

TOWARDS IMPLEMENTATION OF
QUANTUM NOISE REDUCTION SCHEMES
IN 3RD GENERATION GRAVITATIONAL WAVE
DETECTORS

by
PHILIP JONES

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



Institute for Gravitational Wave Astronomy
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
23rd September 2021

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This unpublished thesis/dissertation is copyright of the author and/or third parties.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

Since the initial detection of gravitational waves in 2015, a slow and steady stream of events has been recorded. The next step is to detect more and weaker sources, probing deeper into the universe. In order to do so, our interferometers must be pushed to even greater extremes, to suppress the noise that would otherwise obscure such faint signals.

Quantum noise is perhaps the most fundamental source of noise in a laser interferometer. Arising from the inherent uncertainty in the amplitude and phase of any light source, it limits the sensitivity of all current and proposed future interferometers, especially at higher signal frequencies. The second generation detectors currently in use, such as Advanced LIGO and Virgo, are operating around the standard quantum limit—the maximum sensitivity that can be reached by employing classical techniques. For future upgrades and new detectors, advanced quantum techniques must be employed to further increase sensitivity. Advanced LIGO has already demonstrated its ability to surpass the standard quantum limit over a small range of frequencies by making use of squeezed states of light, and future improvements will increase the magnitude and frequency range of this sensitivity enhancement.

In this thesis, I will detail my work on these quantum techniques, aimed particularly at the upcoming third generation of gravitational wave detectors. This covers a wide variety of topics: different arrangements of filter cavities for improved squeezing, an alternate detection scheme that reduces both quantum and more technical noises,

a mechanism that operates via complex optomechanical interactions, and more. In each of these subjects, my role has largely been to perform validation of new ideas and techniques, and to investigate the effects of varying detector parameters on the proposed scheme. To this end, I have made extensive use of a numerical frequency domain interferometer simulation tool known as FINESSE.

Widely used throughout the interferometer community, FINESSE allows the fast and accurate modelling of optical layouts. Certain advanced techniques, however, were beyond FINESSE's capabilities when I started my work. This thesis therefore also describes my contributions towards a full rewrite of this tool, called FINESSE 3. I performed a complete reimplementaion of the quantum noise calculations in FINESSE 3, and added features necessary to model new and exciting ideas. I will explain the mathematics that FINESSE 3 uses to simulate interferometers, before finally giving a broader overview of its design, usage, and the great potential it has to allow investigation of ever more complex techniques for the foreseeable future.

Acknowledgements

To my supervisor, Andreas, who should already know how fundamental he was in my decision to do a PhD, the fun I've had over this past four years, and without who's constant guidance I wouldn't have achieved half of the things I have here.

To my parents, who's love and support I have relied on fully to get me through any hardship I've faced, and who have been far more happy to hear me witter on about 'quantum noise this' and 'interferometer that' than they have any right to.

To my Brother, for his part in making me the person I am today, and to my grandparents, and everyone else in my family, who have always taken great pride and interest in my work.

To my friends and colleagues in ifolab, who made me look forward to coming to work each day, and to my coworkers here and elsewhere, who have always been happy to lend a helping hand or collaborate on new ideas.

To everyone else I've met in the grav-waves community, who've consistently been a thoroughly lovely bunch.

Thank you.

Statement of Originality

This thesis contains my own research work conducted over the course of my PhD between September 2017 and September 2021.

Chapter 1 contains a general introduction to gravitational waves and the methods of their detection. The limits on the sensitivity of these detectors are discussed, and quantum noise is introduced as a core topic of this thesis. A brief introduction to interferometer simulations and the mathematical framework used throughout this work are also provided.

Chapter 2 details work I carried out studying the implications of the quantum noise target for Einstein Telescope on the design of its infrastructure. This work has been published in: *Philip Jones et al. 'Implications of the quantum noise target for the Einstein Telescope infrastructure design', Phys. Rev. D 101 (8 Apr. 2020) [1]*. I demonstrate how the proposed lengths of various cavities within the interferometer cause the quantum noise limited sensitivity to differ from that desired, and present options to correct these differences.

Chapter 3 investigates the application of output path filter cavities—a quantum noise reduction technique that has seen comparatively little discussion in recent years—to the specific case of Einstein Telescope. This work is awaiting submission for publication under the name *Studies of input and output filter cavities for Einstein Telescope*. I show that there is a case to be made for including output filters in future detectors with low loss optics, although only weakly in the specific case of Einstein

Telescope.

Chapter 4 describes my work using numerical simulation software to verify a novel detector design proposed by a colleague. This verification contributed to the publication of: *Teng Zhang et al. ‘Two-Carrier Scheme: Evading the 3 dB Quantum Penalty of Heterodyne Readout in Gravitational-Wave Detectors’, Phys. Rev. Lett. 126 (22 June 2021) [2]*. I also present derivations of the precise requirements on certain cavity lengths in this scheme.

Chapter 5 provides a brief overview of a quantum noise reduction technique called the optomechanical filter. This filter relies on effects that previous versions of FINESSE, the numerical modelling tool used throughout this thesis, could not model. I describe the extensions myself and others made to FINESSE 3, the latest version of FINESSE, that allow modelling of the filter, and verify that the desired effect is achieved.

Chapter 6 details some work I undertook for an interferometer currently under construction, known as ETpathfinder. My task was to determine the requirements on two sources of noise on the optical input path to the interferometer, to inform near-term decisions about the design. I give a background to ETpathfinder and present these requirements, along with information on which detector parameters influence them.

Chapter 7 first summarises the mathematical framework used in existing versions of FINESSE, explaining how the propagation of signals and quantum noise through an interferometer can be simulated in the frequency domain. I then show how I generalised the quantum noise calculations to be applicable to other sources of noise during the course of reimplementing the quantum noise simulation in FINESSE 3.

Chapter 8 gives a broader overview of the development of FINESSE 3, for which I was part of the core team. Work on FINESSE 3 accounted for roughly half the time I spent doing this PhD. I discuss the motivations behind performing a complete

rewrite of FINESSE, and present the many enhancements myself and others made. I also show how I have already used FINESSE 3 for a detector design task, to great effect.

Appendix A presents a mathematical derivation of the behaviour of the opto-mechanical filter described in Chapter 5. This derivation was carried out in parallel with a colleague, and is included to explain the stated result, which differs slightly from that published elsewhere due to the mathematical conventions used.

Appendix B derives analytical expressions for a laser beam interacting with a suspended mirror via radiation pressure effects for a number of scenarios. These expressions are successfully compared to the results of FINESSE 3 simulations, and the radiation pressure calculations in FINESSE 3 are validated.

Contents

Abstract	i
Acknowledgements	iii
Statement of Originality	iv
Contents	vii
List of Figures	xi
List of Tables	xv
List of Source Code Listings	xvi
Abbreviations	xvii
1 Introduction	1
1.1 Gravitational wave detectors	2
1.2 Sources of noise	6
1.3 Quantum noise	8
1.3.1 Squeezed light	9
1.4 Frequency domain interferometer simulation	13
1.5 Mathematical description of light fields	14
1.6 Thesis overview	15
2 Implications of the Quantum Noise Target for the Einstein Telescope Infrastructure Design	18
2.1 Motivation	21
2.2 Signal Recycling Cavity Length	22

