of the lower sideband of the *pump* laser frequency, i.e. it plots $T_1(-\omega)$ as computed by FINESSE, and then as given by equations 5.3 (‘analytical’), 5.10 (‘simplified’) and the exponential form of 5.13 (‘resolved sideband’).

A cavity can then be formed using this ‘mirror’, much in the manner used in section 5.3. This cavity will have a modified transfer function

$$\frac{a_{\text{refl}}}{a_{\text{in}}} = \frac{r_{\text{in}} - r_{\text{eff}}(\Omega)e^{-2i\phi_{\text{cav}}}}{1 - r_{\text{in}}r_{\text{eff}}(\Omega)e^{-2i\phi_{\text{cav}}}} \quad (5.14)$$

and correspondingly altered bandwidth.

Figure 5.9 uses the same setup depicted in figure 5.1, with an optomechanical filter pumped by




 Figure 5.8: $T_1(-\omega)$ at various degrees of approximation.


Figure 5.9: Bandwidth broadening of a LIGO-like arm cavity using an optomechanical filter.

an off-resonance input field in place of the ‘magic’ cavity. The parameters of the filter cavity are those given by table 5.1, and the arm cavity has a length of 4 km and input mirror transmissivity $T_{\text{in}} = 7000$ ppm; this is representative of a LIGO-like arm cavity. The effective reflectivity is approximated to the level given by equation 5.10 in the resolved sideband regime. In this case the bandwidth of the reference arm cavity is $\gamma_{\text{arm}}/2\pi = 39$ Hz, and the broadened bandwidth is $\gamma_b/2\pi = 570$ Hz.

5.4 Conclusions

The quantum-limited sensitivity curve of current interferometric gravitational wave detectors has a fixed gain-bandwidth product, meaning that we must choose to either operate with higher sensitivity over a narrow frequency range, or a reduced sensitivity over a broader range. The ‘optomechanical filter’ is a proposal to overcome this compromise by using a suspended optical cavity that is dominated by radiation pressure. By carefully selecting the relative frequencies of the input field, cavity resonance, and suspended mirror, such a cavity can act as a negative dispersion element, and used to compensate the phase accumulated by the gravitational wave signal in the detector.

I have reviewed the concept of the optomechanical filter, a very new idea, and developed the theoretical framework using a different sign convention that is commonly used in the experimental gravitational wave community. I also compared the lower-order approximations used in [13] to results produced by an equivalent FINESSE model—the first numerical model of such a system—and found good agreement in the relevant parameter regime. These steps were necessary before a physical implementation of the filter, discussed in chapter 6, could be seriously considered.

Chapter 6

Development of a Tabletop Experiment Towards Optomechanical Bandwidth Manipulation

In this chapter I discuss whether an optomechanical filter capable of broadening the bandwidth of a future gravitational wave detector can be realised and outline a research programme to investigate this question.

From past and present gravitational wave detectors, we have seen that the technology development time for new schemes for improving detector sensitivity typically takes 10 to 30 years. This is due to the very strict requirements posed by gravitational wave detectors and their subsystems, in particular regarding any internal noises and the required stability during operation. Before a new scheme can be considered for extensive testing in large-scale interferometers, its implementation will be tested using a combination of numerical models, table-top experiments, and suspended scale-prototype trials.

Figure 6.1, taken from [13], illustrates the initial proposal for including an optomechanical filter in a LIGO-like gravitational wave detector by Miao *et. al.* in 2015. The optical response of dual-recycled gravitational wave detectors is shaped by the signal-recycling cavity tuning. Broadening the bandwidth of this system allows the interferometer to be operated with a tuned SRC, giving a high-sensitivity, narrowband response referred to as ‘signal recycling mode’, which is then broad-

ened by the filter to give a bandwidth similar to that achieved using resonant sideband extraction.

The filter is included directly inside the SRC via a beamsplitter and additional ‘internal SRM’ (iSRM). The iSRM is designed such that the cavity formed between the ITM and iSRM is impedance matched, so the arm cavity pole is removed and the bandwidth of the detector is just the bandwidth of the SRC.




Figure 6.1: The original proposal for implementation of an optomechanical filter in a LIGO-like gravitational wave detector, from [13]. Optical elements highlighted in blue are additional elements added to the existing dual-recycled interferometer configuration. The filter is introduced inside the signal-recycling cavity (SRC) via a beamsplitter and ‘internal signal recycling mirror’, iSRM, which is used to remove the arm cavity pole by impedance matching to the ITM.

The key features of this optical layout can be considered as shown in figure 6.2. As described in chapter 5, the optomechanical filter can be thought of as a mirror which introduces negative dispersion. This time, we also include the arm cavity of the detector, which can be treated as a mirror with an associated positive dispersion governed by the arm cavity length L_{arm} . The SRC can therefore be thought of as a single cavity, with a dispersive mirror at one end and a negatively-dispersive mirror at the other. The optomechanical filter is thus designed to compensate for the dispersion of the arm by requiring the optical damping to be

$$|\gamma_{\text{opt}}| = \frac{c}{L_{\text{arm}}}. \quad (6.1)$$

In this ideal case, the resulting sensitivity curve is as shown in figure 6.3.

More recently, it has been proposed that the filter cavity could be formed by the signal-recycling cavity itself, by using one of the turning mirrors in the SRC telescope as the suspended mirror

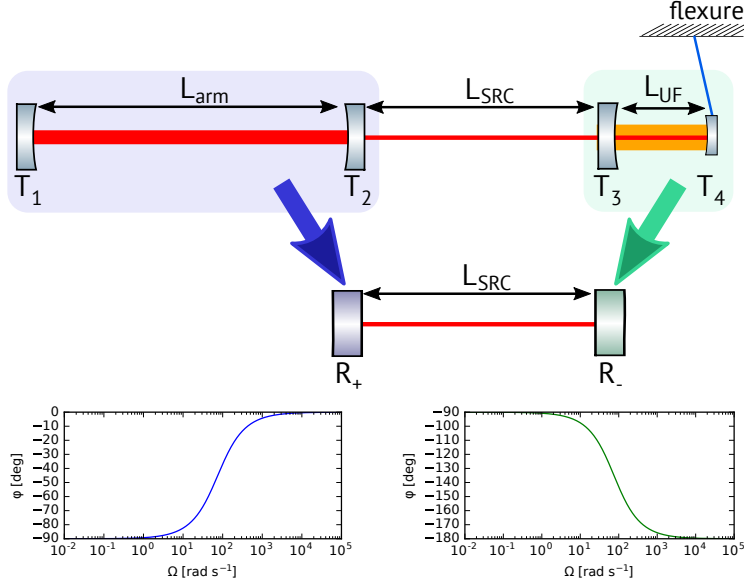


Figure 6.2: Simplified depiction of a gravitational wave detector containing an Unstable optomechanical Filter (‘UF’), which is used to compensate for the dispersion of the Arm Cavity (‘arm’), broadening the overall bandwidth of the optical system formed via the signal-recycling cavity (‘SRC’). The arm and filter cavities can be thought of as frequency-dependent mirrors for the SRC, with opposite phase relations which are designed to compensate for one-another.




Figure 6.3: The shot-noise limited sensitivity of the detector can be changed by introducing an optomechanical filter, as shown in [13]. Including the filter allows the detector to be operated in signal recycling mode resulting in a high sensitivity, narrowband response which is then broadened by the filter. Overall the bandwidth is similar to the current RSE operational mode, but with a peak sensitivity corresponding to SR operation across the band.

[95]. This has the advantage of significantly reducing the number of additional optics required to implement the scheme.

The practical implementation of an optomechanical filter has many challenging aspects. In this chapter I briefly review two key challenges: thermal noise, and the controllability of the filter. I then describe my work designing tabletop optical configurations to demonstrate optomechanical bandwidth manipulation in the classical regime, where we are not limited by thermal noise, and report the status of experimental work towards investigating control systems for unstable optome-

chanical systems.

In sections 6.1 and 6.2 I review the challenges of thermal noise and control of the unstable filter respectively. In section 6.3, I outline the frequency requirements for a tabletop experiment in order for it to be used to demonstrate bandwidth broadening, and in sections 6.4 and 6.5 develop the design for a radiation-pressure-dominated optical cavity within these requirements. Section 6.6 describes two optical layouts which could be coupled to this cavity to demonstrate its broadening effect. Finally, in section 6.7 I provide details of experimental progress so far and outline the next steps towards realising the radiation pressure facility.

6.1 Thermal Noise

For an optomechanical filter to actually broaden the bandwidth of a gravitational wave detector in the manner indicated by figure 6.3, we require the detector to be quantum-limited in the frequency range of interest. At mid to high frequencies, the biggest noise sources are (1) quantum shot noise and (2) phase noise due to Brownian thermal noise in the optics, of which shot noise dominates. We thus require that the additional thermal noise introduced by the optomechanical filter is lower than the detector shot noise. This leads to the requirement [13]

$$\frac{8k_B T}{Q_m} \lesssim \hbar \gamma_{\text{det}}, \quad (6.2)$$

where k_B is the Boltzmann constant, T is the temperature of the environment, Q_m is the mechanical Q-factor of the suspended mirror, and γ_{det} is the bandwidth of the detector without the unstable filter. In the configuration shown in figure 6.1, $\gamma_{\text{det}} = \gamma_{\text{SRC}}$, since the addition of the iSRM means that the detector bandwidth is entirely determined by the SRM. As an order of magnitude estimate, we therefore have

$$\frac{T}{Q_m} \lesssim [1 \times 10^{-12} \text{sK}] \gamma_{\text{det}} \sim [6 \times 10^{-10} \text{K}] \frac{\gamma_{\text{det}}/2\pi}{[100 \text{Hz}]} . \quad (6.3)$$

For a LIGO-like detector with γ_{det} of order kHz, if we assume that the oscillator has a relatively low Q-factor of order $Q_m \sim 100$ we require that the unstable filter be cooled to a sub-nK level. Q-factors exceeding 10^9 have been measured using materials such as cryogenic silicon [96], meaning the temperature requirements are eased to mK but still below liquid helium.

This requirement poses a significant challenge, and it is currently not yet known whether a practical solution will be found. In addition to materials research more exotic schemes are being explored

such as using optical dilution to replace the material restoring force with an optical restoring force to reduce damping [97]. This is an ongoing area of research in the community.

6.2 Controlling unstable systems

The anti-damping term in the optomechanical filter, necessary for it to act as a negative dispersion element, also means that the resulting system is intrinsically unstable. It is therefore essential that the system is stabilised using an external control loop, such that the filter itself retains its critical dispersion properties without driving the complete system to instability.

Whether such a control scheme is possible can be determined by evaluating the observability and controllability matrices for the system. An introduction to these and other general control concepts is provided in appendix E. Using this approach, it was first understood by Miao *et. al.* that the 3-mirror system, such as shown in figure 5.6, would be controllable [13]. The control signal was assumed to be read out on reflection from the stable arm cavity, and approximated as 3 degrees of freedom: the position of the suspended mirror, the amplitude of the probe field in the arm cavity, and the time derivative of this amplitude. It was also assumed that:

1. the filter cavity length is negligible compared to the arm cavity, and thus does not introduce any time delays.
2. the filter is operated in the resolved sideband regime and therefore the sideband generated at $\omega_0 + 2\omega_m$ is negligibly small.
3. the acceleration of the mirror is similarly negligible.
4. there is no time delay between the control signal as generated immediately on reflection from the unstable filter, and the actuator acting on the suspended mirror.

Under these many assumptions, it was found that the system was indeed controllable and observable, and a controller could be designed. However, additional time delays in the system (assumptions 1 and 4) will reduce the system stability, while including the effects of the mirror acceleration and upper sideband (assumptions 2 and 3) may result in the system becoming uncontrollable at high frequencies.

Theoretical work on this topic has so far concentrated on describing the behaviour of an unstable optomechanical filter and whether, once coupled to a stable arm cavity, the resulting system can

be controlled. Future work on this topic will consider the more general question of whether *any* active negative dispersion element is fundamentally observable or controllable in such a way that it still broadens the bandwidth of a detector, as has already been ruled out for passive negative dispersion elements [89].

In parallel to these theoretical studies, we are developing our tabletop facility to test the feasibility of bandwidth control using an optomechanical filter. The aims of the experiment discussed through the remainder of this chapter are:

1. Set up a radiation-pressure-dominated tabletop facility at the University of Birmingham, including a fully digital control scheme for the stand-alone optomechanical cavity. Use this to experimentally demonstrate optomechanical parametric amplification.
2. Develop the control system for the full coupled cavity system to demonstrate bandwidth broadening, testing schemes as these emerge from theoretical discussions.

The first of these goals is known to be achievable and is the foundational step for the Birmingham facility. As the power circulating in gravitational wave detectors is increased, we anticipate radiation pressure effects to become increasingly prevalent. Therefore if theoretical studies find that bandwidth broadening using the filter is not possible, the setup can be used to explore many other features of radiation-pressure-dominated optical systems.

6.3 Frequency Requirements for a Tabletop Optomechanical Filter

In this and the following sections, I develop the design for a tabletop experiment intended to demonstrate optomechanical bandwidth manipulation. This is based on the conceptual coupled optical system described in chapter 5, and it relies on strict frequency relations which I now combine with the related practical requirements for the experiment to work on a tabletop scale.

Equation 5.10 indicates that the bandwidth of the Filter Cavity, γ_{FC} should be small relative to the mechanical frequency, ω_m so that we have negligible contributions from the upper sideband, which will degrade the negative dispersion effect. At the same time, we need to be able to control the final system. We intend to use a fully digital control scheme, and anticipate an achievable control bandwidth of $\Gamma_{\text{control}} \sim 50 \text{ kHz}$. The initial tabletop experiment will be operated at room

6.3. FREQUENCY REQUIREMENTS FOR A TABLETOP OPTOMECHANICAL FILTER

temperature, and use strong test signals such that thermal noise does not mask the key behaviours of interest. These statements set limits on ω_m and γ_{FC} :

$$\gamma_{\text{FC}} \ll \omega_m < \Gamma_{\text{control}}. \quad (6.4)$$

To observe a broadening effect we then require that the bandwidth of the test system, γ_{T} when no optomechanical filter is in operation, is

$$\gamma_{\text{T}} < \gamma_{\text{FC}}, \quad (6.5)$$

This also ensures that pump light at frequency $\omega_p = \omega_0 + \omega_m$ is off-resonance for both the filter cavity and test system.

The mechanical damping of the bare suspended mirror always has a positive value and thus acts to dilute the optical anti-damping responsible for producing negative dispersion (see equation 2.44). We therefore want to use a high Q-factor suspension so that this mechanical damping $\gamma_m = \omega_m/Q_m$ does not limit the system. This is also advantageous in terms of thermal noise.

As shown in equation 2.43, the resonant frequency of an optical spring is also shifted compared to the resonant frequency of the bare mechanical oscillator. When the detuning is large, this frequency shift is small, while the anti-damping effect can still be significant. A large detuning, and therefore a large bare mechanical resonant frequency, is thus preferable both in terms of the resulting negative dispersion and for control, since in that case the detuning that the pump laser requires is not dependent on the optical power (and associated noise) in the system. However, if the mechanical resonance is smaller it may be possible to use a higher input power to shift its frequency closer to the resolved sideband regime, improving the negative dispersion effect.

In addition to the specific needs of this experiment, other requirements can be drawn based on our experience working with tabletop cavities. While the bandwidth of the test cavity should be small, if it is very small it becomes difficult to control and align. Other noise sources must also be considered as the experiment is constructed, for example the effects of air damping and resulting vacuum requirements.

6.4 Design of the Suspended Mirror

Based on the frequency requirements described in section 6.3, we propose a mechanical oscillator with a resonant frequency of $\omega_m/2\pi \approx 10\text{ kHz}$ and Q-factor of $Q_m \approx 10^5$. The filter cavity bandwidth requirement (equation 6.4) then imposes restrictions on the length of the FC and transmissivity of its two mirrors (as in equation 2.16).

In addition to these fundamental requirements, we must also consider that a real mechanical oscillator will have many higher order resonances (as described in section 2.6, and which result in parametric instabilities as studied in chapters 3 and 4). If these are close in frequency, the optomechanical interaction between competing mechanical modes may disrupt the negative dispersion effect, or affect control of the system. We therefore also require that the mechanical mode separation frequency is greater than the bandwidth of the filter cavity, and that the higher order mechanical modes do not also meet the conditions for parametric instability (see, for example, equation 2.59).