2.2.1	The SRC-Arm System as a Coupled Cavity	24
2.3	Optimised Filter Cavities for ET	28
2.3.1	ET-HF	30
2.3.2	ET-LF	32
2.3.3	Coupled Filter Cavities	35
2.4	Summary	37
3	Output path filter cavities	39
3.1	Frequency dependent squeezing	40
3.2	Motivation	42
3.3	Detuned ET-LF	45
3.4	Tuned ET-LF with combined I/O filter cavities	46
3.5	Dissimilar filter cavities	48
3.6	Summary	50
4	A two-carrier scheme: evading the 3 dB quantum noise penalty of heterodyne readout in gravitational wave detectors	52
4.1	Interferometer readout schemes	53
4.2	The two-carrier scheme	56
4.2.1	Resonance conditions	58
4.2.2	Results	63
4.3	Summary	65
5	Broadening the bandwidth of gravitational wave detectors	66
5.1	The optomechanical filter	68
5.2	Modelling the filter	71
5.2.1	Initial verification in FINESSE 2	71
5.2.2	Improved modelling in FINESSE 3	74
5.3	Summary	75
6	Noise requirements for ETpathfinder	78
6.1	ETpathfinder	78
6.2	Noise sources in interferometer input optics	80
6.2.1	Noise coupling and suppression	82
6.3	Modelling noise requirements	83

6.4	Summary	88
7	Quantum noise modelling in FINESSE 3	90
7.1	Frequency domain modelling	91
7.1.1	Simulating signals	94
7.1.2	Radiation pressure	96
7.2	Quantum noise spectra	98
7.3	Noise spectral density	100
7.4	Photodiode demodulation	103
7.5	Propagation of noise sidebands in FINESSE	106
7.6	The MISO method	107
7.7	The quantum matrix M_q	108
7.8	Generalisation to other sources of noise	109
7.9	Summary	113
8	FINESSE 3	115
8.1	State of the field	116
8.2	Motivation for redesigning FINESSE	117
8.2.1	Using FINESSE 2	118
8.2.2	PYKAT	121
8.3	FINESSE 3 design	123
8.3.1	Choice of language	123
8.3.2	Structure	126
8.3.3	Documentation	132
8.3.4	External tooling	132
8.4	Using FINESSE 3: AdVirgo+ modelling	134
8.5	The future of FINESSE	139
8.5.1	Coupling matrix combination	139
8.5.2	Time domain simulation	140
8.5.3	Polarising optics	141
8.6	Summary	142
9	Summary and conclusions	143
	Bibliography	147

A	Optomechanical filter derivation	155
A.1	Notation	155
A.2	Free Space	157
A.3	Tuned Cavity	161
A.4	Detuned Cavity	166
A.5	Optomechanical Filter	170
B	Radiation pressure tests	175
B.1	Mirrors	175
B.1.1	Single-Sided Free Mass	175
B.1.2	Double-Sided Free Mass, One Modulation	180
B.1.3	Double-Sided Free Mass, Two Modulations	183
B.2	Beamsplitter	187
B.2.1	Single-Sided Free Mass	187
B.2.2	Double-Sided Free Mass, One Modulation	189
B.2.3	Double-Sided Free Mass, Two Modulations	192

List of Figures

1.1	Diagram of a simple Michelson interferometer.	3
1.2	Diagram of a dual-recycled Fabry-Perot Michelson interferometer. . .	4
1.3	Noise budget plot for aLIGO.	6
1.4	Plot of quantum noise and the standard quantum limit.	8
1.5	The effect of squeezing on the uncertainty of light measurement. . . .	10
1.6	Ideal application of frequency dependent squeezing.	12
1.7	Diagram of frequency dependent squeezing injection layout.	12
2.1	The triangular layout of interferometers in ET.	19
2.2	Diagram of a dual-recycled Fabry-Perot Michelson with frequency dependent input squeezing and homodyne readout.	19
2.3	Noise budget for ET, with quantum noise highlighted.	20
2.4	The effect of increasing SRC length on ET-HF.	23
2.5	Fields within the ET SRC-arm coupled cavity system.	25
2.6	The effect of SRC length on the strain-output transfer function of ET-HF.	26
2.7	Options for correcting for an increased SRC length in ET-HF.	27
2.8	Behaviour of quantum noise versus power on the central beamsplitter for a fixed SRC length.	28
2.9	The effect of increasing SRC length on ET-LF.	29
2.10	The effect of filter cavity loss on ET-HF.	31
2.11	The effect of filter cavity loss on ET-LF.	33
2.12	A comparison of ET-LF with and without a detuned SRC.	34
2.13	A comparison of two separate filter cavities with one coupled filter cavity.	36
2.14	A comparison of ET-LF with a detuned and tuned SRC, with 10 km total filter cavity length.	37

3.1	Diagrams of input and output path filter cavities.	41
3.2	Comparison of detuned and tuned signal recycling for ET-LF.	45
3.3	Comparison of combined input-output filtering for ET-LF with both the current design and tuned input filtering.	46
3.4	Sensitivity vs. arm cavity power for combined input-output filtering. .	47
3.5	Sensitivity vs. SRM transmissivity for combined input-output filtering.	48
3.6	Effects of different loss sources on combined input-output filtering. . .	49
3.7	Comparison of different filter cavity lengths for ET-LF.	50
4.1	Example layouts of DC, heterodyne and homodyne readout.	53
4.2	Diagram of frequencies present in homodyne and heterodyne readout schemes.	55
4.3	Layout of the two-carrier scheme.	56
4.4	Diagram of frequencies present in the two-carrier scheme.	57
4.5	Plot of differently located dark fringes for different carriers.	59
4.6	Quantum noise limited sensitivity of heterodyne readout and the two-carrier scheme.	64
5.1	Layout of the transmission-readout setup for the optomechanical filter.	69
5.2	Diagram of the frequencies present in the optomechanical filter setup.	72
5.3	Demonstration of negative dispersion exhibited by the optomechanical filter.	73
5.4	The couplings involved at a mirror in FINESSE 3.	75
5.5	Optomechanical filter layout modelled in FINESSE 3.	76
5.6	Negative dispersion observed in FINESSE 3.	76
6.1	Physical layout of ETpathfinder.	80
6.2	Simplified layout of ETpathfinder.	81
6.3	Laser frequency noise requirement as a function of Schnupp asymmetry.	85
6.4	Laser frequency noise requirement as a function of DC offset.	85
6.5	IMC motion requirement as a function of Schnupp asymmetry.	86
6.6	IMC motion noise requirement as a function of DC offset.	86
6.7	Noise requirements for laser frequency and IMC motion noise in ETpathfinder.	87
6.8	IMC motion noise requirements vs ITM R-T imbalance for one arm. .	88

7.1	The coupling matrix for the fields at a mirror.	91
7.2	Basic setup for a laser reflecting from a mirror.	93
7.3	Simplified flowchart of a FINESSE simulations	95
7.4	Diagram of vacuum noise fluctuations with and without a carrier field.	99
7.5	Diagram of the frequencies involved in photodiode demodulation of three carriers with signal sidebands.	103
7.6	Example layout injecting multiple classical noises into one model simultaneously.	112
8.1	Example of a Jupyter Notebook.	135
8.2	KatScript kernel in a Jupyter Notebook.	136
8.3	Black-boxing of multiple components.	140
A.1	The frequencies present in the optomechanical filter setup.	171
B.1	Amplitude modulated laser reflecting from a free mass.	175
B.2	Phase change of a reflected beam in FINESSE.	178
B.3	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.1	181
B.4	Radiation pressure inducing modulation of a laser beam incident on the opposite side of a mirror.	181
B.5	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.2	183
B.6	Radiation pressure caused by laser beams incident on both sides of a mirror.	184
B.7	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.3 for the reflected field on the left of the mirror.	186
B.8	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.3 for the reflected field on the right of the mirror.	187
B.9	Amplitude modulated laser reflecting from a free mass at an angle.	188
B.10	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.1	190
B.11	Radiation pressure inducing modulation of a laser beam incident on the opposite side of a beamsplitter.	190
B.12	Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.2	192
B.13	Radiation pressure caused by laser beams incident on both sides of a beamsplitter.	193

B.14 Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.3 for the reflected field at the top detector.	195
B.15 Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.3 for the reflected field at the right detector.	195