Suspended mirrors for optical spring and radiation pressure experiments have been developed by several groups worldwide. This includes university-based groups such as the University of Western Australia, where they are actively working to develop oscillators with low thermal noise and good mode separation, suitable for use in optomechanical filters in the quantum limit [98]. It also includes specialised companies such as Crystalline Mirror Solutions, led by Garrett Cole, who has developed aluminium gallium arsenide (AlGaAs) as a high-quality optical coating. AlGaAs can also be used to fabricate chips of small micro-oscillator mirrors, such as used in the recent first observation of broadband quantum radiation pressure noise at Louisiana State University (LSU) [99].

AlGaAs is particularly interesting as a coating material due to its extremely low thermal noise properties and correspondingly high mechanical Q-factors [100, 101]. The reflectivity of the coating can also be very high, depending on the number of layers used. This makes AlGaAs a material of interest for future gravitational wave detectors, as well as an appropriate choice for our project. The resonant frequency is then dictated by the mass and geometry of the mirror and its suspension.

Figure 6.4 shows a typical example of an AlGaAs chip of micro-oscillators, developed by Cole for Corbitt, Cripe *et al.* at LSU. As shown each chip typically contains around 100 suspended mirrors. The total size of the chip is approximately $5\text{ mm} \times 5\text{ mm}$; the mirror resonant frequencies are in the range 0.1-1 kHz.




Figure 6.4: Example micro-oscillator chip and mounting developed by G. Cole and used by T. Corbitt, J. Cripe *et. al.* at LSU, and kindly lent to the University of Birmingham.

Each individual mirror has many degrees of freedom which can be used to tune the fundamental and higher order resonant frequencies. As shown in figure 6.4, the most common suspended mirror design is a ‘lollipop’ shape consisting of a round mirror suspended on a single ‘leg’. On a given chip, the maximum material thickness is fixed, therefore the mass of each mirror is determined by its radius. The resonant frequency can then be further tuned by changing the geometric properties of the leg: its length, width, and thickness. A larger number of legs can also be used to increase the resonant frequency further.

Compared to the mirrors in the example shown above, we require at least an order of magnitude increase in mechanical frequency. Through finite element modelling, we have found that a four-legged mirror such as that depicted in figure 6.5 might be able to achieve our target frequency range; example geometries are listed in table 6.1.

Given the large number of mirrors that can fit onto it, one chip can contain mirrors of many different




Figure 6.5: 4-legged mirror design that can be used to achieve resonant frequencies of order 10 kHz

configuration	leg length [μm]	$\omega_m/2\pi$ [kHz]
1 leg	500	3.7
2 legs at 180°	500	13.6
4 legs	500	19.3
4 legs	400	25.6
4 legs	300	37.5

Table 6.1: Mechanical resonant frequencies resulting from various potential mirror geometries. In all cases listed, the mirror diameter is 600 μm, and the leg width is 100 μm. The material is layered AlGaAs, with overall thickness 6.7 μm.

frequencies and geometries. We can therefore use this as an opportunity to explore various different mirror geometries. Care must be taken to ensure that the resonant frequencies of different mirrors do not cross-couple; this can also be ensured using finite element and CAD models of the total chip design.

6.5 Geometric Design of the Filter Cavity




Figure 6.6: Geometry of the Filter Cavity. This cavity is near-hemispherical. The concave input coupler has radius of curvature R_c and focuses the beam down onto the small, flat suspended mirror with diameter D . To minimise clipping losses, the beam spot radius on each mirror should be $w \leq D/5$.

Figure 6.6 shows a schematic diagram of the Filter Cavity. The suspended mirror, M_4 , is a small (sub-mm diameter) disc of AlGaAs suspended from a larger chip which can be clamped securely, as described in section 6.4. Due to the manufacturing method used, this mirror must be flat (infinite radius of curvature). The input coupler, M_3 is a fixed, macroscopically scaled mirror with a relatively tight (of order 10 cm) radius of curvature in order to focus the beam down on to M_4 .

The geometric design of the Filter Cavity is driven primarily by the bandwidth requirement (equation 6.5). The target of $\gamma_{FC}/2\pi \lesssim 10$ kHz is relatively low and can be achieved through a combination of low transmissivity mirrors and large cavity length, as shown in equation 2.16. However,

6.5. GEOMETRIC DESIGN OF THE FILTER CAVITY

the cavity length is limited by the diameter of M_4 , which must be small to achieve high ω_m . Since the Filter Cavity is near-hemispherical, the beam waist will sit at M_4 . Equation 2.49 shows that the radius of the beam waist varies as the square-root of the length ($z - z_0$), and so shorter cavity lengths must be used. To minimise clipping losses, we require that the radius of the beam spot on each mirror is smaller than 1/5th of the mirror diameter (or vice versa), which corresponds to clipping losses of less than 1 ppm per optic.

The beam waist size is also controlled by the radius of curvature of the input coupler, R_C . In this tightly focussed regime, $R_C \approx L$. However, in the case that $R_C = L$, the cavity becomes geometrically unstable, as shown by equation 2.50. This results in the cavity becoming difficult to align and control. We therefore require that the Filter Cavity is geometrically stable, with a g-factor $0.05 < g < 0.95$.

Figure 6.7 plots the resulting waist size as a function of cavity length for a selection of input coupler radii of curvature, modelled analytically using equation 2.49. Here we see that there are many possible ways of achieving the same final spotsize. The grey line marks the threshold at which the cavity stability reaches $g = 0.05$; values above this line are more stable. This line also corresponds to the case with the largest allowed value of L for a given R_C and w , and so marks a locus of minimum-bandwidth cavity designs.




Figure 6.7: Radius of the beam waist (i.e. spot size on the suspended mirror) as a function of cavity length for a selection of input coupler radii of curvature. For a given choice of R_C , two length choices can be used to achieve the same resulting spot size. We choose the longer length, since this results in a narrower cavity bandwidth. The grey line marks a threshold for the geometric stability of the cavity: values *above* this line have stabilities greater than $g = 0.05$ and should therefore be easier to control and align. The final bandwidth of the cavity then depends on the transmissivity of the two mirrors.

The final bandwidth of the cavity is then fully given by equation 2.16. For a given choice of mirror transmissivities and target bandwidth, we therefore have a design for our cavity, since this determines our choice of L , which in turn sets w and R_C via figure 6.7. We then perform minor adjustments to R_C —while maintaining $g \sim 0.05$ and clipping losses below 1 ppm—so that R_C is a round value that can be readily ordered. If the resulting diameter of M_4 is not practical (for example, the dimensions do not allow the target ω_m to be reached, or it cannot be manufactured using established techniques), we restart the design process using adjusted values of T_3 and T_4 so that a different value of L can be used, resulting in a different spot size.

Table 6.2 lists the full resulting set of parameters used for the Filter Cavity modelled throughout this chapter, which is one representative case. Mirror transmissivities are representative of those used by groups performing similar experiments. Using a chip containing many suspended mirrors with different geometries, it will be possible to try many possible combinations. Since in this case we will use the same single input coupler, we can choose to either use a spot size and cavity length suited to the smallest diameter mirror (this will correspond to the broadest cavity bandwidth), or adjust the cavity length to follow one of the lines of constant R_C to achieve the minimum bandwidth for each individual mirror.

Parameter	Symbol	Value
input power	P_{in}	250 mW
mechanical resonant frequency	$\omega_m/2\pi$	25.6 kHz
diameter of M_4	D	600 μm
effective mass of M_4	m	6.2 μg
mechanical Q-factor	Q_m	10^5
transmission of M_3	T_3	50 ppm
transmission of M_4	T_4	350 ppm
spot size on M_4	w	120 μm
RoC of M_3	R_C	20 cm
geometric stability	g	0.05
length of FC	L_f	19.05 cm
bandwidth of FC (HWHM)	$\gamma_f/2\pi$	25.1 kHz

Table 6.2: Geometric and optical parameters of the proposed Filter Cavity.

6.6 Optical Layouts

The filter cavity design above is intended to use parameters we think are achievable in the lab. As a result, it is only marginally in the resolved sideband regime, and several other parameters mean that the approximations used to generate equations 5.10 onwards may begin to break down. A

detailed numerical model of the optical system, without these approximations, should therefore be used to complete the experimental design. This is currently under development in FINESSE¹.

Here, I develop the overall topology of the experiment and suggest initial parameters by designing the system around the filter cavity assuming such approximations are still reasonably valid. As such, if the system was built using three mirrors total (such as shown in figure 5.6) the ideal length of the stable ‘arm’ cavity is given by equation 6.1. For the parameters of the filter design in table 6.2, this means

$$L_{\text{arm}} = \frac{c}{|\gamma_{\text{opt}}|} \sim 2.88 \text{ m}. \quad (6.6)$$

The end mirror reflectivity can then be chosen such that the arm cavity bandwidth meets the frequency requirements outlined in section 6.3.

In the following sections, I outline two possible schemes which couple a stable optical system to the unstable filter designed above. The first design proposes using an optical fibre for the stable cavity. This allows the stable cavity’s length to become very long, at the expense of introducing more optical loss. The second layout instead uses a triple-cavity system, such that the accumulated phase in the stable system is scaled by a factor of the finesse of one of the cavities. This allows the system to become more compact while remaining in free-space.

6.6.1 Fibre Cavity System




Figure 6.8: Schematic layout of a fibre-based cavity system to demonstrate bandwidth broadening.

An optical layout designed to demonstrate bandwidth broadening of a fibre cavity is depicted in figure 6.8. This has the advantage that the fibre cavity can be far longer than a free-space cavity, and therefore the optical system remains similar to the conceptual designs in chapter 5.

The description for a propagating field (see equation 2.7) can be updated to include the additional

¹While FINESSE is able to model the optomechanical filter itself, coupling between sidebands associated with different carrier fields is usually sufficiently small to be ignored in optical models, simplifying and speeding up the computation. In this unusual system, however, it becomes the dominant process, and so such couplings are planned to be including in the future.

losses resulting from using a fibre by describing the electric field as

$$E_{\text{out}} = \sqrt{(1 - \mathcal{L})^D} e^{-i\phi} E_{\text{in}} \quad (6.7)$$

where \mathcal{L} is the fractional power lost per unit length, D is the propagation distance, and $\phi = knD$ is the phase accumulated, with $k = \omega/c$ the wavenumber of the optical field. Note that since we propagate in fibre, rather than the usual assumption of propagation in vacuum, we must now include the refractive index, n when referring to the optical path length. This means that the optical length of the fibre, rather than its physical length, should be matched to the optomechanical filter. This may be significantly different, with consequences for the total in-fibre losses.

As in section 2.3, we can then propagate the fields to find the circulating power and corresponding bandwidth of a cavity with non-unity refractive index and length-dependent losses:

$$\gamma_{\text{fibre}} = 2\pi \times f_{\text{pole}} = \frac{c}{nD} \arcsin \left(\frac{1 - (1 - \mathcal{L})^D r_1 r_2}{2\sqrt{(1 - \mathcal{L})^D r_1 r_2}} \right). \quad (6.8)$$

For the fibre cavity to be useful, we require its bandwidth to be less than that of the filter cavity ($\gamma_{\text{FC}} = 25.1$ kHz). Figure 6.9 plots the bandwidth as a function of cavity length for a free-space cavity compared to cases with length-dependent losses, comparing equation 6.8 to equation 2.16. In this example, the fibre is assumed to have a refractive index of $n = 1$, and to introduce no additional loss sources (such as coupling loss). The additional dependence on the cavity length means that the bandwidth of a fibre cavity is limited, and asymptotes to a minimum value dependent on \mathcal{L} before increasing as further loss is introduced. For the bandwidth of the fibre cavity to be lower than γ_{FC} , we see that the loss per unit length in the fibre must be extremely low, of order 100 ppm per metre. Increasing the reflectivity of the input mirror does not change the minimum bandwidth, and changes the length for which this bandwidth is achieved. This would mean that it was no longer ideally designed for the optomechanical filter to broaden.

Typical off-the-shelf 1064 nm fibres have losses that are higher than this. Common values are 2.5 dB/km \sim 400 ppm per metre (e.g. Thorlabs PM980-XP fibre [102]) or higher. 1550 nm fibres, which are more common, have been made with losses of 0.4 dB/km \sim 90 ppm per metre [103], making the design more achievable. Coupling into and out of the fibre cavity is also typically lossy. Both of these properties would need to be optimised for a fibre cavity to be used in our experimental setup.




Figure 6.9: Bandwidth of fibre cavities compared to a free space cavity of equivalent length. Length-dependent loss in fibres means that fibre cavities have a minimum bandwidth. Changing the reflectivity of the end mirrors of the cavity only changes the length at which this minimum bandwidth is achieved, and for what range of lengths it is maintained. In this example the end mirror of the fibre cavity is the input mirror of the optomechanical filter with transmission $T_3=50$ ppm. To meet the requirement that the fibre cavity bandwidth is less than the filter cavity, it must have less than 10^{-3} loss per unit length.

6.6.2 Triple Cavity System

Figure 6.10 shows a free-space optical configuration designed to demonstrate bandwidth broadening of the central cavity. This is a coupled system of 3 cavities, serving in the following roles:

- **Test Cavity** - a low-dispersion cavity where bandwidth measurements will be made.
- **Arm Cavity** - acting as a highly dispersive ‘mirror’. This is the main source of dispersion that we wish to compensate.
- **Filter Cavity** - the optomechanical filter.

This design is based on the scheme depicted in figure 6.2. It is therefore similar to the initial design for implementation in gravitational wave detectors from [13] described at the start of this chapter. The test cavity in this case takes the role of the signal-recycling cavity.

Inclusion of the arm cavity means that the overall length of the optical system can be much less than a simpler two-cavity design. Figure 6.11 illustrates the difference. When the stable optical system is formed from two mirrors, the difference in phase between the fields incident on and directly reflected from the end mirror is simply the propagation time in this cavity, which can be described using the free spectral range (FSR). The unstable filter should therefore be designed




Figure 6.10: Proposed layout of a triple-cavity free-space optical system designed to demonstrate optomechanical bandwidth manipulation. The arm cavity and filter cavity act as sources of positive and negative dispersion respectively, shaping the overall bandwidth of the central test cavity. Additional ‘input’ and ‘output’ mirrors enable the pump, probe, and control fields to be injected into the system and test cavity bandwidth to be measured. Finally, a third highly-reflective mirror in the test cavity is used as the actuator for the test cavity length control. This results in a ‘W’-shaped design which is scaled to fit inside our vacuum tank.




Figure 6.11: Phase accumulation in 2-mirror versus 3-mirror cavity systems.

such that

$$-e^{-\frac{2i\Omega}{\gamma_{\text{opt}}}} = -e^{-2i\phi} = -e^{-i\frac{\Omega}{\text{FSR}}}, \quad (6.9)$$

where $\phi = \tau\Omega$ as usual for a tuned cavity, so we require $\gamma_{\text{opt}} = L/c$. If instead the stable optical system is a coupled cavity formed from three mirrors, the difference in phase takes a more complicated form due to the additional propagation in the second cavity. Equation 2.18 describes the input-output relation for a two mirror cavity. Through Taylor expansion, we can further approximate this relation:

$$\frac{a_2}{a_1} \approx \frac{-\Omega - i\gamma}{-\Omega + i\gamma} \sim -e^{-2i\frac{\Omega}{\gamma}} = -e^{-i\frac{\Omega}{\text{FSR}}\frac{2\mathcal{F}}{\pi}}, \quad (6.10)$$

where γ is the bandwidth of cavity B. The phase reflected from the cavity is approximately scaled by the cavity finesse, \mathcal{F} . The fields at the end mirror are then related by

$$a_i = -a_o e^{-2i\phi_B} \left(r_2 + e^{-2i\phi_A \frac{2\mathcal{F}_A}{\pi}} \right). \quad (6.11)$$

The relative phase between the two fields is approximately scaled by the number of round-trips that light makes in cavity A. When cavity A has a high finesse, and correspondingly a much smaller bandwidth than cavity B, the phase accumulated in B becomes negligible. In this case the unstable filter should be designed such that

$$-e^{-\frac{2i\Omega}{\gamma_{\text{opt}}}} \sim -e^{-2i\phi_A \frac{2\mathcal{F}_A}{\pi}} = -e^{-i\frac{\Omega}{\text{FSR}_A} \frac{2\mathcal{F}_A}{\pi}}. \quad (6.12)$$

In this manner, we can replicate the high dispersion of a long cavity using a shorter, high-finesse cavity. The entire setup can then be scaled to fit inside a 60×50 cm vacuum tank which is already present in the group research laboratories. Proposed parameters are listed in table 6.3.

The unstable filter produces the negative dispersion response described in section 5.3 on reflection, therefore an input beamsplitter with transmission T_i is introduced in the test cavity, through which the high-power pump laser is injected via a window in the vacuum tank. We can also inject the low-power control laser field here, either by phase locking two lasers (see, for example, [78]), or using Acousto-Optic Modulators to shift part of the pump laser to the control frequency (the resonant frequency of all cavities, ω_0) or vice versa, e.g. as in [32].

Introducing the input beamsplitter also provides an additional port, which could be used in the control scheme. We also include a second output beamsplitter of transmissivity T_o , which can

Parameter	Value
<i>Arm Cavity</i>	
Length, L_A	40 cm
Mirror transmissions: T_1	10 ppm
T_2	0.4
Finesse, \mathcal{F}_A	12
<i>Test Cavity</i>	
Length, L_C	27 cm
Mirror transmissions: $T_i T_o$	550 ppm
<i>Filter Cavity</i>	
Length, L_F	19.05 cm
Mirror transmissions: T_3	50 ppm
T_4	350 ppm
Bandwidth, γ_F	25.1 kHz

Table 6.3: Parameters proposed for a triple-cavity free-space bandwidth broadening experiment

be used to measure both error signals for the control of the cavity system, introduced by phase modulating the control laser using electro-optic modulators at radio frequencies, and to measure the bandwidth of the system. Error signals at this location will be affected by all three cavities.

We envisage using a hierarchical control scheme. The arm cavity has a narrow bandwidth and therefore reflects all of the pump field. The control laser can therefore be locked to the arm cavity by looking at the error signal transmitted or reflected from the arm. The control laser can then be used to lock both the test and filter cavities on resonance. For this reason, a third, highly reflective mirror has been included in the test cavity, which can be actuated on away from points of error signal extraction. Finally, the difference in frequency between the control and pump lasers must be set to the mechanical resonance of the suspended mirror, ω_m . This is discussed in section 6.2.

6.7 Current Experimental Status: Chip Characterisation

Work has begun to develop a radiation-pressure-dominated tabletop optical facility in our lab. The experimental goal for the first stage of the project is to build a cavity containing a suspended mirror, produce an optical spring, and control it. This is the foundational step necessary for the group.

Current progress has focused on procurement and characterisation of the AlGaAs mirror chip shown in figure 6.4, which has been kindly donated to our group by T. Corbitt, J. Cripe *et. al.* from LSU. The micro-oscillators on this chip have resonant frequencies far lower than we would like for an experiment to demonstrate bandwidth broadening, however we are using this one first

to become familiar with handling the chip and to achieve our primary experimental goals, which do not have the same strict frequency requirements as described in section 6.3.

We first need to understand the behaviour of the micro-oscillators on the chip when they are not driven by radiation pressure. In particular, we wish to measure each mirror's mechanical Q-factor and resonant frequency. This is achieved by using a simple Michelson to measure the in-axis displacement of the mirror after a short excitation was applied. We have developed the setup in-air to perform initial tests.

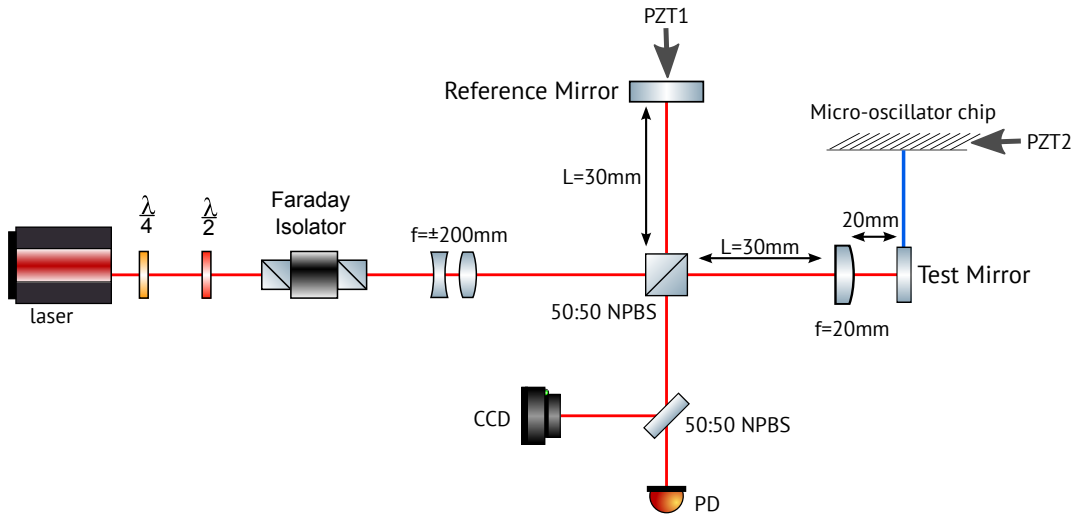


Figure 6.12: Schematic diagram of the optical layout used to characterise mirrors on the AlGaAs chip.

The core optical layout is shown schematically in figure 6.12 and photographed in figure 6.13. A quarter-wave-plate ($\lambda/4$) linearises the polarisation of the input beam from the 500 mW 1064 nm laser. A half-wave-plate ($\lambda/2$) is used in conjunction with a faraday isolator to prevent back-reflection into the laser. We also use this to attenuate the input beam power, together with several low-reflectivity turning mirrors in the input path (not depicted), producing a final input power of ~ 10 mW. This is sufficiently low that the radiation pressure force on the test mirror will be negligible. A pair of ± 200 mm lenses are used to collimate the beam, resulting in a spot size of approximately 1 mm.

The Michelson interferometer is formed from a 50:50 non-polarising beamsplitter (NPBS), a reference arm with a high-reflectivity reference mirror, and the test arm containing the micro-oscillator chip. The arms of the Michelson interferometer are identical in length and short, so that it is easy to align and the resulting interference pattern has a fundamental Gaussian intensity distribution. The smallest micro-mirror on the chip has a diameter of $50 \mu\text{m}$, therefore to be able to characterise




Figure 6.13: Core Optics of the in-air Michelson setup. The mirror shown in the test arm was replaced with the mirror chip, mounted on a 3-axis manual translation stage, once the setup had been pre-aligned. The optical path is marked in red.

all mirrors, the spot size on the mirror should be significantly smaller than this; ideally $w \sim 10 \mu\text{m}$. We achieve this by including a 20 mm lens in the arm of the test mirror and aligning the chip so that it sits exactly at the focal point. This results in a ‘2f system’ in the test arm that means the beam spot size returning to the beamsplitter is unaffected by the lens. It also results in the Michelson becoming immune to angular misalignment of the test mirror; this feature is used to ensure that the mirror is at the focal point.

A second NPBS is used to split the output of the Michelson so that it can be measured using both a photodiode (PD) and CCD camera. The photodiode is used for our final measurements; the CCD is useful for determining if the chip is correctly aligned in the direction transverse to the beam axis, since if the beam is clipped by hitting the edge of one of the mirrors, its shape will be distorted.

The chip is mounted on a 3-axis manual translation mount, formed from three linear stages, so that different micro-oscillators can be selected and the distance between the chip and focusing lens fine-tuned. Both the reference mirror and chip are also mounted to linear piezoelectric crystals (PZTs). PZT2, mounted to the chip, is used to excite the micro-mirrors by driving the chip near their expected resonant frequencies. PZT1 is used to tune the Michelson so that it sits at the middle of a fringe. The motion of the test mirror in the beam axis is smaller than the laser wavelength, therefore the voltage measured on the photodiode is directly proportional to the displacement of the mirror. Finally a ‘ring-down’ measurement is made by cutting the input excitation at PZT2 and recording the amplitude of the mirror displacement as it returns to its un-excited state.




Figure 6.14: In-air ‘ringdown measurement’ of micro-mirror ‘A’ (see figure 6.4). The raw data is fit to a decaying sinusoid with tight constraints, based on well known values for the mirror resonant frequency f_0 , initial amplitude a and offset c , resulting in a Q-factor of 50 ± 12.5 .

Figure 6.14 shows the result for one of the largest mirrors on our chip, marked ‘A’ in figure 6.4, which has a diameter of $400\ \mu\text{m}$. The resonant frequency is known to be approximately 320 Hz, and the initial amplitude and offset can be well identified, therefore a decaying sinusoid was fit to the data using tight constraints, resulting in a mechanical Q-factor measurement of $Q_m = 50.0 \pm 12.5$. This is very low compared to the order- $10^4 - 10^5$ values expected of AlGaAs micro-oscillators, but not unexpected. It is highly likely that this is caused by significant air damping. As planned, we are currently migrating the setup to our vacuum tank so that these measurements can be repeated, and extended to all micro-oscillators of interest for our experiment.

6.8 Future Work

The next immediate steps are to characterise the micro-mirrors in-vacuum, measure their mechanical transfer functions, and then to construct an optical cavity containing the chip using the $T = 50\text{ppm}$, $R_c = 10\text{cm}$ input coupler provided with it. This can start by using a low input power for the initial alignment of the cavity so that radiation pressure effects are negligible. This is expected to be challenging due to the cavity’s small length and high finesse. A control scheme can then be developed for the cavity. The setup can then be used to achieve the first experimental goal of demonstrating parametric amplification using an optomechanical system.

The second goal of the experiment—development of a control scheme for a coupled cavity system

including an unstable filter, and demonstration of optomechanical bandwidth broadening—will depend on the outcome of theoretical studies regarding the controllability of the system, numerical models to explore potential experimental designs in detail, and experimental tests of potential control schemes.

Key milestones in this process could therefore include:

- An experimental demonstration of parametric amplification and control of a radiation-pressure-dominated system in different stability regimes.
- A ‘go/no-go’ statement resulting from theoretical studies on the fundamental controllability of systems containing active negative dispersion elements.
- A full numerical model of a minimal optical system containing an optomechanical filter coupled to a stable arm cavity using minimal approximations, showing that bandwidth broadening can be achieved prior to control considerations.
- A control scheme proposal for the minimal optical system, to be tested numerically.
- An updated optical design for a tabletop-scale experiment to demonstrate bandwidth broadening using parameters informed by a full numerical model, including full geometric considerations and noise budget. This should also take into account developments in thermal noise studies.
- Construction of the tabletop configuration in-vacuum.
- A control scheme proposal for the tabletop experiment, to be demonstrated experimentally.
- A measurement of radiation-pressure-driven behaviour in the tabletop experiment, including demonstrating manipulation of the optical system’s bandwidth.

Several of these tasks can occur in parallel.

6.9 Conclusions

The implementation of any new idea to improve the sensitivity of gravitational wave detectors typically takes several decades due to the rigorous requirements it must meet. The optomechanical filter, intended to broaden the bandwidth of future interferometric detectors, faces many challenges. In particular, it has very strict thermal noise requirements, and it is not yet known whether the

overall optical system including the filter is controllable.

We have begun to establish a radiation-pressure-dominated optical facility in our labs, and procured an AlGaAs micro-oscillator chip. We have built an optical setup that can be used to characterise these oscillators through ring-down measurements. This has been successfully demonstrated with an initial in-air test. The measured Q -factor of the oscillator used in the test is far lower than expected of AlGaAs ($Q \sim 50$ rather than $Q \sim 10^5$); this is attributed to air damping. As planned, the next step for the facility is to repeat these measurements in-vacuum. The first major milestone for the facility will be to demonstrate optomechanical parametric amplification using the chip.

In the longer term, the facility is intended to test potential control schemes for a coupled optical system including the optomechanical filter. This will be used to explore the practical issues associated with the setup, and developed in conjunction with theoretical and numerical models of the system. I have developed a design for a tabletop optomechanical filter that can be used to demonstrate negative dispersion. Based on this cavity design, I have proposed optical configurations which can be used to verify bandwidth broadening in the classical regime and test control schemes. The first uses a long fibre-cavity to achieve a narrow initial bandwidth. This bandwidth will be limited by the loss-per-unit-length of the fibre; for the proposed design this loss should be less than 100 ppm/m, which is lower than is typically available commercially. The second design instead uses free-space optics, achieving the required effective length instead by introducing an additional coupled cavity.

Finally, I have described the key milestones of a research programme towards a conclusive ‘go/no-go’ outcome for optomechanical bandwidth broadening. This statement could emerge as a result of fundamental controllability studies, unsolvable thermal noise requirements, or due to practical issues meaning any proposed scheme is unworkable as part of a gravitational wave detector. Alternatively, if such a scheme is shown to be practically achievable, it could provide a relatively low-cost route to broadening the bandwidth of future detectors, allowing exciting new scientific measurements—such as measurements of the neutron star equation of state—to be carried out.

Chapter 7

Summary and Outlook

Radiation pressure plays an increasingly significant role in the performance of gravitational wave detectors. This is the central theme of my thesis.

I have focused on improving the gravitational wave community’s understanding of parametric instability (PI), and in particular the role of the interferometer configuration. I have studied the behaviour of parametric instabilities in a complete model of the Advanced LIGO interferometer core optics. Common thinking previously concentrated on the optical response of mechanical modes in the arm cavities of the detector, where the instabilities originate. I demonstrated that the power- and signal-recycling cavities have a significant effect on the parametric gain—the figure of merit for determining whether a mechanical mode will result in a PI. PIs thus depend on a significantly expanded parameter space.

By increasing the complexity of the interferometer in stages, I identified that PIs in Advanced LIGO could be separated into ‘common’ and ‘differential’ modes, depending on whether instability happened through the power- or signal-recycling cavities respectively. Including dual-recycling approximately doubled the number of mechanical modes that might be expected to result in PIs when compared to an arm cavity model. Using the example of the signal-recycling cavity, I then showed that changing the tuning or accumulated Gouy phase in the recycling cavities will change the number of expected PIs. This work is published in my paper ‘The influence of dual recycling on parametric instabilities at Advanced LIGO’ (2017) [1].

The techniques I have described can be used as a method to characterise parametric instabilities in the LIGO detectors in a way that hasn’t been possible with other more general models. We

have developed FINESSE models for each of the gravitational wave detectors which use parameters that are matched to those measured at each site. These could be used to optimise the operational mode of the detector to minimise PI effects.

I have also highlighted the potential opportunities and consequences for PIs that may be created by future upgrades to the LIGO detectors. Proposed plans will also have consequences for PIs that should be taken into account when considering the PI mitigation strategy. Careful choices of PRC and SRC accumulated Gouy Phases can be used to reduce the number of PIs, while a change in SRC tuning, in order to modify the shape of the sensitivity curve, will change which mechanical modes require mitigation.

Future detectors, such as the Einstein Telescope (ET), can be designed taking into account PIs. I have demonstrated how several key design decisions will affect PIs in ET. I first compared simple finite-element models for the test masses suggested for the ‘ET-HF’ and ‘ET-120K’ detector designs, finding that while both mirror designs have a similar total mass, different material choices and radius-to-thickness ratios mean that the ET-HF mirrors have a much higher mechanical mode density, which could be expected to correspond to a higher number of potential PIs. I then compared interferometer options for ET-120K, considering cases that are optimised to (a) minimise clipping loss, (b) reduce coating Brownian noise (at the expense of shot noise), and (c) specifically target gravitational wave signals in the 1-4 kHz frequency band, which is of particular interest for exploring neutron star physics. Due to the high circulating power and high mass mirrors in ET-120K, any design can be expected to experience more PIs than have been observed at LIGO. However, these three designs demonstrate that parameters such as clipping loss and the relative lengths of the recycling cavities, chosen primarily to optimise the detector sensitivity, will also shape the requirements of ET’s PI mitigation strategy. As such, I have proposed an action plan for integrating PI considerations into the ET design.

The second part of my thesis discusses the potential implementation of an ‘optomechanical filter’ to broaden the bandwidth of future gravitational wave detectors. I showed that, within the relevant parameter regime, a numerical model of the system shows good agreement with a lower-order analytical model of the system, and verified the optomechanical filter concept. This is, to my knowledge, the first numerical model of such a setup. Through this work I have also compared the sign, phase and other notational conventions typically used within the ‘theoretical’ and ‘experimental’ branches of the instrumental gravitational wave community, which can sometimes act as a barrier to communication between the two groups. There is now active research in modelling

interferometer configurations based on the optomechanical filter both analytically and numerically. Comparison of these results is a vital step prior to experimental validation of the designs.