List of Tables

2.1	Optimal filter cavity parameters for ET-HF.	32
2.2	Optimal filter cavity parameters for ET-LF.	32
2.3	Optimal filter cavity parameters for ET-LF, for both two filter cavities and a coupled filter cavity.	35
4.1	Phase contributions to the fields in the PRC in the two-carrier scheme.	62
4.2	Conditions on various parameters necessary for the two-carrier scheme.	63

List of Source Code Listings

8.1	Kat file for a simple Michelson setup with curved, suspended mirrors.	119
8.2	Example usage of PYKAT.	122
8.3	Output of Cython when run on unannotated Python code.	125
8.4	Output of Cython when run on fully type-annotated Python code. . .	125
8.5	FINESSE 3 KatScript translation of the file shown in Listing 8.1. . . .	127
8.6	Comparison of relative parameter definition in FINESSE 2 and FINESSE 3.	128
8.7	Minimal example of chaining actions in FINESSE 3.	131
8.8	Example usage of actions in FINESSE 3 to pretune an interferometer.	138
B.1	KatScript representation of the setup shown in Figure B.1.	180
B.2	KatScript representation of the setup shown in Figure B.4.	182
B.3	KatScript representation of the setup shown in Figure B.6.	186
B.4	KatScript representation of the setup shown in Figure B.9.	189
B.5	KatScript representation of the setup shown in Figure B.11.	191
B.6	KatScript representation of the setup shown in Figure B.13.	194

Abbreviations

aLIGO	Advanced LIGO
au	Arbitrary Units
DRFPMI	Dual-Recycled Fabry-Perot Michelson Interferometer
ET	Einstein Telescope
ET-HF	ET High Frequency
ET-LF	ET Low Frequency
ETM	End Test Mass
FSR	Free Spectral Range
IMC	Input Mode Cleaner
ITM	Input Test Mass
MISO	Multiple-Input Single-Output
OPO	Optical Parametric Oscillator
PRC	Power Recycling Cavity
PRM	Power Recycling Mirror
PSD	Power Spectral Density
RF	Radio Frequency
RHS	Right-Hand Side
SQL	Standard Quantum Limit
SRC	Signal Recycling Cavity
SRM	Signal Recycling Mirror

Chapter 1

Introduction

Gravitational waves are ripples in space-time, predicted by Einstein's general theory of relativity, and first detected in 2015 [3]. This detection marked the beginning of a new era for astronomy in which we can directly observe objects invisible to traditional electromagnetic telescopes, such as black holes. These observations allow us to test general relativity and quantum mechanics at their most extreme, pushing the boundaries of known physics, and deepening our understanding of the universe.

Gravitational waves are emitted by almost any body undergoing acceleration. They are, however, extraordinarily weak. A leading order post-Newtonian approximation to the amplitude h of a gravitational wave is given by the quadrupole formula [4],

$$h = \frac{2G}{c^4 r} \frac{d^2 I}{dt^2}, \quad (1.1)$$

where G is the gravitational constant, c the speed of light, r the distance from the observer to the source, and I is the mass quadrupole moment of the source. h is unitless, and is also known as the gravitational wave *strain*; a passing wave stretches and squashes a space perpendicular to the direction of travel of length L

by a maximum amount δL , such that

$$h = \frac{\delta L}{L}. \quad (1.2)$$

The term at the beginning of Equation (1.1) should give anyone cause for concern, as $G/c^4 \approx 8.26 \times 10^{-45} \text{ s}^2 \text{ kg}^{-1} \text{ m}^{-1}$. The hopes of detecting gravitational waves caused by ‘everyday’ or man-made objects are therefore nil. Instead, we must turn to systems many orders of magnitude more massive, and undergoing extreme accelerations—the inspirals and mergers of binary black holes and neutron stars. The first direct detection of gravitational waves came from the merger between two orbiting black holes of masses 30–40 M_{\odot} [3]. Just before merging, they were orbiting at around 250 Hz, moving at over half the speed of light relative to each other. During the event, 3 M_{\odot} was radiated away as gravitational waves. At our distance of ~ 410 Mpc, this resulted in a peak strain of 1×10^{-21} . Even for these high energy events, then, we require detectors with exquisite sensitivity.

1.1 Gravitational wave detectors

All gravitational wave detectors currently in operation [5–8] are based on the well-known Michelson interferometer, shown in Figure 1.1. This acts as a differential length sensor. Any relative change in the distances from the central beamsplitter to each of the end mirrors will result in a change in the phase of the returning light, and thus the amount of light incident on the photodiode. A passing gravitational wave will alternately stretch and squash each space in the interferometer, with a component of the effect being out of phase in each arm (except in the special case of a wave approaching diagonally in the plane of the detector). For a sinusoidal gravitational wave signal with frequency Ω , then, the power detected by the photodiode will oscillate at 2Ω . It is instructive to treat the gravitational wave as modulating the




Figure 1.1: Basic Michelson interferometer. In (a), the distances from the beamsplitter to each mirror are the same, and the light will recombine at the beamsplitter such that all of the power is sent to the output photodiode. In (b), one of the distances is longer by a quarter of the laser wavelength λ . This results in a relative phase change of 180° in the reflected light, causing destructive interference on the photodiode side of the beamsplitter, and all of the power is sent back towards the laser.

laser beam, creating two extra fields called *sidebands* at $\pm\Omega$ relative to the laser's frequency. The oscillating power at the output is then the beating between these two sidebands. Some more detail on this is given in Section 1.5

The basic Michelson design shown above is insufficient, however. Equation (1.2) tells us that the longer the arms of the Michelson, the more sensitive to a passing gravitational wave it will be. To detect a strain of 1×10^{-21} , the arms would have to be impractically long, at hundreds or even thousands of kilometres. Modern gravitational wave detectors therefore make a number of additions to the Michelson design, some of which are shown in Figure 1.2. Each of these additions serve a specific purpose, which I will detail here.

Fabry-Perot arm cavities

To avoid requiring arms that are extremely long, we can introduce an additional low-transmissivity mirror in each one. This converts the arms into Fabry-Perot cavities. Light inside these cavities bounces back and forth many times on average, effectively travelling a greater distance. As long as the time spent in the arm cavities is less than half the period of the gravitational wave, this will result in a greater




Figure 1.2: A simplified diagram of the type of Michelson interferometer used at aLIGO [5] and AdVirgo [6]. An incoming gravitational wave alternately stretches and squashes the space in each arm out of phase with the other, changing the optical path length, and resulting in a signal incident on the output detector. The main extensions over a basic Michelson shown are the conversion of the arms into Fabry-Perot cavities, which increase the effective path length and power of the light in the arms, the *Power Recycling Mirror* (PRM), which further increases the circulating power, and the *Signal Recycling Mirror* (SRM), which amplifies signals at a desired frequency at the cost of detection bandwidth. This design is therefore known as a dual-recycled Fabry-Perot Michelson interferometer (DRFPMI). ITM and ETM stand for input and end test mass respectively, with the arms labelled as X and Y.

phase change due to a passing gravitational wave, and therefore a stronger signal. In Advanced LIGO [5], the arm cavities are 4 km in length, and light makes somewhere around 100 round trips.

The arm cavities also serve to increase the power of the beam that is being modulated by a gravitational wave, resulting in a stronger signal arriving at the output. This is explained in more detail in Section 1.3.

Power recycling mirror

For various reasons that will be explained in later sections, gravitational wave interferometers are operated at, or close to, the dark fringe—the point at which no light is sent to the output photodiode. Correspondingly, the vast majority of the power reflected from the arms is sent back towards the laser, and is effectively ‘wasted’. The power recycling mirror (PRM) does exactly what its name suggests; it ‘recycles’ the waste power, by reflecting it back towards the main interferometer, further increasing the power in the arm cavities and improving sensitivity. The cavity formed by the PRM and the input test masses (ITMs) is known as the power recycling cavity (PRC).

Signal recycling mirror

The operation of the signal recycling mirror (SRM) is somewhat complex, but the motivation is fairly straightforward. As mentioned, the strength of the signal depends on how long light circulates in the arm cavities in relation to the period of an incident gravitational wave. By inserting a mirror into the output path of the interferometer, signal light can be sent back into the arm cavities to circulate further, altering sensitivity. By carefully choosing the position and properties of the SRM, the detector can be tuned to be most sensitive to a specific frequency of gravitational wave. The cavity formed by the SRM and the ITMs is known as the signal recycling cavity (SRC).

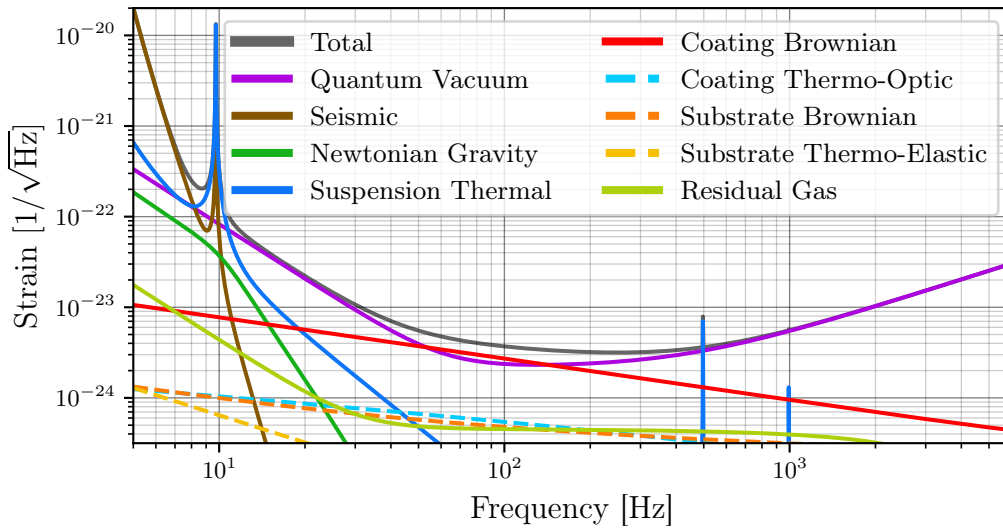


Figure 1.3: Noise budget plot for aLIGO [5], showing many of the noise sources present. The x-axis shows the gravitational wave frequency. The y-axis represents the equivalent strain amplitude spectral density of the noise—that is, the square root of the power distribution of the noise into each frequency, normalised such that the value is equivalent to that produced by a gravitational wave with the given strain. In very broad terms, the ‘Total’ line represents the weakest gravitational wave strain that could be observed, although other factors such as length of observation time are important. This plot was produced using `pygwinc` [9], a noise calculator for gravitational wave interferometers.

1.2 Sources of noise

Unfortunately, gravitational waves are not the only thing that will cause an oscillating signal to be detected at the photodiode in Figure 1.2. Many other effects will also contribute, resulting in a background of noise. In order to detect a gravitational wave, we must see a signal rise above this noise. A useful tool to quantify how much different sources of noise will limit the sensitivity of a detector to gravitational waves is a noise budget, an example of which is shown in Figure 1.3. From this plot, we see that different noise sources limit the sensitivity of the detector to different gravitational wave frequencies. Some of the key noises shown here are:

- Quantum vacuum or quantum noise — this is noise arising from the inherent uncertainty in the amplitude and phase of the laser beam as a result of the Heisenberg uncertainty principle. It is the limiting source of noise at high frequencies for all Michelson interferometers, and can also be in the low-mid

frequency range. The majority of this thesis is focused on quantum noise, and a more detailed background is provided in Section 1.3 and Chapter 7.

- Coating Brownian noise — thermal motion of the atoms in the surface of the arm cavity mirrors results in modulation of the reflected beam. As shown in Figure 1.3, aLIGO is limited by this thermal noise in the middle of its frequency range around 10 Hz, where it would otherwise be most sensitive. The latest detector to come online, KAGRA [7], and future detectors such as Einstein Telescope (ET) [10], make use of cryogenically cooled test masses to reduce thermal noise. ET will also use larger diameter laser beams in the arm cavities than current detectors, in an effort to reduce the effects of thermal motion.
- Seismic noise — ground motion shakes the mirrors in the arm cavities, changing the length of the arms. To combat this, modern interferometers hang the arm cavity mirrors on multiple suspension stages [11]. These greatly suppress the effect of ground motion at high frequencies, apparent from the steep curve in Figure 1.3. Active control systems are also used to cancel out the effects of ground motion. At very low frequencies, however, seismic noise remains the limiting factor.
- Suspension thermal noise — similarly to coating Brownian noise, thermal motion of the wires used to suspend the mirrors causes them to vibrate, modulating the reflected laser beam.

Many more noise sources than those shown in Figure 1.3 are also present in a gravitational wave interferometer. I discuss a couple of these, namely laser frequency noise and input mode cleaner motion noise, in Chapter 6. The majority of this thesis, however, is devoted to quantum noise and its reduction, and so we should gain a little more background before proceeding.

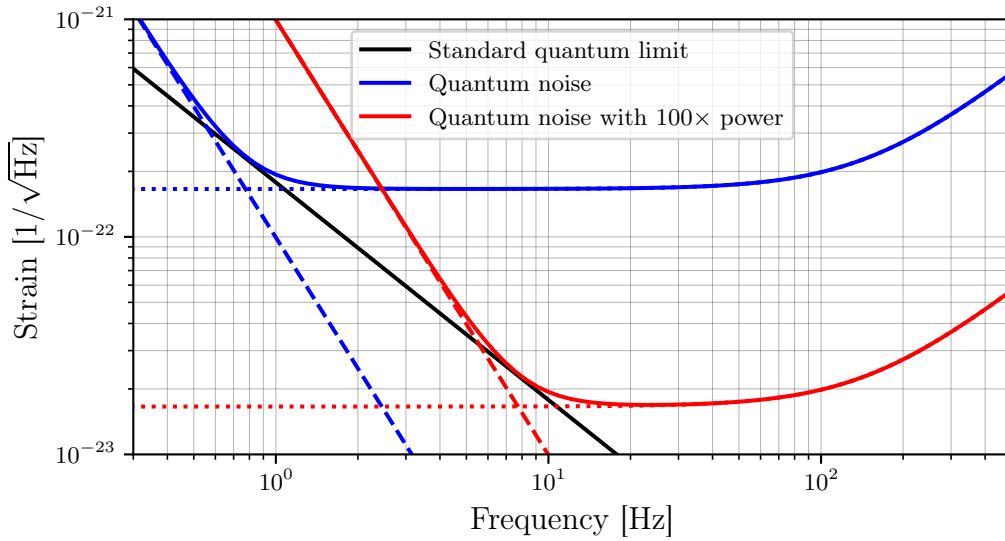


Figure 1.4: Quantum noise limited sensitivity for an example interferometer with varying arm cavity power. The overall quantum noise curve can be understood as the combination of two effects: radiation pressure noise (dashed lines), which has a $1/\Omega^2$ frequency dependence; and shot noise (dotted lines), which is flat with a high-frequency tail due to the limited bandwidth of the arm cavities. The point at which they are equal lies on a line known as the standard quantum limit (SQL), which cannot be beaten with classical techniques such as changing circulating arm cavity power or tuning of mirrors.

1.3 Quantum noise

Quantum noise is a result of the inherent uncertainty in the amplitude and phase of the laser beam. This results in different effects at different parts of the interferometer, and so is frequently broken down into the two contributions shown in Figure 1.4, known as shot noise and radiation pressure noise.