Finally, I have presented a research program towards understanding whether the optomechanical filter concept can be practically implemented. Two particular challenges have so far been identified for the filter: its thermal noise requirements, and its controllability. In both cases, it is not yet known whether a viable solution will be found. These questions can be tackled from both a theoretical and experimental perspective in tandem, and so I have proposed designs for a tabletop experiment which can be used to develop and test control schemes in a coupled cavity system that is intended to demonstrate optomechanical bandwidth broadening.

We have started to set up a tabletop facility which can be used for this experiment, and I have reported our current status, which has focused on procurement and characterisation of an AlGaAs micro-oscillator chip. In the first instance, this facility will be used to experimentally demonstrate optomechanical parametric amplification and digital control of a radiation-pressure-dominated system in different stability regimes. It is then intended to be used to test control schemes as stated above. A final statement on the feasibility of the optomechanical filter, and of the predicted SNR improvements over different frequency ranges, will thus emerge from a combination of theoretical studies on the fundamental controllability of the filter, consideration of schemes to alter the thermal noise requirements, and practical experiences of working with the filter in the lab. This represents a high-risk, high-gain program of research that could enable new discoveries in gravitational wave astronomy.

Appendix A

Advanced LIGO Model

This appendix provides the key parameters used to model parametric instabilities as described in chapter 3. I provide a list of material properties used in COMSOL, and a copy of the FINESSE file, which includes all of the optical parameters used in the model. It has been adapted, as noted through the file, from the core file for Advanced LIGO by removing input and output optics as well as any other additional components outside of the core optics that are not considered in the study. This speeds up the simulation. The original file was developed by members of the LIGO FINESSE modelling team, including Daniel Brown, Charlotte Bond, Antonio Perreca, Paul Fulda and Andreas Freise. This file version was adapted by Daniel Töyrä, Andreas Freise and myself.

A description of how a file like this can be developed is given in Appendix B. The core FINESSE files for the global network of interferometric gravitational wave detectors are constantly evolving as the designs develop and to reflect the state of the real detectors. The current versions of these files are now included in the ‘ifo’ folder of PyKat [104], the python wrapper for FINESSE.

A.1 Test Mass Geometry and Material Properties

In table A.1 I summarise the properties of the LIGO test masses as used in a COMSOL model to identify their eigenmodes and associated surface motion.

Parameter	Value
Radius	17.0 cm
Thickness	20.0 cm
Height of flats	9.49 cm
Half-width to flats	16.325 cm
HR face radius of curvature	2245 m
Material	Fused Silica
– Density	2203 kg/m ³
– Poisson’s ratio	0.17
– Young’s Modulus	72.6 GPa

Table A.1: Parameters used to generate our COMSOL model of a LIGO test mass, from [2] and [59]. The geometry of each test mass is based on a cylinder; the sides are flattened to provide a surface that the suspension fibres can be bonded to.

A.2 Finesse Model

```
#-----
# an aLIGO file to introduce a PI into one arm
# based on
# aLIGO_IFO_AWC_tuning_DCOff_maxtem4.kat (details in the original file header)
#
# Anna Green and Daniel Toyra, 22nd September 2015
#-----

%%% FTblock Laser
#####
# Laser
l LO 125 0 n0
#s lmod1 1 n0 nMC1in
s lmod1 1 n0 nREFL
#removed all modulators
#####
%%% FTend Laser

# removed blocks: IMC, HAM2

%%% FTblock PRC
#####
```

A.2. FINESSE MODEL

```

# PRM
# AR surface
m2 PRMAR 35u 4.5u $phi_PRM nREFL nPRMARb

# Substrate
s sPRMsub1 0.0737 $nsilica nPRMARb nPRMHRa

# HR surface
m1 PRMHR 0.03 8.5u $phi_PRM nPRMHRa nPRMHRb
attr PRMHR Rc 11.009

# Distance between PRM and PR2
s lp1 $Lpr1 nPRMHRb nPR2a

# PR2
bs1 PR2 250u $Mloss 0 -0.79 nPR2a nPR2b nPOP dump11
attr PR2 Rc -4.545

# Distance from PR2 to PR3
s lp2 $Lpr2 nPR2b nPR3a

# PR3
bs1 PR3 0 $Mloss 0 0.615 nPR3a nPR3b dump12 dump13
attr PR3 Rc 36.027

# Distance from PR3
s lp3 $Lpr3 nPR3b nPRBS

#####
%% FTend PRC

%% FTblock BS
#####

# BS beamsplitter
##-----
## BS
##
##          to IMY      |
##          |           | ,'-
##          |           | + '
##          |           | , '
##          nYBS |       | , '
##          |       | +i1 +
##          |       | ,:._ i2 ,
##          ----->
##          from the PRC      nPRBS + \ ' - . + nXBS

```

A.2. FINESSE MODEL

```

##          , ' i3\      , ' ----->
##          +      \ +      to IMX
##          , '      i4.'
##          '._      ..
##          '._ , ' |nSRBS
##          -      |
##          |to the SRC
##          |
##          v
##-----
bs1 BS 0.5 $Mloss $phi_BS 45 nPRBS nYBS nBSi1 nBSi3
s BSsub1 0.0685 $nsilica nBSi1 nBSi2
s BSsub2 0.0684 $nsilica nBSi3 nBSi4
bs2 BSAR1 50u 0 0 -29.195 nBSi2 dump14 nXBS nPOX
bs2 BSAR2 50u 0 0 29.195 nBSi4 dump15 nSRBS dump16

#####
%%% FTend BS

%%% FTblock Yarm
#####

# Distance from beam splitter to Y arm input mirror
s ly1 5.0279 nYBS nITMY11

# Thermal lens correction
lens ITMYTL $TL_f nITMY11 nITMYTLtrans
s ITMYTL_null 0 nITMYTLtrans nITMYconstL_in
lens ITMYconstL inf nITMYconstL_in nITMYconstL_trans
s ITMYTL_null2 0 nITMYconstL_trans nITMY1

# Y arm input mirror
m2 ITMYAR 20u 0 $phi_ITMY nITMY1 nITMYs1
s ITMYsub 0.2 $nsilica nITMYs1 nITMYs2
m1 ITMYHR 0.014 $Mloss $phi_ITMY nITMYs2 nITMY2
#m1 ITMYHR 0 $Mloss $phi_ITMY nITMYs2 nITMY2
attr ITMYHR Rc -1934

# Y arm length
s LYarm 3994.5 nITMY2 nETMY1

# Y arm end mirror
m1 ETMYHR 5u $Mloss $phi_ETMY nETMY1 nETMYs1
s ETMYsub 0.2 $nsilica nETMYs1 nETMYs2
m2 ETMYAR 500u 0 $phi_ETMY nETMYs2 nPTY

```

A.2. FINESSE MODEL

```
attr ETMYHR Rc 2245
attr ETMYHR mass 40
attr ITMYHR mass 40

#####
%%% FTend Yarm

%%% FTblock Xarm
#####

# Distance from beam splitter to X arm input mirror
s lx1 5.0082 nXBS nITMX11

# Thermal lens correction
lens ITMXTL $TL_f nITMX11 nITMXTLtrans
s ITMXtl_null 0 nITMXTLtrans nITMXconstL_in
lens ITMXconstL inf nITMXconstL_in nITMXconstL_trans
s ITMXTL_null12 0 nITMXconstL_trans nITMX1

# X arm input mirror
m2 ITMXAR 20u 0 $phi_ITMX nITMX1 nITMXs1
s ITMXsub 0.2 $nsilica nITMXs1 nITMXs2
m1 ITMXHR 0.014 $Mloss $phi_ITMX nITMXs2 nITMX2
#m1 ITMXHR 0 $Mloss $phi_ITMX nITMXs2 nITMX2
attr ITMXHR Rc -1934

# X arm length
s LXarm 3994.5 nITMX2 nETMX1

# X arm end mirror
m1 ETMXHR 5u $Mloss $phi_ETMX nETMX1 nETMXs1
s ETMXsub 0.2 $nsilica nETMXs1 nETMXs2
m2 ETMXAR 500u 0 $phi_ETMX nETMXs2 nPTX
attr ETMXHR Rc 2245
attr ETMXHR mass 40
attr ITMXHR mass 40

#####
%%% FTend Xarm

%%% FTblock SRC
#####

# Distance to SR3
s ls3 $Lsr3 nSRBS nSR3b
```

A.2. FINESSE MODEL

```
# SR3
bs1 SR3 0 $Mloss 0 0.785 nSR3b nSR3a dump17 dump18
attr SR3 Rc 35.972841

# Distance from SR3 to SR2
s ls2 $Lsr2 nSR3a nSR2b

# SR2
bs1 SR2 0 $Mloss 0 -0.87 nSR2b nSR2a dump19 dump20
attr SR2 Rc -6.406

# Distance from SR2 to SRM
s ls1 $Lsr1 nSR2a nSRMHRa

# Signal recycling mirror SRM-08
m1 SRMHR $T_SRM $L_SRM $phi_SRM nSRMHRa nSRMHRb
s SRMsub 0.0749 $nsilica nSRMHRb nSRMARa
m2 SRMAR 50n 0 $phi_SRM nSRMARa nSRMARb
attr SRMHR Rc -5.6938

#####
%% FTend SR

##removed blocks: FI, OMC path, OMC, powers, errsig

%% FTblock Lengths
#####
# Calculate lengths of variables so Schnupp asymmetry can be changed easily.
# Cavity lengths
const Lprc 57.656
#const Lprc 57.645
const Lsrc 56.008
const Lschnupp 0.08

# Individual lengths
# PRC
const Lpr1 16.6107
const Lpr2 16.1647
const Lpr3 19.5381

# SRC
const Lsr1 15.7586
const Lsr2 15.4435
const Lsr3 19.3661
```

A.2. FINESSE MODEL

```
# Arms
const BStickness 0.06873

func Laver = $Lprc - $Lpr1 - $Lpr2 - $Lpr3
noplot Laver

# x length between BS and ITM
func Lmx = $Laver + 0.5*$Lschnupp - $BStickness * $nsilica - 0.2*$nsilica
noplot Lmx
put lx1 L $Lmx

# y length between BS and ITM
func Lmy = $Laver - 0.2*$nsilica - 0.5*$Lschnupp
noplot Lmy
put ly1 L $Lmy

# Lsr3
# func Lasrc = $Laver + $BStickness * $nsilica
# noplot Lasrc
# func Lsr3 = $Lasrc - $Lsr1 - $Lsr2 - $Lasrc
# noplot Lsr3
# put ls3 L $Lsr3

#####
%%% FTend Lengths

%%% FTblock HOMs
#####
maxtem 4
#cav cavIMC MC2 nMC2in MC2 nMC2refl
cav cavPRX PRMHR nPRMHRb ITMXHR nITMXs2
cav cavPRY PRMHR nPRMHRb ITMYHR nITMYs2
cav cavSRX SRMHR nSRMHRa ITMXHR nITMXs2
cav cavSRY SRMHR nSRMHRa ITMYHR nITMYs2
cav cavXARM ITMXHR nITMX2 ETMXHR nETMX1
cav cavYARM ITMYHR nITMY2 ETMYHR nETMY1
# removed OMC

#####
%%% FTend HOMs

%%% FTblock Reflectivities
#####
const Mloss 37.5u
```

A.2. FINESSE MODEL

```
const T_SRM 0.35
const L_SRM 8.7u
#####
%% FTend Reflectivities

%% FTblock Constants
#####
const nsilica 1.44963098985906
const nTGG 1.954
const nCalcite 1.65846
const f1 9099471
const nf1 -9099471
const f2 45497355
const nf2 -45497355
const f3 24000000
const nf3 -24000000

const fM 36.397884M
const fP 54.596826M
const TL_f 34.5k
#####
%% FTend Constants

%% FTblock Tunings
#####
const phi_SRM 90.0068420962415
const phi_PRM 0.000212467087295627
const phi_ITMX 4.36684131706263e-05
const phi_ITMY -4.36684131706263e-05
const phi_ETMX 0.00171608637012845
const phi_ETMY -0.00175327499066803
const phi_BS 0

#####
%% FTend Tunings

%% FTblock commands
%% FTend commands
```


Appendix B

Modelling Complex Interferometers in Finesse

In this appendix I provide an overview of how FINESSE can be used to model the core optics of a Dual-Recycled Michelson Interferometer with Fabry-Perot arm cavities (DRFPMi). The file is tuned and mode-matched as it is built, beginning with a single arm cavity and finishing with the quantum-limited sensitivity curve of the interferometer using DC readout. It is assumed that key parameter values, such as cavity lengths and mirror curvatures, have been pre-determined here, however one may explore each of these parameters as they are integrated into the model. In this case, it can be useful to first build a plane-wave model of the detector, then include Gaussian beams and associated properties later.

I first provide an overview of the procedure, then demonstrate the process using the example of an ET-120K design that has 6690 m radii of curvature on all test masses (see chapter 4 for details of this design).

B.1 Overview

In order to explore the behaviour of current detectors and test out proposed future detector designs, a base FINESSE file (with extension `.kat`) is written, containing all of the main optics in their optimal design configuration. For example, there are currently 3 main `kat` files for LIGO: one matching the Advanced LIGO design study, and one for each of the Livingston and Hanford sites.

While the first shows how the LIGO detectors were intended to work in the original design, the site files are updated to match parameters at the sites as closely as possible. The aLIGO file is therefore useful for understanding general behaviours in a more ideal interferometer case, while the site files are useful diagnostic tools.

We build the interferometer model in stages so that the behaviour of each element can be understood and checked along the way:

0. Before beginning a complex FINESSE model, it is useful to plan the file: collect together all needed constants, draw a sketch of the layout such as shown in figure B.1, and name the key components, lengths and nodes in a way that is memorable and adaptable.
1. **Model a single arm of the interferometer** including any optics in the path from the beamsplitter (BS) to the cavity. By default the laser wavelength is set to 1064 nm; this can be overridden using the `lambda` command. If the laser beam shape is not the fundamental Gaussian mode (or plane wave), this can be set using `tem` commands. Define the arm cavity using the `cav` command. By default the model is plane-wave; you can specify the order of higher order modes (HOMs) to include in the model using the `maxtem` command. At any time, the model can be swapped back to the plane waves case by setting `maxtem off`. If using Gaussian beams, check the cavity stability and other values of interest such as the spot sizes on the test masses by using `cp` commands or `trace 2`. Additional properties of the test masses, such as their mass, suspension transfer function, and radius (to include clipping effects) should be added as attributes using the `attr` command once the core optical model is built, so that their effects do not get confused with the base optical response.
2. **Introduce the Michelson and second arm cavity.** In an ideal case the two arms are identical copies of each other, so here we introduce the beamsplitter and link the two together. Now the behaviour of the Michelson with Fabry-Perot arms (FPMi) can be explored. The *CARM* and *DARM* degrees of freedom can be introduced as functions of the input mirror tunings, then used to plot the power at the exit port and in the arms as the common and differential displacements of the test masses are scanned. We can now also set the Michelson at the operating point by checking the beamsplitter tuning is such that for zero displacement, no light is detected at the exit port of the interferometer. If the beamsplitter is realistic (i.e. thick – modelled as a set of three beamsplitters connected by spaces with the optical path length and refractive index of the beamsplitter material, see figure B.2), the spaces from the BS to the arm cavity input mirrors may need to be adjusted to ensure that the the power

circulating in the two arms is matched. This also improves the mode matching in the FPMi.

3. **Add the power recycling mirror (PRM)**, and define the additional cavities this introduces with the X and Y -arms. This is the Power-Recycled FPMi (PRFPMi). The interferometer should now reach (or exceed, given that this is a ‘perfect’ model) the design arm circulating power. The effect can be observed by plotting powers as $CARM$ and $DARM$ are swept, as before, or by directly outputting the circulating power on resonance. If Gaussian beams are used, the power recycling cavity (PRC) can be mode matched to the FPMi by matching its curvature to the curvature of the field reflected from FPMi at the target position of the PRM. Further optimisation can be achieved by adjusting the curvature of the PRM or length of the space from the PRM to BS such that the arm power is maximised.

4. **Add the signal recycling mirror (SRM)** in the same way, to make the Dual Recycled FPMi (DRFPMi). In order to see the effect of signal recycling, it is more useful to look at the quantum-limited sensitivity curves. The masses of the input and end mirrors, and that they are suspended, must now be taken into account so that the radiation pressure curve can be calculated. A differential frequency signal, mimicking a gravitational wave, is then applied to the two arm lengths and the quantum-limited sensitivity calculated as this frequency is scanned. At this stage, we assume DC readout will be used, and therefore apply a small offset to the BS tuning (the ‘DC offset’) so that we have a detectable signal. Broadband sensitivity is achieved using RSE, when the signal recycling cavity (SRC) is tuned to be anti-resonant with the arms. Usually this occurs for an SRC tuning of 90° . In cases where higher order optical modes are included in the model, the BS tuning should be set such that the mode containing the gravitational wave signal is the dominant power source at the dark port; if other modes contain more power, these will dilute the signal and decrease the detector SNR.

The model now contains the ‘core optics’ of a DRFPMi gravitational wave detector. In reality, there will be many additional parts and features. These can be added to the model in a similar stepwise fashion when they become important to the behaviour of the detector. Potential additions include:

- Sensors, including photodiodes (`pd`) to measure the power at various locations, and demodulated (`pd[n]`) and partitioned photodiodes (via `pdtype`) like quadrant or bullseye sensors to measure error signals;
- Modulators (`mod`), usually added in the input path between the laser and PRM, so that error

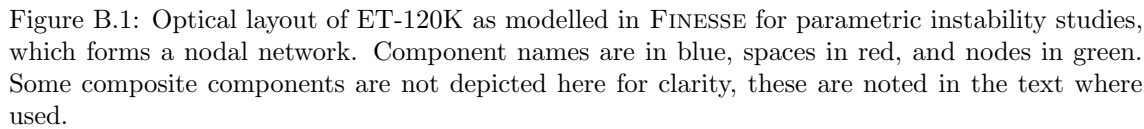
signals for the linear, angular and mode matching degrees of freedom can be modelled and control schemes designed;

- Schnupp Asymmetry – deliberately including a macroscopic length difference between the BS-ITMX and BS-ITMY path lengths so that modulated optical fields are partially transmitted to the output port of the detector;
- Input Mode Cleaner(s) (IMCs), included before the PRM to ensure that only the desired beam shape enters the detector¹;
- Output Mode Cleaner(s) (OMCs) to separate the optical mode containing the gravitational wave signal from HOMs generated in the detector;
- Isolators (`isol`), for injection of squeezed light and to prevent unwanted back-reflections;
- Squeezers and filter cavities to enable frequency-dependent squeezing;
- Homodyne detection – removing the DC offset and instead beating the field at the output port with a local oscillator at the same frequency as the main laser to extract the gravitational wave signal;
- Other new optical technologies for testing.

B.2 Example: ET-120K

This ET model is based on the parameters listed in table 4.1. The target for this file is to model how parametric instabilities will behave in ET-120K. From [1] we know that the properties of the arm cavities and the recycling cavities influence PI, therefore the model will include all the core optics, and must include high orders of optical modes and therefore be well mode matched (representing an ideal interferometer). The overall layout modelled is shown in figure B.1, including the key length and node names needed to create the file. In this example, I use a radius of curvature of 6690 m for all test masses, as described in section 4.4. We then approach the model by tuning and mode-matching the layout to the arm cavities at each stage.

¹In FINESSE the beam is ‘clean’ by default, however in reality the beam from the laser will include some HOM content that should be removed before entering the detector. Including the IMC in the detector model means that effects such as mode mismatching between the IMC and core detector can be explored



1. Single Arm

First I create the block of FINESSE code which models the X-arm, i.e. the optical path from the beamsplitter to ETMX:

```
%%% FTblock Xarm
#####

s lx1 300 nBSx nIMX1

lens lensIMX 303 nIMX1 nIMX2          %Input Mirror including focussing element
s simx1 0 nIMX2 nIMXi
m2 ITMXAR 0 $Mloss 0 nIMXi nIMXs1
s ITMXsub 0.55 $nsilicon nIMXs1 nIMXs2
m1 ITMX 0.003 $Mloss 0 nIMXs2 nIMXo
attr ITMX Rc -6690
attr ITMX mass $Mtm

s LX 10k nIMXo nEMXi

m1 ETMX 5u $Mloss 0 nEMXi nEMXs1      %End Mirror
s ETMXsub 0.55 $nsilicon nEMXs1 nEMXs2
m2 ETMXAR 0 $Mloss 0 nEMXs2 nEMXo
attr ETMX Rc 6690
attr ETMX mass $Mtm
#####
%%% FTend Xarm

cav cavXarm ITMX nIMXo ETMX nEMXi
```

Particular features to note are:

- the input and end mirrors are modelled as ‘composite’ optics, i.e. they are composed of two mirror components, forming the highly-reflective and anti-reflective (AR) surfaces of the mirror, and a **space** component with the refractive index of the test mass material (silicon).
- the input mirror in this design also directly includes a focusing element inside it. This is modelled by adding a lens (**lensIMX**) component followed by a null space of zero length (**simx1**).
- I have wrapped the code for the arm in a ‘block’ containing all core properties specific to the arm—this is useful longer-term for understanding the final file and manipulating it.
- The arm cavity is declared using the **cav** command, meaning that its eigenmodes will be computed. We will later collect all of these together in another block that can be removed

in plane-wave modes that do not require this feature.

The arm cavity code above did not include a laser, or specify general properties of the detector such as the laser wavelength, or number of higher order modes to include in calculations. This means that we cannot yet run this code. We will introduce the general interferometer properties and laser instead at the Michelson stage.

2. Fabry-Perot Michelson

The code block for the beamsplitter is:

```
%%% FTblock BS
#####
bs1 BS 0.5 $Mloss 0 45.0 nBSi nBSy nBSi1 nBSi3
s BSsub1 0.0687 $nsilica nBSi1 nBSi2
s BSsub2 0.0687 $nsilica nBSi3 nBSi4
bs2 BSAR1 0 $Mloss 0.0 -29.195 nBSi2 dump nBSx nPOX
bs2 BSAR2 0 $Mloss 0.0 29.195 nBSi4 dump nBSo dump
#####
%%% FTend BS
```

The beamsplitter is again a composite component, this time made of *three* angled mirror surfaces, using **bs** components, and *two* spaces with the material (fused silica) refractive index, due to the two paths light follows afterwards. This is shown in figure B.2.




Figure B.2: The beamsplitter is modelled as a composite component made from three angled reflective surfaces and two spaces. Component names are in blue, spaces in red, and nodes in green.

I now assemble these three blocks into a Fabry-Perot Michelson and introduce code to describe

general detector properties, add a laser input and define key constants:

```
1 LO 250 0 nBSi      % temporary laser component for testing
lambda 2e-06         % override the default laser wavelength and set it to 2 microns
maxtem 6             % include higher order optical modes up to 6th order
phase 2              % set the phase reference for HG00 to zero
const nsilicon 3.43  % list some constants, accessed with '$' tags
const nsilica 1.44963098985906
const Mloss 10u
const Mtm 203.7
```

We first check the behaviour of the common and differential degrees of freedom in the plane waves case by setting `maxtem off` and plotting the power circulating in the X- and Y-arms, and at the transmission port `nBSo` as the tuning of the arm cavities is changed in phase (CARM) and antiphase (DARM), as shown in figure B.3. Here, the beamsplitter is chosen so that the detector sits at its operating point, i.e. the power at `nBSo` is null for zero DARM detuning. We will later adjust this to enable DC readout.




Figure B.3: Power circulating in the X- and Y-arms, and transmitted from the Michelson, as the arm cavity tunings are changed in common (CARM) and differential (DARM) modes. When the beamsplitter is correctly tuned, signals that are common to both arms are reflected, while differential signals are transmitted from the Michelson. The power at the output node of the beamsplitter is therefore negligible for CARM detunings but produces a signal for DARM. In both cases, the X- and Y-arm cavities are moved away from resonance, therefore the circulating power drops.

Introducing the composite beamsplitter results in some small path-length differences between the otherwise-identical X- and Y-arms. This means that there is a small amount of mode mis-match between the arms, that can be seen by running the above test with higher order modes enabled. To correct this, I choose to match the Y-arm mode to the X-arm by slightly adjusting length `ly1` until

the power in the Y-arm is maximised This results in `ly1` becoming 4 cm longer than `lx1`, reducing mismatches to below 1 ppm (this can be output using the `mismatches` command).

3. Power-recycled Michelson

To introduce the power-recycling mirror (PRM), we first propagate the reflected beam from the Michelson to the position intended for the PRM and calculate the beam curvature at that location. This sets the radius of curvature of the PRM. I choose to set this to the closest integer value and then adjust the cavity length slightly to achieve good mode matching. The code for the PRC begins as

```
%%% FTblock PRC
m2 PRMAR 0.0 $Mloss 0.0 nPRMi nPRMs1
s sPRMsub 0.0737 $nsilica nPRMs1 nPRMs2
m1 PRM 0.033 $Mloss 0 nPRMs2 nPRMo
attr PRM Rc 51

s lpr1 10 nPRMo nBSi
%%% FTend PRC

cav cavPRX PRM nPRMo ITMX nIMXs2
cav cavPRY PRM nPRMo ITMY nIMYs2
```

Where once again, the PRM is a composite mirror, and I have introduced two new commands describing the cavities formed with the X-and Y- input mirrors.

After tuning length `lpr1` to maximise the power in the X-and Y-arms, its length is slightly reduced to 9.9115 m and the arm circulating power reaches the target value of 3 MW.

4. Signal-Recycled Michelson & Sensitivity Curves

Introduction of the SRC follows the same procedure as the PRC. In this case the SRC and PRC designs are identical, except that we set the signal recycling cavity tuning to 90°, i.e. RSE. The SRC code is thus

```
%%% FTblock SRC
s lsr1 10 nBSi nSRMi

m1 SRM 0.027 $Mloss 90 nSRMi nSRMs1 %RSE: SRM tuning 90 degrees
attr SRM Rc -51
s sSRMsub 0.0737 $nsilica nSRMs1 nSRMs2
m2 SRMAR 0.0 $Mloss 0.0 nSRMs2 nSRMo
%%% FTend SRC
```

```
cav cavSRX SRM nSRMi ITMX nIMXs2
cav cavSRY SRM nSRMi ITMY nIMYs2
```

and I again adjust the length to 9.872227 m to mode-match the SRC to the arms by maximising the arm power.

The quantum-limited sensitivity of this detector design is calculated by differentially modulating lengths L_X and L_Y . In the case of DC readout, a `qnoisedS` detector is used, and a small DC offset applied to the beamsplitter tunings so that some light leaks through to the output port (`nSRMo`). This tuning should be small, but sufficiently large that the fundamental mode (HG_{00}), which contains the gravitational wave signal, is the dominant optical field at the output. If other modes dominate then the signal-to-noise ratio is decreased, since they do not contain useful information. This is tested by placing amplitude detectors, sensitive to particular higher order modes, at the output power and measuring the power in each mode as the DARM tuning is adjusted. Figure B.4 shows the mode content up to 6th order at the port with a DC offset of 0.02° : the fundamental mode dominates by several orders of magnitude.




Figure B.4: Power in HG modes at the output node of the dual-recycled interferometer, as the arm cavity tunings are differentially changed. Each line is the modulus-squared output of an `ad`-type detector, which reads out the field amplitude of the HG_{nm} mode indicated. In this case, the DARM signal in the HG_{00} mode is ~ 7 orders of magnitude higher than the next largest modes (HG_{20} and HG_{02}), indicating that the file is very well mode matched.

Finally, figure B.5 shows the quantum-noise limited sensitivity curve for this detector design in three cases: a plane waves model, showing the ideal sensitivity curve given the other parameters of the detector; and the sensitivity curve before and after mode-matching. In this case, the detector is considered very well mode matched, so higher order mode leakage is very low, and the peak sensitivity more than doubles as a consequence.




Figure B.5: Quantum-noise-limited sensitivity curve of the ET-120K detector design using 6690 m test mass radii of curvature. The final mode-matched file (red) has the same sensitivity as expected from a plane-waves model of the optical system (black dashed). If instead the lengths `ly1`, `lpr1` and `lsr1` (or equivalent optic curvatures) are not adjusted optimally the detector signal-to-noise ratio decreases since higher order modes, which do not contain the gravitational wave signal, leak to the output port.

Appendix C

Input-Output Relation for a Radiation Pressure Limited Cavity

In this appendix the input-output relation of an optical cavity with a suspended end mirror and significant radiation pressure is derived, as used in chapter 5. Typically this relation has been derived in the quadrature picture [94]; here I present an alternative method which describes the behaviour in the sideband picture. Figure C.1 shows the optical configuration used throughout.




Figure C.1: Optical layout used to derive the input-output relation of a radiation-pressure-dominated cavity.

C.1 Linearised Time Domain Equations of Motion

In systems with dynamic behaviour, the steady-state complex field amplitudes used to describe the optical fields such as in section 2.3 become insufficient to describe the system. Instead, we describe the field amplitudes more generally, taking into account their time dependence. In the case that the end mirror is movable, the phase accumulated on propagation becomes dependent on the end

mirror position $x_{\text{RP}}(t)$, so the field incident on the suspended mirror is described by

$$\begin{aligned} a_2(t) &= a_1(t - \tau) e^{-ik(L+x_{\text{RP}}(t))} \\ &= it_1 a_{\text{in}}(t - \tau) e^{-ik[L+x_{\text{RP}}(t)]} + r_1 r_2 a_2(t - 2\tau) e^{-ik[2L+x_{\text{RP}}(t)+x_{\text{RP}}(t-2\tau)]}, \end{aligned} \quad (\text{C.1})$$

where $k = \omega_p/c$ is the wavenumber of a pump input field with angular frequency ω_p , L is the static cavity length, $\tau = L/c$ is the one-way propagation time in the cavity, and $r_{1,2}$ and $t_{1,2}$ are the amplitude reflectivities and transmissivities of the two mirrors.

We then assume that all of the functions are slow enough that their values can be represented as a first-order Taylor series expansion:

$$a(t + dt) = \sum_{n=0}^{\infty} \frac{a^{(n)}(t)}{n!} (dt)^n \approx a(t) + dt \dot{a}(t). \quad (\text{C.2})$$

This is the *single mode approximation*, and requires that $|\dot{a}(t) \cdot 2\tau| \ll |a(t)|$, or $|f_{\text{max}}| \ll f_{\text{FSR}}$, where f_{max} is the highest characteristic frequency of $a(t)$ and $f_{\text{FSR}} = \frac{c}{2L}$ is the free spectral range of the cavity.

The pump laser frequency ω_p is defined such that $\omega_p = \omega_0 + \Delta$, where ω_0 is the resonant frequency of the cavity ($\frac{\omega_0}{c}L = N\pi$, where N is an integer), and Δ is the detuning. If $2\Delta\tau \ll 1$, i.e. $\Delta \ll f_{\text{FSR}}$, then

$$e^{-ikL} \approx 1 - i\Delta\tau; \quad e^{-2i\frac{\Delta}{c}L} \approx 1 - 2i\Delta\tau. \quad (\text{C.3})$$

Assuming also that $t_{1,2} \ll 1$, we also write

$$r_1 r_2 \approx 1 - \frac{t_1^2 + t_2^2}{2} = 1 - \frac{\kappa}{2}. \quad (\text{C.4})$$

Here κ represents total power loss of the optical beam upon a single round trip inside the cavity.

Using approximations C.2, C.3 and C.4 in C.1 and keeping only zero and first order terms we can write a differential equation describing the dynamics of the field amplitude inside the cavity:

$$\dot{a}_2(t) + \left(\frac{\kappa f_{\text{FSR}}}{2} + i[\Delta + G_0 x_{\text{RP}}] \right) a_2(t) = -\frac{it_1}{2} \dot{a}_{\text{in}}(t) + \frac{t_1}{2} [G_0 x_{\text{RP}} + \Delta + 2if_{\text{FSR}}] a_{\text{in}}(t), \quad (\text{C.5})$$

where $G_0 = \frac{\omega_0}{L}$.