Shot noise can be thought of as due to the photodiode counting the rate of incoming photons at the output port of an interferometer. This rate is subject to Poisson counting statistics, so there will always be an error $\propto 1/\sqrt{P}$ [12], where P is the power in the signal sidebands at the detector. While shot noise is constant at all frequencies, the arm cavities have a finite bandwidth. For laser frequencies outside this bandwidth, the circulating power will fall like $1/\Omega^2$. At higher gravitational wave signal frequencies, the sidebands created by the modulation of the laser beam will move out of the arm cavity resonance, therefore experiencing this $1/\Omega^2$ drop and increasing shot noise.

Radiation pressure noise arises from the circulating light in a gravitational wave detector's arm cavities exerting a force on the suspended end mirrors. As this light is subject to random fluctuations, the force exerted will also fluctuate randomly, leading to oscillations of the end mirrors, and therefore noise in the output. The motion of the mirrors increases with $\sqrt{P_{\text{circ}}}$ [12], with P_{circ} being the circulating power in the arms. This motion also falls off as frequency increases, as for very small oscillations, the mirrors act like free masses with a mechanical susceptibility that scales like $1/\Omega^2$.

From these descriptions, we can see that they follow opposite scaling relations with the power in the detector. Adjusting the power therefore leads to a trade-off between shot noise and radiation pressure noise, with a minimum that follows the standard quantum limit (SQL). aLIGO is close to operating at this point, with few classical upgrades remaining. The SQL cannot be surpassed with classical techniques, and so quantum effects must be used to further increase the sensitivity of gravitational wave detectors.

1.3.1 Squeezed light

One of the most promising techniques for reducing quantum noise below the SQL is the use of squeezed states of light [13–16]. A coherent laser beam can be described in terms of two quadratures: its amplitude and phase. The fluctuations that cause shot and radiation pressure noise arise from the laser beam interacting with the quantum vacuum state. This vacuum state has some inherent uncertainty in its amplitude and phase quadratures. Normally there is no correlation between these uncertainties, corresponding to equal uncertainty in both. While the total uncertainty cannot be reduced, as it is limited by the Heisenberg uncertainty principle, the uncertainty in one quadrature can be reduced (squeezing) at the expense of increased uncertainty (anti-squeezing) in the other. This is shown in Figure 1.5. By phase squeezing light, we can reduce the error in the arrival rate of photons at the detector, reducing high-




Figure 1.5: The effect of squeezing on the uncertainty of light measurement. The ellipse in the centre of each diagram represents the vacuum state, and the offset ellipse a coherent state. $\Delta x_{1,2}$ represent arbitrary conjugate quadratures of light. In (a), before squeezing, the noise is uncorrelated and there is equal uncertainty in both amplitude and phase quadratures. In (b), the vacuum state has been amplitude squeezed, reducing the uncertainty in the amplitude of the coherent state, but increasing the uncertainty of the phase.

frequency shot noise. The amplitude fluctuations of the field will increase, however, increasing radiation pressure noise. The opposite is true for amplitude squeezed light. The amount of phase vs. amplitude squeezing in a field is characterised by the angle of the squeezed light, corresponding to the rotation angle of the ellipse relative to the laser beam in Figure 1.5.

Squeezing is often achieved via the use of a device such as an optical parametric oscillator (OPO) [13, 17]. Here, a nonlinear medium such as a crystal is pumped with light at twice the desired laser frequency, which then undergoes parametric down-conversion, producing pairs of photons at different frequencies around the desired frequency [18]. These photons are correlated, such that there is squeezing in their subtraction. As the amount of squeezing generated on a single pass is small, the nonlinear medium is placed in a cavity, to feed light through multiple times, and remove undesired spatial modes. The limiting factor on the amount of squeezing is the loss of the mirrors used for the cavity, as any loss reintroduces the unsqueezed vacuum state, leading to a reduction in squeezing. Recently, it has been proposed that the cavity surrounding the medium is unnecessary, creating a mirrorless OPO (MOPO) [19]. If investigations into MOPOs prove fruitful, they could represent a

simpler and more compact source of squeezed light for gravitational wave detectors.

So far, OPOs have been successfully used to inject squeezed light at GEO600 [8] for the past 11 years, and more recently at LIGO [15] and AdVirgo [20]. At GEO600, the peak improvement in shot noise limited sensitivity has been ~ 4 dB, or a factor of ~ 2.5 , with a duty of $\sim 85\%$. At aLIGO, improvements of up to 2.15 dB (a factor of ~ 1.6) in sensitivity have been observed. These results demonstrate that squeezing is a both a viable and worthwhile improvement to gravitational wave detectors.

In all of these cases, the squeezing used has been frequency independent, where the squeezing angle is constant for all frequencies. This works well at very high and very low frequencies, where quantum noise is dominated by either radiation pressure noise or shot noise, and so the corresponding increase in noise of the anti-squeezed quadrature is of little consequence. However, the key area of interest for detectors is the middle frequency range at the SQL, where radiation pressure noise and shot noise are of roughly equal amplitudes. Here, frequency independent squeezing is insufficient, so we must turn to frequency dependent squeezing [21, 22].

Frequency dependent squeezing is a technique where the squeezing angle of the squeezed vacuum is adjusted depending on the gravitational wave frequency range being investigated. This allows us to reduce the quantum noise present across the entire detection frequency range, by amplitude squeezing at low frequencies and phase squeezing at high frequencies. The ideal application of frequency dependent squeezing is shown in Figure 1.6. We can implement frequency dependent squeezing by reflecting a frequency independent squeezed state from a carefully tuned Fabry-Perot cavity, known as a filter cavity [21]. The reflection causes a frequency dependent change in the phase of each squeezed sideband, thereby changing the squeezing angle ϕ . The reflected sidebands are then injected into the detector from the output port via the use of a Faraday isolator or similar, as shown in Figure 1.7.

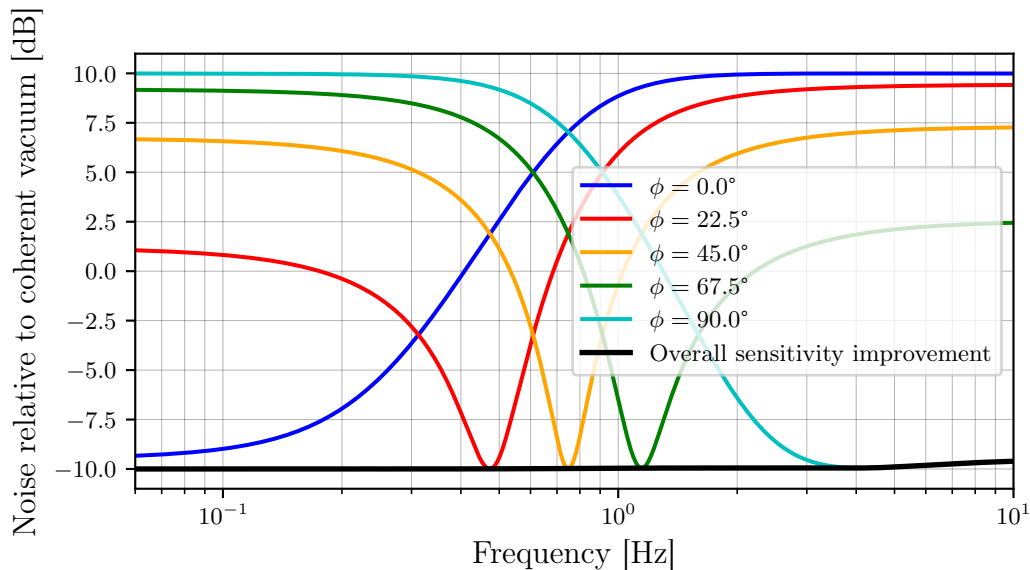


Figure 1.6: Sensitivity improvement from the ideal application of frequency dependent squeezing. 10 dB of squeezing is injected. The coloured lines show the change in quantum noise at a fixed squeezing angle ϕ , where improvements are made only within a limited frequency range, and anti-squeezing is observed elsewhere. The black line shows the ideal case for frequency dependent squeezing, where the optimum squeezing angle is used for each frequency. Curves such as these can be observed via measurement of the injected frequency dependent squeezed state with a balanced homodyne detector and varying the homodyne angle; see [23, 24] for practical demonstrations.