The dynamics of the field amplitude reflected from the cavity are thus described by

$$a_{\text{out}}(t) \approx it_1 r_2 \left[\left(1 - i\Delta\tau - i\frac{\omega_0}{c} x_{\text{RP}}(t) \right) a_2(t) - \tau \dot{a}_2(t) \right] + r_1 a_{\text{in}}(t). \quad (\text{C.6})$$

We can similarly write down the equation of motion for the suspended mirror, which is a simple harmonic oscillator driven by radiation pressure:

$$\ddot{x}_{\text{RP}}(t) + \gamma_m \dot{x}_{\text{RP}}(t) + \omega_m^2 x_{\text{RP}}(t) = \frac{2r_2^2 |a_2|^2}{mc} + \frac{F_{\text{ext}}}{m}, \quad (\text{C.7})$$

where ω_m and γ_m are the mechanical frequency and damping rate, $|a_2|^2$ is the power incident on the mirror, m is the mirror's mass, and F_{ext} describes any external force acting on the mirror.

Equations C.5 and C.7 form a system of nonlinear coupled equations. If the time dependence of the variables in these equations can be treated as perturbations, the system can be linearised to obtain an approximated system of equations describing the evolution of the system around the steady state. We therefore write

$$a_2(t) = A_2 + \alpha_2(t), \quad |\alpha_2(t)| \ll |A_2|; \quad (\text{C.8a})$$

$$x_{\text{RP}}(t) = X_{\text{RP}} + \xi(t), \quad |\xi(t)| \ll |X_{\text{RP}}|; \quad (\text{C.8b})$$

$$a_{\text{in}}(t) = A_{\text{in}} + \alpha_{\text{in}}(t), \quad |\alpha_{\text{in}}(t)| \ll |A_{\text{in}}|, \quad (\text{C.8c})$$

and consider the behaviour of the small perturbations to the steady state variables.

Linearisation is performed by substituting C.8 into C.5 and C.7, keeping only terms linear in ξ , α_2 and α_{in} :

$$\dot{\alpha}_2(t) + \left[\frac{\kappa f_{\text{FSR}}}{2} + i\Delta' \right] \alpha_2(t) = -iG_0 A_2 \xi(t) + it_1 f_{\text{FSR}} \alpha_{\text{in}}(t), \quad (\text{C.9})$$

$$\ddot{\xi}(t) + \gamma_m \dot{\xi}(t) + \omega_m^2 \xi(t) = \frac{2r_2^2}{mc} [A_2 \alpha_2^*(t) + A_2^* \alpha_2(t)] + \frac{F_{\text{ext}}}{m}, \quad (\text{C.10})$$

where we introduced the modified detuning

$$\Delta' = \Delta + G_0 X_{\text{RP}}, \quad (\text{C.11})$$

which is the frequency detuning of the optical carrier field A_2 from the resonant frequency of the cavity ω_0 in the presence of this carrier.

C.2 Frequency Domain Input-Output Relation

We take the Fourier transform of C.9 and C.10, remembering that $\widetilde{\alpha^*}(\omega) = \tilde{\alpha}^*(-\omega)$, to describe the equations of motion in the frequency domain:

$$\tilde{\alpha}_2(\omega) = \frac{t_1 f_{\text{FSR}} \tilde{\alpha}_{\text{in}}(\omega) - G_0 A_2 \tilde{\xi}(\omega)}{\omega + \Delta' - i \frac{\kappa f_{\text{FSR}}}{2}}, \quad (\text{C.12})$$

$$\tilde{\xi}(\omega) = \frac{\frac{2r_2^2}{c} [A_2 \tilde{\alpha}_2^*(-\omega) + A_2^* \tilde{\alpha}_2(\omega)] + \tilde{F}_{\text{ext}}}{m [-\omega^2 + i\omega\gamma_m + \omega_m^2]}. \quad (\text{C.13})$$

C.13 can be used to write down the mechanical susceptibility of the mirror:

$$\tilde{\chi}_m(\omega) = \frac{\tilde{\xi}(\omega)}{\tilde{F}_{\text{ext}}(\omega)} = \frac{1}{\tilde{\chi}_{m0}^{-1}(\omega) + \tilde{K}_{\text{opt}}(\omega)} \quad (\text{C.14})$$

where

$$\tilde{\chi}_{m0}(\omega) = \frac{1}{m [-\omega^2 + i\omega\gamma_m + \omega_m^2]} \quad (\text{C.15})$$

is the mechanical susceptibility of the oscillator in the absence of optical feedback, while

$$\tilde{K}_{\text{opt}}(\omega) = \frac{4r_2^2 G_0 |A_2|^2}{c} \frac{\Delta'}{\Delta'^2 + \left(i\omega + \frac{\kappa f_{\text{FSR}}}{2}\right)^2} \quad (\text{C.16})$$

is the optical rigidity, or complex optical spring, and accounts for the effect of radiation pressure on dynamics of the mechanical oscillator.

In the case of no external force F_{ext} , we can combine equations C.12, C.13 and C.14 to write the field amplitude incident on the mirror as

$$\tilde{\alpha}_2(\omega) = \frac{t_1 f_{\text{FSR}} \tilde{\alpha}_{\text{in}}(\omega)}{\omega + \Delta' - i \frac{\kappa f_{\text{FSR}}}{2}} - \frac{G_0 A_2 \tilde{F}_{\text{in}}(\omega) \tilde{\chi}_m(\omega)}{\left(\omega + \Delta' - i \frac{\kappa f_{\text{FSR}}}{2}\right)}, \quad (\text{C.17})$$

where

$$\tilde{F}_{\text{in}}(\omega) = \frac{2r_2^2 t_1 f_{\text{FSR}}}{c} \left[\frac{A_2^* \tilde{\alpha}_{\text{in}}(\omega)}{\omega + \Delta' - i \frac{\kappa f_{\text{FSR}}}{2}} + \frac{A_2 \tilde{\alpha}_{\text{in}}^*(-\omega)}{-\omega + \Delta' + i \frac{\kappa f_{\text{FSR}}}{2}} \right] \quad (\text{C.18})$$

describes the radiation pressure force acting on the mirror in terms of the input field.

Assuming that the field reflected from the cavity can also be treated perturbatively, and linearising

C.6, we can describe the reflected field at sideband frequency ω as:

$$\tilde{\alpha}_{\text{out}}(\omega) = t_1 r_2 \left[i + \frac{\omega + \Delta}{2f_{\text{FSR}}} \right] \tilde{\alpha}_2(\omega) + t_1 r_2 A_2 \frac{G_0 \tilde{\xi}(\omega)}{2f_{\text{FSR}}} + r_1 \tilde{\alpha}_{\text{in}}(\omega). \quad (\text{C.19})$$

Substituting C.17 and C.18 into C.19, we can directly express $\tilde{\alpha}_{\text{out}}$ as a linear function of $\tilde{\alpha}_{\text{in}}$.

C.3 Simplified form

We define $t_2 = 0$ and introduce several new variables to make the derivation more compact:

- $\gamma = \frac{ct_1^2}{4L}$ (the cavity bandwidth)
- $g = A_2 \sqrt{\frac{2\omega_0}{\hbar L c}}$, so that $\hbar|g|^2 = 2|A_2|^2 G_0/c$
- $\tilde{U} = G_0 A_2 \tilde{\xi}(\omega) = G_0 A_2 \tilde{\chi}_m(\omega) \tilde{F}_{\text{in}}(\omega)$

Substituting C.17 and C.18 into C.19 with these new variables gives:

$$\tilde{\alpha}_{\text{out}}(\omega) = \frac{t_1}{2f_{\text{FSR}}} \frac{2i f_{\text{FSR}} + \omega + \Delta}{\omega + \Delta - i\gamma} \left(t_1 f_{\text{FSR}} \tilde{\alpha}_{\text{in}} - \tilde{U} \right) + \frac{t_1 \tilde{U}}{2f_{\text{FSR}}} + \tilde{\alpha}_{\text{in}} \quad (\text{C.20})$$

$$\approx \left[1 + \frac{2i\gamma}{\omega + \Delta - i\gamma} \right] \tilde{\alpha}_{\text{in}} + \frac{t_1}{2f_{\text{FSR}}} \left[1 - \frac{2i f_{\text{FSR}} + \omega + \Delta}{\omega + \Delta - i\gamma} \right] \tilde{U} \quad (\text{C.21})$$

$$\approx \left[1 + \frac{2i\gamma}{\omega + \Delta - i\gamma} \right] \tilde{\alpha}_{\text{in}} - \frac{it_1}{\omega + \Delta - i\gamma} \tilde{U} \quad (\text{C.22})$$

$$= \left[1 + \frac{2i\gamma(\omega + \Delta - i\gamma - \hbar|g|^2 \tilde{\chi}_m(\omega))}{(\omega + \Delta - i\gamma)^2} \right] \tilde{\alpha}_{\text{in}}(\omega) - \frac{2i\gamma \hbar g^2 \tilde{\chi}_m(\omega)}{\Delta^2 - (\omega - i\gamma)^2} \tilde{\alpha}_{\text{in}}^*(-\omega). \quad (\text{C.23})$$

Then using C.16 and C.14 we can write

$$\tilde{\chi}_m(\omega) = \frac{\tilde{\chi}_{m0}(\omega)}{1 + \tilde{\chi}_{m0}(\omega) \frac{2\hbar|g|^2 \Delta}{\Delta^2 - (\omega - i\gamma)^2}} \quad (\text{C.24})$$

$$\therefore \frac{2i\gamma \hbar g^2 \tilde{\chi}_m(\omega)}{\Delta^2 - (\omega - i\gamma)^2} = \frac{2i\gamma \hbar g^2 \tilde{\chi}_{m0}(\omega)}{\Delta^2 - (\omega - i\gamma)^2 + 2\hbar|g|^2 \Delta \tilde{\chi}_{m0}(\omega)}. \quad (\text{C.25})$$

Substituting C.25 into C.23 and rearranging, we find the simplified input-output relation for a sideband reflecting from a radiation pressure limited cavity:

$$\begin{aligned} \tilde{\alpha}_{\text{out}}(\omega) &= \left[1 + \frac{2i\gamma [i\gamma + \Delta - \omega + \hbar|g|^2 \tilde{\chi}_{m0}(\omega)]}{(i\omega + \gamma)^2 + \Delta [\Delta + 2\hbar|g|^2 \tilde{\chi}_{m0}(\omega)]} \right] \tilde{\alpha}_{\text{in}}(\omega) \\ &\quad - \frac{2i\gamma \hbar g^2 \tilde{\chi}_{m0}(\omega)}{(i\omega + \gamma)^2 + \Delta [\Delta + 2\hbar|g|^2 \tilde{\chi}_{m0}(\omega)]} \tilde{\alpha}_{\text{in}}^*(-\omega). \end{aligned} \quad (\text{C.26})$$

Appendix D

Sign conventions for ‘Theorists’ and ‘Experimentalists’

Equation 5.3 uses the same sign conventions as [14] and FINESSE, amongst others. These conventions, though commonly used by many, are not fundamental, so long as they are self-consistent. Indeed many others, for example [94], use alternate self-consistent conventions.

In comparing the results produced by FINESSE to an analytical derivation of equation 5.3 which used the latter set of conventions (here nicknamed the ‘Theorists’ conventions’ since they are often used in the theoretical quantum optics community), we have identified three key convention differences:

- Definition of the electric field
- Definition of phase on reflection and transmission through optics
- Definition of the Fourier Transform

These definitions are summarised in table D.1.

Convention:	‘Theorist’	‘Experimentalist’
Electric Field	$E_1(t) \cos(\omega t - kz)$ $+ E_2(t) \sin(\omega t - kz)$	$E_1(t) \cos(\omega t - kz)$ $- E_2(t) \sin(\omega t - kz)$
Reflection/Transmission	direction-dependent	$E_{\text{refl}} = r E_{\text{in}}; E_{\text{trans}} = it E_{\text{in}}$
Fourier Transform	$f(\omega) = \int_{-\infty}^{\infty} f(t) e^{i\omega t} dt$	$\tilde{f}(\omega) = \int_{-\infty}^{\infty} f(t) e^{-i\omega t} dt$

Table D.1: Sign conventions used by theorists and experimentalists. Typically derivations in the quadrature picture (such as [94]) and sideband picture (such as [14] and FINESSE) use the conventions listed, however the results can be used to compare derivations using any self-consistent combination of these conventions.

The net result is that equation 5.3 (written there in the experimentalist’s convention) changes as:

$$\begin{aligned}
E_{\text{out}}^{\text{Experimentalist}}(\omega) &= \left\{ 1 - \frac{2i\gamma [-i\gamma - \Delta + \omega - \hbar|g|^2\chi_{xx}(\omega)]}{(\gamma + i\omega)^2 - \Delta [-\Delta - 2\hbar|g|^2\chi_{xx}(\omega)]} \right\} E_{\text{in}}(\omega) \\
&\quad + \frac{2i\hbar g^2 \gamma \chi_{xx}(\omega)}{(\gamma + i\omega)^2 - \Delta [-\Delta - 2\hbar|g|^2\chi_{xx}(\omega)]} E_{\text{in}}^*(-\omega); \quad (\text{D.1})
\end{aligned}$$

$$\begin{aligned}
E_{\text{out}}^{\text{Theorist}}(\omega) &= \left\{ -1 + \frac{2i\gamma [-i\gamma + \Delta - \omega + \hbar|g|^2\chi_{xx}(\omega)]}{(\gamma - i\omega)^2 + \Delta [\Delta + 2\hbar|g|^2\chi_{xx}(\omega)]} \right\} E_{\text{in}}(\omega) \\
&\quad + \frac{2i\hbar g^2 \gamma \chi_{xx}(\omega)}{(\gamma - i\omega)^2 + \Delta [\Delta + 2\hbar|g|^2\chi_{xx}(\omega)]} E_{\text{in}}^*(-\omega), \quad (\text{D.2})
\end{aligned}$$

and we also find the definition of g is slightly altered by the change in reflectivity/transmissivity convention, such that $g^{\text{Experimentalist}} = ig^{\text{Theorist}}$. These sign changes mean that the phase of the reflected sideband has the opposite sign to the ‘Experimentalists’ version’. However, the conventions also change the sign of the phase accumulated in a stable, fixed cavity, for example meaning that equation 2.18 (assuming a perfectly reflective end mirror) changes as

$$E_{\text{out}}^{\text{Theorist}}(\omega) = -\frac{\omega + \Delta - i\gamma}{\omega + \Delta + i\gamma} E_{\text{in}}(\omega); \quad (\text{D.3})$$

$$E_{\text{out}}^{\text{Experimentalist}}(\omega) = \frac{-\omega - \Delta - i\gamma}{-\omega - \Delta + i\gamma} E_{\text{in}}(\omega) \quad (\text{D.4})$$

This means that the relative sign between the radiation-pressure-limited and fixed-cavity response (see figure 5.5), and therefore the physical effect of interest, is maintained.

Appendix E

Controlling MIMO systems: State-Space Approach

For stable single-input-single-output (SISO) dynamical systems, we can often describe the system, and design its control scheme, entirely in the frequency domain. This is described well in [15]. However, it can be more intuitive, particularly for multi-input-multi-output (MIMO) systems including sources of delay, to first consider the system in the time-domain. As described in [105] and summarised here, we can describe the dynamical system using a *state-space* approach, and then determine whether the system is fundamentally *observable* and *controllable*. If these dual properties are true, a control scheme can, in principal, be built for the system.

E.1 Describing Dynamical Systems

We can compare the two notation styles using a simple example: a driven, damped harmonic oscillator. A mass, m is at a position $q(t)$. It is attached to a spring with spring constant k and damping γ . The other end of the spring is driven, so that its position is described by $q_0(t)$. The equation of state of the system is therefore

$$\ddot{q}(t) + 2\zeta\dot{q}(t) + q(t) = q_0(t), \quad (\text{E.1})$$

where we have rescaled the equation so $\zeta = \frac{\gamma}{\sqrt{mk}}$ is a dimensionless damping term.

E.1.1 Frequency Domain

In the *frequency* domain, we use the Laplace transform to then directly write the transfer function, $G(s)$ from the input, q_0 to the output, q :

$$\begin{aligned}\mathcal{L}(\ddot{q}(t) + 2\zeta\dot{q}(t) + q(t)) &= \mathcal{L}(q_0(t)) \\ &= (s^2 + 2\zeta s + 1)q(s) = q_0(s)\end{aligned}\tag{E.2}$$

so

$$G(s) = \frac{q(s)}{q_0(s)} = \frac{1}{1 + s^2 + 2\zeta s}.\tag{E.3}$$

s is the complex number frequency parameter. We can thus find the frequency response by evaluating the transfer function (TF) at $s = i\omega$, where ω is the angular frequency:

$$G(i\omega) = \frac{1}{1 - \omega^2 + 2i\omega\zeta}.\tag{E.4}$$

Transfer functions can then be understood by producing a *Bode Plot*, which plots the amplitude $|G(i\omega)|$, and phase $\angle G(i\omega)$ of the TF as functions of frequency.

$G(s)$ tends to infinity when the denominator is zero. In this case, the denominator of $G(s)$ equals zero when $s = -\zeta \pm \sqrt{\zeta^2 - 1}$. For our simple harmonic oscillator, $0 < \zeta < 1$ means that s has complex values: the system exhibits oscillatory behaviour which is underdamped. if $\zeta > 1$, s takes real values and corresponds to an exponential decay (overdamping).

Zeros in the denominator are referred to as *poles*, while zeros in the numerator are referred to as *zeros*. We can plot the values of s which correspond to these in the complex plane; this is referred to a pole-zero (P-Z) plot.

A real pole with a negative value corresponds to an exponential decay, while complex poles, which typically come in conjugate pairs, correspond to oscillatory behaviour. If the real part of the pole is negative, the oscillation again decays exponentially. If instead the real part of the pole is positive (a '*right half plane pole*'), the oscillation grows with time and is therefore unstable. This would happen for our simple harmonic oscillator if the damping term became negative.

E.1.2 State-Space Approach

In the *time* domain, we define the *state variables* of the system and use these to reduce the higher order equation of state into a set of first order equations in those variables.

For the simple harmonic oscillator we choose these variables to be the position and velocity of the end mass, and define a vector \vec{x} which describes the internal states of the system:

$$\vec{x} = \begin{pmatrix} q \\ \dot{q} \end{pmatrix}. \quad (\text{E.5})$$

In this case we have a single input to the system, $\vec{u} = u(t) = q_0(t)$, and a single output, $\vec{y} = y(t) = q(t)$, i.e. we measure only the position of the mirror.

We can now rewrite the dynamics in terms of the new variables:

$$\dot{\vec{x}} = \frac{d}{dt} \begin{pmatrix} q \\ \dot{q} \end{pmatrix} = \mathbf{A}\vec{x} + \mathbf{B}\vec{u} \quad (\text{E.6})$$

$$= \begin{pmatrix} 0 & 1 \\ -1 & -2\zeta \end{pmatrix} \begin{pmatrix} q \\ \dot{q} \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u(t), \quad (\text{E.7})$$

where by definition $\frac{d}{dt}q = \dot{q}$, and \ddot{q} is given by the original equation of state.

The output of the system can also be described in terms of the new variables:

$$\vec{y} = \mathbf{C}\vec{x} + \mathbf{D}\vec{u} \quad (\text{E.8})$$

$$= \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} q \\ \dot{q} \end{pmatrix} + 0u \quad (\text{E.9})$$

If the elements of \mathbf{B} are such that the inputs \vec{u} cannot cross-couple to all of the states in \vec{x} , the system is not *controllable*. Likewise, if the elements of \mathbf{C} result in restricted cross-couplings then the available outputs \vec{y} will not contain information about the corresponding state in \vec{x} , so the system is not considered *observable*.

Controllability and observability are dual properties. We can confirm whether a system is controllable by defining a controllability matrix with elements of form $\mathbf{A}^i\mathbf{B}$ for $i \in (0, n-1)$ where n is the number of elements in \vec{x} , i.e. $\mathbf{U}_{\text{control}} = (\mathbf{B} \ \mathbf{A}\mathbf{B} \ \mathbf{A}^2\mathbf{B} \ \dots \ \mathbf{A}^{n-1}\mathbf{B})$. We form an observability

matrix in a similar manner. The system is controllable if this matrix has full rank. For single-input-single-output systems, this requirement reduces down to requiring that $\mathbf{U}_{\text{control}}$ is invertible. A system is then considered controllable if we can always choose some $\vec{u}(t)$ to take the system from some arbitrary starting state to an arbitrary final state in finite time.

In the example of the simple harmonic oscillator, $n = 2$ so the observability matrix is

$$\mathbf{U}_{\text{obs}} = \begin{pmatrix} \mathbf{C} & \mathbf{AC} \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} 1 & 0 \end{pmatrix} & \begin{pmatrix} -1 & 0 \end{pmatrix} \end{pmatrix}, \quad (\text{E.10})$$

while the controllability matrix is

$$\mathbf{U}_{\text{control}} = \begin{pmatrix} \mathbf{B} & \mathbf{AB} \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ -1 \end{pmatrix} \end{pmatrix}, \quad (\text{E.11})$$

both of which are invertible. The simple harmonic oscillator can therefore, as expected, be observed and controlled.

E.2 Controlling Dynamical Systems

Once it has been determined that a system is controllable, we can introduce some controller which reads in the output signal(s) $\vec{y}(t)$ and acts on the system such that it follows reference signal(s) $\vec{r}(t)$. We do this by introducing feedback to the system. The general strategy for controlling a system using negative feedback is:

1. construct an error signal, $e(t)$ which compares the measured output $y(t)$ to the target value $r(t)$:

$$e(t) = r(t) - y(t) \quad (\text{E.12})$$

2. pass this error signal to a controller K which aims to minimise $e(t)$ by driving the input, $u(t)$.

E.2.1 describing systems with feedback

For SISO systems, we can use the frequency domain directly and build an intuitive picture using block diagrams, like figure E.1. The original system has a transfer function $G(s)$ which relates

an input $u(s)$ to output $y(s)$. We now introduce a controller, $K(s)$ which drives $u(s)$ based on some control law that aims to minimise the error signal $e(s)$. In the time domain, such a series




Figure E.1: Block diagram for a system $G(s)$ and controller $K(s)$.

of compound systems results in a convolution of functions; taking the Laplace transform of a convolution, however, simply results in a multiplication: $\mathcal{L}[G(t) * K(t)] = G(s)K(s)$. From the diagram we then see that the output can be related to the reference signal:

$$\begin{aligned} y(s) &= G(s)K(s)e(s) = L(s)e(s) \\ \rightarrow y(s) &= \frac{L(s)}{1 + L(s)}r(s) = T(s)r(s) \end{aligned} \quad (\text{E.13})$$

where we have defined the *open loop gain*, $L(s)$ and *closed loop gain*, $T(s)$.

In state-space, we consider the internal dynamics of the system first. Once again we describe the dynamics and outputs of the system as in equations E.15 and E.8. We introduce a negative feedback mechanism by driving the inputs in a manner that relates to the states of the system:

$$\vec{u} = -\mathbf{K}\vec{x}, \quad (\text{E.14})$$

where \mathbf{K} is the feedback matrix. The dynamics of the system are therefore modified:

$$\dot{\vec{x}} = \mathbf{A}\vec{x} + \mathbf{B}\vec{u} \quad \rightarrow \quad \dot{\vec{x}} = (\mathbf{A} - \mathbf{BK})\vec{x}. \quad (\text{E.15})$$

We can then transform this description of the internal dynamics into a physical description of the controlled system by taking the Laplace Transform and finding the transfer function of the controlled system in the frequency domain.

E.2.2 Stability

There are several ways in which a system with feedback can become unstable. As described in section E.1.1, the original physical system (referred to commonly as the *plant*) can in itself be

unstable if its transfer function contains a right-half-plane pole. We can generalise this statement to the state-space description by noting that the equation of motion for the plant has typical form $\dot{\vec{x}} = \mathbf{A}\vec{x}$. The eigenvalues of \mathbf{A} thus relate to the poles and zeros of the system. If one of these eigenvalues has a positive real part, we again get exponential growth. If all of the eigenvalues have negative real parts, the system is described as ‘linearly stable’ at $\vec{x}(t) = 0$.

We can extend this to consider the stability of the *controlled* system. Recall that the closed-loop transfer function is given by

$$T(s) = \frac{K(s)G(s)}{1 + K(s)G(s)}. \quad (\text{E.16})$$

$T(s)$ tends to infinity, and the system is therefore unstable, when $K(s)G(s) = L(s) = -1$. We can thus understand whether a controlled system will be stable by looking at the Bode plot of the open loop gain. Instability occurs when $|L(s)| = 1$ and *simultaneously* $\angle L(s) = 180^\circ$. Typically any control loop will include some proportional gain, so it is very likely that at some frequency the gain will be 1; we refer to this as the *unity gain frequency* (UGF). For the system to be stable, we therefore require that the phase lag at the UGF has not reached 180° .

One common source of instability is time delays in the system. For example, if there is a time delay δt in the sensor reading $y(t)$, the error signal will actually be composed of the reference signal $r(t)$ and the delayed output signal, $v(t) = y(t - \delta t)$. The shift theorem for Laplace transforms means that we can describe the sensor dynamics using a transfer function $H(s) = e^{-s\delta t}$ so that $v(s) = H(s)y(s)$. The closed loop transfer function for a controlled system with a sensing lag is then

$$T(s) = \frac{K(s)G(s)}{1 + H(s)K(s)G(s)}, \quad (\text{E.17})$$

i.e. the loop stability now depends on whether $HKG = -1$ anywhere. Since $|H(s)| = 1$, the time delay in sensing will change the frequency at which the phase of the open loop gain reaches 180° .

A second source of instability occurs when a transfer function contains zeros in the right half plane, since these also affect the phase of the transfer function. This is referred to as a ‘non-minimum phase system’. This can make the system more difficult to control. Sometimes, it is possible to design a simpler control scheme by changing which features of the system are observed. For example, it is easier to balance a vertical ruler on your hand by looking at the top of the ruler rather than the bottom. The transfer function between your hand and both ends of the ruler contains right half plane poles, however the TF from your hand to the bottom of the ruler also contains a right half plane zero. While both systems are therefore unstable, only one is also a non-minimum

phase system.

The exponential growth of an unstable mode $y(t)$ means that non-linear behaviour of a system becomes important in determining the overall dynamics. For example, a system may be stable up to a certain magnitude of gain, and then become unstable. This is referred to as *bifurcation*. The behaviour of the instability is first driven by the exponentially growing linear term. However, if the equation of motion also includes a nonlinear term, this will either speed up or slow down the rate of growth. For example if a mode of a system is described by $\dot{x} = \lambda x - ax^3$, where λ is an eigenvalue of the linear system and a is a constant, the cubic term will determine the value at which the unstable system will saturate. If instead we have $\dot{x} = \lambda x + ax^3$, the cubic term reinforces the instability and so the saturation point will be determined by the next leading order term.

E.2.3 Designing simple controllers

A controller is effective if it results in the system following the desired behaviour and does not have instabilities: i.e. the eigenvalues of matrix \mathbf{A} all have negative real parts, or in other words all of the poles and zeros of the open loop transfer function are in the left hand plane.

For example, a underdamped harmonic oscillator with damping $\zeta = 0.5$ is described in state-space as

$$\begin{aligned}\dot{\vec{x}} &= \begin{pmatrix} 0 & 1 \\ -1 & -1 \end{pmatrix} \begin{pmatrix} q \\ \dot{q} \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u(t), \\ \vec{y} &= \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} q \\ \dot{q} \end{pmatrix}.\end{aligned}$$

If instead we want the system to be critically damped, we require a controller which can change the behaviour so that $\zeta = 1$ and the eigenvalues become $\lambda = (-1, -1)$.

Adding the controller modifies the dynamics such that

$$\dot{\vec{x}} = (\mathbf{A} - \mathbf{BK})\vec{x} = \mathbf{A}'\vec{x} \tag{E.18}$$

so we can find a controller that performs as required by equating the coefficients of the characteristic

equations for the controlled and target stability matrices:

$$\begin{aligned}(1 + \lambda)^2 &= \lambda^2 + (1 + k_2)\lambda + 1 + k_1 \\ \therefore \mathbf{K} &= \begin{pmatrix} k_1 & k_2 \end{pmatrix} = \begin{pmatrix} 0 & 1 \end{pmatrix}\end{aligned}\tag{E.19}$$

Many methods have been developed to design controllers for more complex systems. Common examples include Linear Quadratic (LQR and LQG), H_2 , H_∞ and μ -controller optimisations.

Bibliography

- [1] AC Green, DD Brown, M Dovale-Álvarez, et al. The influence of dual-recycling on parametric instabilities at Advanced LIGO. *Classical and Quantum Gravity*, 34(20):205004, 2017. [ii](#), [7](#), [29](#), [41](#), [44](#), [131](#), [145](#)
- [2] J Aasi, BP Abbott, R Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015. [viii](#), [1](#), [2](#), [42](#), [45](#), [52](#), [58](#), [135](#)
- [3] BP Abbott, R Abbott, TD Abbott, et al. Observation of gravitational waves from a binary black hole merger. *Physical review letters*, 116(6):061102, 2016. [1](#)
- [4] BP Abbott, R Abbott, TD Abbott, et al. GW151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Physical Review Letters*, 116(24):241103, 2016. [1](#)
- [5] BP Abbott, R Abbott, TD Abbott, et al. GW170104: Observation of a 50-solar-mass binary black hole coalescence at redshift 0.2. *Physical Review Letters*, 118(22):221101, 2017. [1](#)
- [6] BP Abbott, R Abbott, TD Abbott, F. Acernese, et al. GW170608: Observation of a 19 solar-mass binary black hole coalescence. *The Astrophysical Journal Letters*, 851(2):L35, 2017. [1](#)
- [7] F Acernese, M Agathos, K Agatsuma, et al. Advanced virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, 2015. [1](#)
- [8] BP Abbott, R Abbott, TD Abbott, et al. GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Physical Review Letters*, 119(14):141101, 2017. [1](#)
- [9] BP Abbott, R Abbott, TD Abbott, et al. GW170817: observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119(16):161101, 2017. [1](#), [63](#), [87](#)
- [10] M Evans, S Gras, P Fritschel, et al. Observation of parametric instability in advanced ligo. *Phys. Rev. Lett.*, 114:161102, Apr 2015. [4](#), [32](#), [42](#), [47](#), [64](#)
- [11] VB Braginsky, SE Strigin, and SP Vyatchanin. Parametric oscillatory instability in fabry-perot interferometer. *Physics Letters A*, 287:331–338, September 2001. [5](#), [27](#), [29](#), [30](#), [64](#)
- [12] ET science team. Einstein gravitational wave telescope conceptual design study. Technical Report ET-0106C-10, ET, 2010. [5](#), [62](#), [63](#), [65](#), [67](#), [70](#), [71](#), [80](#), [86](#)

- [13] H Miao, Y Ma, C Zhao, and Y Chen. Enhancing the bandwidth of gravitational-wave detectors with unstable optomechanical filters. *Phys. Rev. Lett.*, 115(21):211104, 2015. 6, 70, 94, 95, 104, 107, 108, 109, 110, 111, 112, 122
- [14] C Bond, D Brown, A Freise, and KA Strain. Interferometer techniques for gravitational-wave detection. *Living Reviews in Relativity*, 19(1):3, 2017. 9, 11, 13, 19, 24, 26, 158, 159
- [15] A Freise. *The Next Generation of Interferometry: Multi-Frequency Optical Modelling, Control Concepts and Implementation*. PhD thesis, Universität Hannover, 2003. (Appendix B). 13, 160
- [16] TR Corbitt. *Quantum Noise and Radiation Pressure Effects in High Power Optical Interferometers*. PhD thesis, Massachusetts Institute of Technology, 2008. 17, 29, 38, 39, 40
- [17] H Kogelnik and T Li. Laser beams and resonators. *Appl. Opt.*, 5(10):1550–1567, 1966. 24, 26, 75
- [18] C Bond. *How to stay in shape: Overcoming beam and mirror distortions in advanced gravitational wave interferometers*. PhD thesis, University of Birmingham, June 2014. 24
- [19] T Corbitt, D Ottaway, E Innerhofer, J Pelc, and N Mavalvala. Measurement of radiation-pressure-induced optomechanical dynamics in a suspended fabry-perot cavity. *Physical Review A*, 74(2):021802, 2006. 27, 39, 40
- [20] M Evans, L Barsotti, and P Fritschel. A general approach to optomechanical parametric instabilities. *Physics Letters A*, 374(4):665 – 671, 2010. 28, 29, 31, 44, 48, 60
- [21] C Blair. Three mode parametric instability and their control for advanced gravitational wave detectors. *Science China Physics, Mechanics and Astronomy*, 58(1):014201, 2015. 29
- [22] C Blair. *Parametric Instability in Gravitational Wave Detectors*. PhD thesis, University of Western Australia, 2017. 29, 42, 43, 62, 72
- [23] A Freise, G Heinzl, H Lueck, et al. Frequency-domain interferometer simulation with higher-order spatial modes. *Classical and Quantum Gravity*, 21(5):S1067–S1074, 2004. 31
- [24] DD Brown and A Freise. Finesse, May 2014. You can download the binaries and source code at <http://www.gwoptics.org/finesse>. 31
- [25] C Bond, P Fulda, L Carbone, K Kokeyama, and A Freise. Higher order laguerre-gauss mode degeneracy in realistic, high finesse cavities. *Phys. Rev. D*, 84:102002, Nov 2011. 31
- [26] L Carbone, C Bogan, P Fulda, A Freise, and B Willke. Generation of high-purity higher-order laguerre-gauss beams at high laser power. *Physical Review Letters*, 110:251101, 2013. 31
- [27] DD Brown, H Miao, C Collins, et al. Broadband sensitivity enhancement of detuned dual-recycled michelson interferometers with epr entanglement. *Phys. Rev. D*, 96:062003, Sep 2017. 31
- [28] D Töyrä, DD Brown, M Davis, et al. Multi-spatial-mode effects in squeezed-light-enhanced interferometric gravitational wave detectors. *Phys. Rev. D*, 96:022006, Jul 2017. 31
- [29] D Brown, R Smith, and A Freise. Fast simulation of gaussian-mode scattering for precision interferometry. *Journal of Optics*, 18(2):025604, 2016. 32