Figure 1.7: (a) Injection of frequency dependent squeezing at the output port of an interferometer. Squeezed light is first reflected off a filter cavity to induce a frequency-dependent rotation of the quadratures, and then sent towards the interferometer via a Faraday isolator. (b) Enlarged view of a Faraday isolator. Polarising optics are used to send light in the indicated directions.

Frequency dependent squeezing has been demonstrated recently in the audio band necessary for gravitational wave detection in a 300 m experiment [25]. Improvements in quantum noise range from 1–3 dB across the entire frequency range, representing a significant increase in detector sensitivity. Further investigation is required to scale this up to the sizes required for detectors, such as the proposed 1 km filter cavities in ET [10], and to reduce the squeezing degradation due to filter cavity losses and noises.

1.4 Frequency domain interferometer simulation

Numerical simulations play a critical role in the design and investigation of modern gravitational wave detectors. They can be used to test new ideas before committing to a physical implementation, and to see the effects of changes to existing designs.

Very broadly speaking, there are two main types of general-purpose simulation used for gravitational wave detectors. These are time domain simulations, in which the evolution over time of light fields in the interferometer are modelled, and frequency domain simulations, which assume the system is in a steady state. Frequency domain simulations cannot capture complex dynamics and non-linear effects, but are still suitable for a wide range of modelling tasks. The treatment of fields as being in a steady state also allows them to be much more computationally efficient than time domain simulations.

Multiple examples of each type of simulation exist, along with other more specialised tools. For the purposes of this thesis, however, we only need to consider one: the frequency domain modelling tool FINESSE [26]. Chapters 2 to 6 all present results of simulations performed with FINESSE, demonstrating the wide applicability of the tool. Chapter 7 gives a detailed mathematical background for the approach FINESSE takes to modelling interferometers, and Chapter 8 provides a general overview of the features in and improvements made in the latest version of FINESSE.

1.5 Mathematical description of light fields

A general description for the electric field component of an electromagnetic wave must take into account many factors: the direction of travel, the polarisation, the transverse spatial distribution of the field, and more. For many purposes, however, this level of detail is unnecessary. Throughout this thesis, the description of the electric field is simplified by making a number of approximations that allow for easier modelling and derivation of various effects. The exact model used is that of scalar plane waves propagating only along a well-defined optical axis. The electric field at a point z on the optical axis at time t can then be expressed as

$$E(z, t) = E_0 \cos(\omega t - kz + \varphi), \quad (1.3)$$

where E_0 is the amplitude of the field, $\omega = 2\pi f$ is the angular frequency, $k = \omega/c$ is the wavenumber, and φ the phase of the field relative to some reference field. This can also be written in complex notation [27],

$$E = \Re\{E'\} \quad \text{where} \quad E' = E_0 e^{i(\omega t - kz)}, \quad (1.4)$$

Here, the scalar amplitude and phase are combined into a complex amplitude $E'_0 = E_0 e^{i\varphi}$. It is important to note that the complex amplitude is not a real amplitude—only $\Re\{E'\}$ is the true field amplitude. With that in mind, the prime will be dropped from here on for simplicity. This complex notation is the form used in the frequency domain interferometer simulation tool FINESSE [26], the use and development of which is a core theme of this thesis. Both the forms for the electric field from Equations (1.3) and (1.4) will be used in this thesis at different points; which one is in use shall be clear during usage.

A passing gravitational wave stretching and squashing the length of a space at a frequency Ω will cause phase modulation of a laser beam propagating through that space,

$$E = E_0 e^{i(\omega t - kz + m \cos(\Omega t))}, \quad (1.5)$$

where m is the modulation depth. For small modulations, we can expand this to leading order in m , giving

$$E = E_0 e^{i(\omega t - kz)} \left(1 + i \frac{m}{2} (e^{i\Omega t} + e^{-i\Omega t}) \right). \quad (1.6)$$

The notion of signal sidebands then appears naturally from the three frequency components present: the main carrier at ω , and two small sidebands at $\omega \pm \Omega$.

1.6 Thesis overview

Gravitational waves were first detected only a few years ago by the current generation of interferometers, marking the beginning of a new era for astronomy. To explore deeper into the universe, these detectors must be upgraded, reducing the effects of the different noise sources that limit sensitivity. Looking further ahead, the next generation of gravitational wave detectors already need to be developed, with the focus currently on ET [10] in Europe and Cosmic Explorer [28] in the U.S. This development is a combination of multiple areas of work: new conceptual designs that propose the use of novel technologies and layouts, instrument research and development to improve current technologies and enable new ones, and the construction of prototype systems to test new ideas and build technical expertise. During the course of my PhD, I have focused on one particular aspect of this development work; the reduction of quantum noise in future detectors.

Chapters 2 and 3 take a look at this quantum noise reduction in the specific case of ET. In Chapter 2, the effect on quantum noise of practical limits in the proposed

ET design are studied. I show that the effect on quantum noise from increasing the length of the SRC compared to current detectors, which was previously neglected, is significant enough to warrant consideration. I also discuss the practicality of the design of the filter cavities used for frequency dependent squeezing. In Chapter 3, I investigate an alternative squeezing scheme applied to the specific case of ET, and show the motivations for and pitfalls of the scheme in this context.

Chapters 4 and 5 feature more novel noise reduction techniques, which I have used FINESSE to verify. In Chapter 4, I discuss a new method for reading out the signal from a gravitational wave interferometer with improved noise performance, known as the two-carrier scheme, and present the modelling work I performed to verify the idea. In Chapter 5, we see a quantum noise reduction technique using complex optomechanics to avoid the high-frequency shot noise tail, and I describe the extensions I made to the latest version of FINESSE to allow modelling of the technique.

Before ET is built, many of the proposed technologies will be tested in the prototype interferometer ETpathfinder [29]. This is currently under construction, with some of the finer details of the optical layout still to be decided. Chapter 6 details modelling work I undertook for ETpathfinder. My task was to determine the requirements on two sources of noise on the optical input path to the interferometer, to inform near-term decisions about the design. I present these requirements, along with information on which detector parameters influence them.

Finally, Chapters 7 and 8 dive into the development of FINESSE, the tool used throughout the rest of this thesis. Chapter 7 takes a more mathematical approach, describing how FINESSE models interferometers in general, and how this approach is extended to quantum noise. I then present further extensions I made when performing a complete reimplementaion of the quantum noise calculations in FINESSE 3, the latest version of FINESSE. These extensions permit the modelling of other incoherent

sources of noise in FINESSE, increasing the range of simulations for which FINESSE is suitable. Chapter 8 provides a broader overview of the my role in the development of FINESSE 3 and the tools around it. The motivations behind a complete rewrite of FINESSE are given, and I present the many enhancements made during this rewrite, opening the doors for even more advanced simulation work in years to come.

Chapter 2

Implications of the Quantum Noise Target for the Einstein Telescope Infrastructure Design

This chapter is a modified version of my paper ‘Implications of the Quantum Noise Target for the Einstein Telescope Infrastructure Design’ [1]. I am the principal author of this paper, and carried out all of the simulation work required. A brief introduction to Einstein Telescope has been added at the beginning, including Figure 2.1 and Figure 2.3 which were not present in the paper. Additional minor format and wording changes have been made to better fit this thesis.

As we have so far seen, current gravitational wave detectors such as aLIGO [5] and Advanced Virgo [6] make use of a dual-recycled Fabry-Perot Michelson interferometer (DRFPMI) design. Current detectors also make use of frequency-independent squeezing to increase quantum-noise limited sensitivity [30–32].

In this chapter, we will be considering the design of a future gravitational wave detector, called Einstein Telescope (ET) [33–36], shown in Figures 2.1 and 2.2. ET contains a few key additions over today’s detectors with regards to quantum




Figure 2.1: The triangular arrangement of interferometers in ET. Each corner represents one of the three detectors, with each detector comprised of two interferometers. This figure is a recreation of Figure 1.8 from [10].

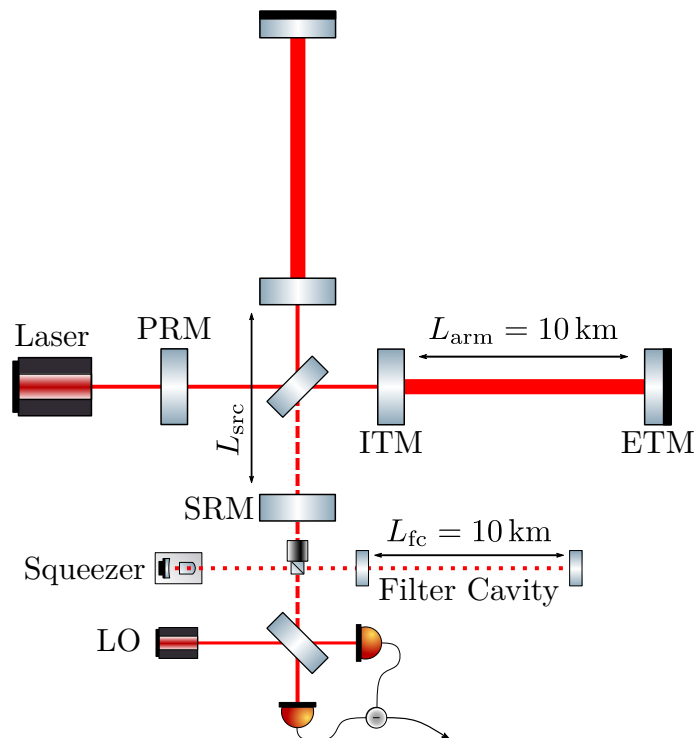


Figure 2.2: Simplified interferometer design considered in this chapter, with parameters shown for ET-HF. The signal recycling cavity length L_{src} is the total distance between the signal recycling mirror (SRM) and the input test mass (ITM). The squeezer is a frequency-independent squeezed vacuum source. ET makes use of a balanced homodyne detection scheme, comprised of two photodiodes and a local oscillator (LO). For ET-LF the only addition is an extra filter cavity.




Figure 2.3: The total noise budget for ET presented in the original design report [35] (dash-dotted line), with the quantum-noise limited sensitivity for each of ET-HF and ET-LF highlighted. The broadband response of ET-HF and narrowband response of ET-LF can be seen.

noise. The largest change is that ET is not a single detector; instead it is made of three identical detectors, each comprised of two interferometers, for a total of six interferometers. In each of these detectors, one interferometer is aimed at high frequencies of $10\text{--}10^4$ Hz, and the design is known as ET-HF. The other interferometer is aimed at low frequencies of $3\text{--}30$ Hz, and is known as ET-LF. This combination of a high and low frequency interferometer in one detector with partially overlapping frequency ranges is referred to as a xylophone configuration [37].

As in current detectors, ET-HF will operate with its signal recycling mirror (SRM) tuned, giving a broadband response [38]. In contrast, ET-LF will operate with a detuned SRM, and will have a narrow-band response. The effects of this can be seen in Figure 2.3. The use of a detuned SRM also has implications for the use of frequency-dependent squeezed light in ET. Whereas ET-HF can achieve the necessary rotation of squeezed input with only one filter cavity, ET-LF requires two [21] due to its detuned SRM.

There are some aspects of ET’s design which, although not motivated by quantum noise considerations, do have an impact on the options available to improve the

quantum-noise limited sensitivity. The rest of this chapter will introduce these factors, and explore the extent to which they affect the quantum-noise limited sensitivity of ET, as well as available mitigation strategies.

2.1 Motivation

A fundamental property of interferometers is the trade-off between bandwidth and peak sensitivity, known as the Mizuno Limit [39]. The finite bandwidth of the arm cavities arises due to the gravitational wave signal light gaining extra phase with increasing frequency (phase dispersion), and eventually no longer resonating. This is a property of all cavities, including the signal recycling cavity (SRC). Up until now, the length of the SRC has been largely ignored with regards to quantum noise, as it is often negligible compared to the length of the arms (~ 55 m in aLIGO, compared to 4 km arm cavities). This allowed us to treat the response of the SRC as practically instantaneous relative to the arms, and thus the whole ITM-SRM system can be thought of as a single compound mirror. For future detectors, especially ET, this may no longer be the case; the initial proposed SRC length is 300 m [35]. We therefore need to understand the consequences of a non-negligible SRC length on detector design.

Another variable worth investigating is the length of the filter cavities in ET. Whereas the ET design study assumed 10 km long filter cavities, in the more recent design update a reduction of this length to 1 km for ET-LF and 300 m for ET-HF has been considered [36]. As the performance of filter cavities is determined solely by their loss per unit length [21, 40], it is necessary to understand how this reduction in length would affect the quantum-noise limited sensitivity of ET.

In the rest of this chapter, we shall see that increasing SRC length leads to a reduction in quantum-noise limited sensitivity at high frequencies, impacting

the performance of ET-HF. We will also see that attempting to compensate for this reduced sensitivity leads to increased laser power on the central beamsplitter and through the ITMs, which gives rise to thermal lensing and other undesirable effects [41]. For these reasons, a reduction of SRC length from 300 m to 100 m for ET-HF is proposed, and a range of options for compensating for the decreased performance given. We shall also see that the reduction of filter cavity length from 10 km has little effect on the performance of ET-HF, but is much more severe for ET-LF, leading to a factor of 2 reduction in quantum-noise limited sensitivity at 7 Hz for currently achievable optical losses. Two possible alternatives with improved quantum noise performance are given, based on the possibility of having a single 10 km filter cavity per ET-LF interferometer. These are the use of a tuned SRM in ET-LF, for which only a single filter cavity is required, and the use of a coupled cavity in place of two separate filter cavities for a detuned ET-LF.

2.2 Signal Recycling Cavity Length

There are a few motivating factors for an increased SRC length in ET compared to that of current detectors. In the arm cavities, a relatively large beam radius is required in order to reduce coating thermal noise [42]. This is especially important for ET-HF, which is almost entirely limited by coating thermal noise around 40–200 Hz. At the central beamsplitter, however, a small beam radius is desirable; it would allow smaller optics and better control of scattered light in the central interferometer. In order to achieve such a change in the beam sizes, a lens or telescope must be placed between the ITMs and the beamsplitter. A short distance between the ITMs and beamsplitter, and hence a short SRC, would require stronger focusing elements with more stringent optics requirements to avoid introducing aberrations and noise. Another factor leading to a long SRC is the use of cryogenic mirrors in ET-LF. To achieve sufficient cooling of the ITMs, cryoshields along the vacuum tubes are




Figure 2.4: Effect of increasing L_{src} on ET-HF, without changing other interferometer parameters. For the value $L_{\text{src}} = 300$ m proposed in the ET Design Study [35], quantum noise deviates from the design curve significantly above 2 kHz, reaching a factor of 2 increase at ~ 7.5 kHz.

required to reduce the solid angle under which the cold ITMs are exposed to room temperature parts of the instrument. The lengths of the cryoshields (several tens of meters) also add to the SRC length.