- [30] DD Brown. *Interaction of light and mirrors: Advanced techniques for modelling future gravitational wave detectors*. PhD thesis, University of Birmingham, 2015. 32, 34, 38, 39, 40, 41
- [31] A Freise and DD Brown. *FINESSE 2.0 manual*. 32, 34, 74
- [32] T Corbitt, Y Chen, E Innerhofer, et al. An all-optical trap for a gram-scale mirror. *Phys. Rev. Lett.*, 98(15):150802, Apr 2007. 37, 39, 124
- [33] DV Martynov, ED Hall, and BP Abbott et. al. The sensitivity of the advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93(112004), 2016. 42
- [34] C Blair, S Gras, R Abbott, et al. First demonstration of electrostatic damping of parametric instability at advanced ligo. *Phys. Rev. Lett.*, 118:151102, Apr 2017. 42, 43
- [35] C Zhao, L Ju, J Degallaix, S Gras, and DG Blair. Parametric Instabilities and Their Control in Advanced Interferometer Gravitational-Wave Detectors. *Physical Review Letters*, 94(12):121102–+, April 2005. 43
- [36] C Zhao, DG Blair, P Barrigo, et al. Gingen high optical power test facility. *Journal of Physics: Conference Series*, 32(1):368, 2006. 43
- [37] J Degallaix, C Zhao, L Ju, and D Blair. Thermal tuning of optical cavities for parametric instability control. *J. Opt. Soc. Am. B*, 24(6):1336–1343, June 2007. 43
- [38] C Zhao, L Ju, Y Fan, et al. Observation of three-mode parametric interactions in long optical cavities. *Physics Review A*, 78(2):023807, 2008. 43
- [39] S Susmithan, C Zhao, F Qi, L Ju, and D Blair. Thermal tuning the optical cavity for 3 mode interaction studies using a co 2 laser. *Journal of Physics: Conference Series*, 363(1):012018, 2012. 43
- [40] J Miller, M Evans, L Barsotti, et al. Damping parametric instabilities in future gravitational wave detectors by means of electrostatic actuators. *Physics Letters A*, 375(3):788 – 794, 2011. 43
- [41] M Hewitson, K Danzmann, H Grote, et al. Charge measurement and mitigation for the main test masses of the geo 600 gravitational wave observatory. *Classical and Quantum Gravity*, 24(24):6379, 2007. 43
- [42] S Gras, P Fritschel, L Barsotti, and M Evans. Resonant dampers for parametric instabilities in gravitational wave detectors. *Phys. Rev. D*, 2015. 43
- [43] P Fritschel, S Biscans, T Hardwick, C Blair, and S Gras. Test of acoustic mode dampers for parametric instability control. Technical Report G1800454, LIGO, 2018. 43
- [44] Y Fan, L Merrill, C Zhao, et al. Testing the suppression of opto-acoustic parametric interactions using optical feedback control. *Classical and Quantum Gravity*, 27(8), 2010. 44
- [45] DD Brown, AC Green, M Dovale-Álvarez, et al. Optical suppression of parametric instabilities with extraction cavities. (*in preparation*), 2017. (*in preparation*). 44

- [46] C Zhao, L Ju, Q Fang, et al. Parametric instability in long optical cavities and suppression by dynamic transverse mode frequency modulation. *Physical Review D*, 91(9):092001, 2015. 44
- [47] VB Braginsky, SE Strigin, and SP Vyatchanin. Analysis of parametric oscillatory instability in power recycled LIGO interferometer. *Physics Letters A*, 305:111–124, December 2002. 44
- [48] SE Strigin and SP Vyatchanin. Analysis of parametric oscillatory instability in signal recycled LIGO interferometer with different arms. *Phys. Letters A*, 2006. 44, 60
- [49] AG Gurkovsky, SE Strigin, and SP Vyatchanin. Analysis of parametric oscillatory instability in signal recycled LIGO interferometer. *Physics Letters A*, 362(2-3):91–99, 2007. 44
- [50] S Gras, C Zhao, DG Blair, and L Ju. Parametric instabilities in advanced gravitational wave detectors. *Classical and Quantum Gravity*, 27(20):205019, 2010. 44
- [51] W Kells. Parametric instability with general recycling. Technical report, LIGO, 2006. 44
- [52] SL Danilishin, SP Vyatchanin, DG Blair, J Li, and C Zhao. Time evolution of parametric instability in large-scale gravitational-wave interferometers. *Phys. Rev. D*, 90(12):122008, Dec 2014. 44
- [53] R Abbott, R Adhikari, S Ballmer, et al. Advanced LIGO Length Sensing and Control Final Design. Technical Report T1000298, LIGO, 2010. 45
- [54] MA Arain and G Mueller. Design of the advanced ligo recycling cavities. *Optics Express*, Vol. 16, Issue 14:10018–10032, 2008. 46
- [55] BJ Meers. The frequency response of interferometric gravitational wave detectors. *Physics Letters A*, 142(8 - 9):465 – 470, 1989. 52
- [56] H Wang, M Dovale-Álvarez, C Collins, et al. On the edge: the feasibility of near-unstable cavities for future gravitational wave detectors. *ArXiv e-prints*, (1711.05177), 2017. 57
- [57] D Toyra et. al. (tbc). Technical Report T1700286, LIGO, 2017. 57
- [58] C Bond, P Fulda, D Brown, et al. FINESSE input files for llo advanced LIGO. technical report. Technical Report T1300901, LIGO, 2014. 60
- [59] G Billingsley. LIGO core optics reference page. <https://galaxy.ligo.caltech.edu/optics/>, Accessed 2017.12.08. 60, 135
- [60] LIGO Scientific Collaboration. Instrument science white paper 2017-8. Technical Report T1700231, LIGO, 2017. 62, 69
- [61] BP Abbott, R Abbott, RX Adhikari, et al. Multi-messenger observations of a binary neutron star merger. *Astrophysical Journal Letters*, 848(2):L12, 2017. 63, 87
- [62] X Chen, C Zhao, S Danilishin, et al. Observation of three-mode parametric instability. *Phys. Rev. A*, 91(3):033832, Mar 2015. 64
- [63] SE Strigin. The effect of parametric oscillatory instability in a fabry-perot cavity of the einstein telescope. *Optics and Spectroscopy*, 112(3), 2012. © Pleiades Publishing, Ltd., 2012. Original Russian Text © S.E. Strigin, 2012, published in *Optika i Spektroskopiya*, 2012, Vol. 112, No. 3. 64, 91

- [64] SE Strigin. Parametric oscillatory instability in a fabry-perot cavity of the einstein telescope with different mirror's materials. *Phys. Letters A*, 2017. 64, 91
- [65] SE Strigin. Suppression of parametric oscillatory instability in third generation gravitational wave detectors. *Phys. Letters A*, 379(28-29), 2015. 64
- [66] J Zhang, C Zhao, L Ju, and D Blair. Study of parametric instability of gravitational wave detectors using silicon test masses. *ArXiv*, (1609.01505), 2016. 64, 73, 83, 91
- [67] RX Adhikari, N Smith, A Brooks, et al. LIGO voyager upgrade: Design concept. Technical Report LIGO-T1400226, LIGO Scientific Collaboration, 2018. (accessed 6 Feb 2018). 67, 70, 71, 72
- [68] W G Unruh. *Quantum Optics, Experimental Gravitation, and Measurement Theory*, chapter Quantum Noise in the Interferometer Detector, page 647. Plenum Press, New York, 1982. 69, 95
- [69] HJ Kimble, Y Levin, AB Matsko, KS Thorne, and SP Vyatchanin. Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics. *Physical Review D*, 65(2):022002, January 2002. 69, 95
- [70] J Harms, Y Chen, S Chelkowski, et al. Squeezed-input, optical-spring, signal-recycled gravitational-wave detectors. *Physical Review D*, 68(4):042001, August 2003. 69, 95
- [71] S Chelkowski, H Vahlbruch, B Hage, et al. Experimental characterization of frequency-dependent squeezed light. *Physical Review A*, 71(1):013806+, January 2005. 69, 95
- [72] J Miller, L Barsotti, S Vitale, et al. Prospects for doubling the range of advanced LIGO. *Phys. Rev. D*, 91(062005), 2015. 69
- [73] S Hild, H Grote, J Degallaix, et al. DC-readout of a signal-recycled gravitational wave detector. *Classical and Quantum Gravity*, 26(5):055012 (10pp), 2009. 69
- [74] P Fritschel, M Evans, and V Frolov. Balanced homodyne readout for quantum limited gravitational wave detectors. *Optics Express*, 22(4), 2014. 69
- [75] H Yu, D Martynov, S Vitale, et al. Prospects for detecting gravitational waves at 5 hz with ground-based detectors. *Phys. Rev. Lett.*, 120(14), 2018. 69
- [76] R Nawrodt, L Cunningham, R Flaminio, et al. Mirror thermal noise calculation for et. Technical Report ET-027-09, Einstein Telescope, 2009. 71
- [77] J Read, C Markakis, M Shibata, et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys. Rev. D*, 2009. 73
- [78] H Wang. *Beware of Warped Surfaces: Near-Unstable Cavities for Future Gravitational Wave Detectors*. PhD thesis, University of Birmingham, 2017. 79, 81, 124
- [79] J-Y Vinet. On special optical modes and thermal issues in advanced gravitational wave interferometric detectors. *Living Reviews in Relativity*, 12(5), 2009. 80
- [80] B Mours, E Tournefier, and J-Y Vinet. Thermal noise reduction in interferometric gravitational wave antennas: using high order TEM modes. *Classical and Quantum Gravity*, 23:5777–5784, October 2006. 81

- [81] P Fulda, K Kokeyama, S Chelkowski, and A Freise. Experimental demonstration of higher-order laguerre-gauss mode interferometry. *Phys. Rev. D*, 82(1):012002, Jul 2010. 81
- [82] C Bond, P Fulda, L Carbone, K Kokeyama, and A Freise. The effect of mirror surface distortions on higher order laguerre-gauss modes. *Journal of Physics: Conference Series*, 363:012005, 2012. 81
- [83] D Martynov, H Miao, H Yang, et al. Prospects of achieving high sensitivity around kilohertz with current interferometer topology. (*in preparation*), 2018. 87
- [84] H Miao. Private email from Haixing Miao, April 2018. Requested parameters for a kHz-optimised 10km detector; his reply “For 10km facility, the numbers for SRC, which I got from my optimisation, are $T_{itm} = 0.029$, $T_{srm} = 0.024$, $L_{src} = 91.1(m)$. The optimisation is to maximise the sensitivity from 1kHz to 4kHz with an arm cavity power of 3MW and 12dB squeezing.”. 87
- [85] J Mizuno, KA Strain, PG Nelson, et al. Resonant sideband extraction: a new configuration for interferometric gravitational wave detectors. *Physics Letters A*, 175(5):273 – 276, 1993. 95
- [86] A Wicht, K Danzmann, M Fleischhauer, et al. White-light cavities, atomic phase coherence, and gravitational wave detectors. *Optics Communications*, 134:431–439, 1997. 95
- [87] A Wicht, M Müller, R-H Rinkleff, A Rocco, and K Danzmann. Experimental demonstration of negative dispersion without absorption. *Optics Communications*, 179:107 – 115, 2000. 95
- [88] GS Pati, M Salit, K Salit, and MS Shahriar. Demonstration of a tunable-bandwidth white-light interferometer using anomalous dispersion in atomic vapor. *Phys. Rev. Lett.*, 99(13), 2007. 95
- [89] Y Ma, H Miao, C Zhao, and Y Chen. Quantum noise of a white-light cavity using a double-pumped gain medium. *Physical Review A*, 92(023807), 2015. 95, 113
- [90] S Wise, G Mueller, D Reitze, DB Tanner, and BF Whiting. Linewidth-broadened fabry-perot cavities within future gravitational wave detectors. *Classical and Quantum Gravity*, 21(5), 2004. 95
- [91] S Wise, V Quetschke, AJ Deshpande, et al. Phase Effects in the Diffraction of Light: Beyond the Grating Equation. *Physical Review Letters*, 95(1):013901, 2005. 95
- [92] HN Yum and J Scheuer. Demonstration of white light cavity effect using stimulated brillouin scattering in a fiber loop. *Journal of Lightwave Technology*, 31(23), 2013. 95
- [93] M Zhou, S Kim, G Bahl, and SM Shahriar. Realization of a negative dispersion medium at 1064nm in an optomechanical resonator for enhancing the sensitivity-bandwidth product in a gravitational-wave detector. *ArXiv e-prints*, (1610.07255), 2016. 95
- [94] H Miao. *Exploring Macroscopic Quantum Mechanics in Optomechanical Devices*. PhD thesis, The university of western Australia, 2010. 98, 153, 158, 159
- [95] J Bentley. Reducing the shot noise of cosmic explorer using an unstable optomechanical filter. Technical Report G1800440, LIGO, 2018. 110

- [96] DF McGuigan, CC Lam, RQ Gram, et al. Measurements of the mechanical Q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30(5):621–629, Mar 1978. 111
- [97] H Miao, H Yang, and D Martynov. Towards the Design of Gravitational-Wave Detectors for Probing Neutron-Star Physics. *ArXiv e-prints*, December 2017. 112
- [98] M Page, J Qin, JL Fontaine, C Zhao, and D Blair. Enhanced gravitational-wave detection of black hole quasi-normal modes using optically diluted optomechanical filters. *ArXiv e-prints*, (1711.04469), 2017. 115
- [99] J Cripe, N Aggarwal, R Lanza, et al. Observation of a room-temperature oscillator’s motion dominated by quantum fluctuations over a broad audio-frequency band. *ArXiv e-prints*, February 2018. 115
- [100] GD Cole, S Gröblacher, K Gugler, S Gigan, and M Aspelmeyer. Monocrystalline $\text{Al}_x\text{Ga}_{1-x}\text{As}$ heterostructures for high-reflectivity high-Q micromechanical resonators in the megahertz regime. *Applied Physics Letters*, 92(26):261108, June 2008. 115
- [101] GD Cole, W Zhang, MJ Martin, J Ye, and M Aspelmeyer. Tenfold reduction of brownian noise in high-reflectivity optical coatings. *Nature Photonics*, 7(8), 2013. 115
- [102] Thorlabs. Polarization-maintaining single mode optical fiber. www.thorlabs.com - accessed 14th May 2018. 121
- [103] ITU. Characteristics of a single-mode optical fibre and cable. www.itu.int/rec/T-REC-G.652-201611-I/en - accessed 14th May 2018. 121
- [104] DD Brown and A Freise. Pykat, July 2017. <http://www.gwoptics.org/pykat>. 134
- [105] J Bechhoefer. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.*, 77:783–836, Aug 2005. 160