Previous models used throughout the collaboration, and in the ET Design Study [35], assumed that the SRC length can be neglected. This is no longer valid. We therefore need to model and understand what effects SRC length has on the sensitivity of ET, and how to choose optimal SRC parameters for a given length. Figure 2.4 shows the effects of different SRC lengths on the quantum noise performance of the example ET-HF setup from Figure 2.2, compared to that given in the ET Design Study (known as ET-D). All modelling throughout this chapter was performed using the frequency-domain modelling software FINESSE [26, 43]. It should be noted that our investigations do not yet include the effects of higher-order

modes and beam shapes. These effects should be studied further in the future, as changing the length of the SRC impacts the detector design in other ways as well, for example, with respect to avoiding higher-order mode resonances and parametric instabilities [44].

From Figure 2.4, we see that increasing the SRC length L_{src} leads mostly to a change in the gradient of the quantum noise performance at high frequencies, with a small increase in sensitivity at certain other frequencies. This arises due to a change in the coupled cavity dynamics of the interferometer, and to understand how to compensate for it, we must first understand exactly why this effect occurs.

2.2.1 The SRC-Arm System as a Coupled Cavity

For $L_{\text{src}} \ll L_{\text{arm}}$, where L_{arm} is the interferometer arm length, we can treat the SRM-ITM system as a kind of compound mirror, the only effect of which is to alter the output from the arm cavities. This is described in Buonanno & Chen [45], which for the tuned SRM case gives the half-bandwidth of the SRC-arm system as

$$\gamma_0 = \frac{1 + r_{\text{sr}}}{1 - r_{\text{sr}}} \gamma_{\text{arm}}, \quad (2.1)$$

where r_{sr} is the amplitude reflectivity of the SRM, and $\gamma_{\text{arm}} = cT_{\text{itm}}/4L_{\text{arm}}$ is the half-bandwidth of the arm cavity, with T_{itm} the power transmissivity of the ITMs. When the SRC is comparable in length to the arm cavities, complicated coupled cavity effects come into play. We should therefore have a brief look at the basic properties of coupled cavities before proceeding.

The distinguishing feature of a coupled cavity is the presence of a split resonance. A single cavity exhibits an infinite number of equally spaced resonances, where the frequency difference between consecutive resonances is known as the free spectral range (FSR). A coupled cavity consists of two cavities, each with their own FSR, and




Figure 2.5: Simplified ET coupled cavity setup. Each of the E_n are the full field at that point, including all frequencies. Here we are interested only in observing how the value of L_{src} affects the shape (not absolute value) of the strain-output transfer function of the interferometer, in a numerical simulation. In this model the only purpose of the laser is to provide power in the arm cavity, so it can be placed at either end of the setup.


will exhibit resonance peaks whenever a field is resonant in either of these cavities. For a field that is resonant in both cavities (i.e. every common multiple of both FSRs), a split resonance can occur, where two closely-spaced resonance peaks are observed instead of one. In the case where the two cavities have the same length, a derivation of the frequency difference between the two peaks is given by Thüring, Lück and Danzmann [46]. If we then assume that $T_{\text{itm}} \ll 1$, Martynov *et al.* [47] provides formulas for the frequency difference between the two peaks ω_s , and the bandwidth of each peak γ_s :

$$\omega_s = \frac{c\sqrt{T_{\text{itm}}}}{2\sqrt{L_{\text{src}}L_{\text{arm}}}}, \quad \gamma_s = \frac{cT_{\text{srcm}}}{4L_{\text{src}}}, \quad (2.2)$$

where T_{srcm} is the power transmissivity of the SRM.

For sample ET-HF parameters of $T_{\text{srcm}} = 0.1$, $T_{\text{itm}} = 0.007$, $L_{\text{src}} = 100$ m & $L_{\text{arm}} = 10$ km, we obtain $\gamma_s = 12$ kHz, $\omega_s = 2$ kHz. As the bandwidth of the peaks is much greater than the separation, they are indistinguishable, and the whole system has a single resonance peak per FSR. However, as L_{src} increases, γ_s decreases faster than ω_s , and the two peaks become more resolved. For the simplified setup in Figure 2.5, the effects of increasing L_{src} are shown in Figure 2.6.

It should be noted that Equations (2.1) and (2.2) are only correct in their relevant extremes. In reality, interferometers operate somewhere between the two, where although calculating the response of the setup in Figure 2.5 is simple, there is no




	L_{src} [m]	γ_s [Hz]	$\omega_s/2\pi$ [Hz]
—	1	1.2×10^6	2×10^4
—	100	1.2×10^4	2×10^3
—	500	2.4×10^3	893
—	1000	1.2×10^3	631
—	5000	240	282
—	10,000	120	200

Figure 2.6: Effect of SRC length on the strain-output transfer function for the simplified ET-HF setup shown in Figure 2.5, normalised to $L_{\text{src}} = 1$ m at 0 Hz. The absolute value of the y-axis is unimportant, and is therefore in arbitrary units (au). As L_{src} increases, ω_s grows in relation to γ_s until they are comparable in size and the split resonance becomes visible.

analytical solution for the bandwidth of the one or two peaks present. However, Equations (2.1) and (2.2) are useful as a starting point for investigating the behaviour of a long SRC with numerical simulations.

It should now be clear why we see a decrease in sensitivity at high frequency for longer SRC lengths in Figure 2.4, as increasing L_{src} reduces the bandwidth of the coupled cavity resonance and decreases the magnitude of the frequency-splitting. To combat this, we can restore ω_s & γ_s to their original values, or as close as possible. For a change in SRC length from L_{src} to L'_{src} , we should therefore increase T_{srin} & T_{itm} by the same ratio $L'_{\text{src}}/L_{\text{src}}$. By increasing T_{itm} , however, we change both the finesse of the arm cavities, and the gain of the power recycling cavity (PRC). This reduces the circulating arm power, and also redistributes power in the interferometer from the arm cavities to the PRC. If we then increase input power to restore the arm cavity circulating power, we can recover the original quantum noise sensitivity curve



	L_{src} [m]	P_{BS} [kW]	P_{in} [W]	T_{itm} [%]	T_{srm} [%]
—	30	10	504	0.70	10.0
—	100	25	614	1.68	21.9
—	300	71	1141	4.62	50.2
—	1000	196	2661	12.25	86.8

Figure 2.7: Options for correcting for an increased SRC length in ET-HF. Parameters were found by numerically optimising to give the minimum beamsplitter power possible, without decreasing sensitivity relative to ET-D. Compared to the uncorrected curves in Figure 2.4, we no longer see a decrease in sensitivity at high frequency with increasing L_{src} , at the cost of increased power on the central beamsplitter—the minimum power required to keep below the original curve scales slightly less than linearly with SRC length. It is noteworthy that SRM transmissivity places a physical limit on how much γ_s can be increased for a given value of L_{src} . For $L_{\text{src}} > 1000$ m, this would start to constrain the design more strictly. It is also notable that the scaling relations from Equation (2.2) do not hold over the whole range of lengths shown; for $L_{\text{src}} = 1000$ m, the required mirror transmissivities are much lower than one would expect.

with a larger L_{src} , at the cost of increased power incident on the central beamsplitter and transmitted through the ITMs. This is undesirable as absorbed laser power causes thermal distortion of the optics, creating a thermal lens which can lead to mode mismatches and losses. Compensation for different values of L_{src} , along with power incident on the beamsplitter, is shown in Figure 2.7. A good compromise for ET-HF, including the beam expansion telescope, can be achieved with L_{src} of around 100 m. Figure 2.8 shows how the quantum noise at high frequencies scales with power incident on the central beamsplitter for this length.

So far we have only discussed the effect of increasing L_{src} on ET-HF. This is because, for any practical value of $L_{\text{src}} \lesssim 1$ km, the effect on the frequency range of




Figure 2.8: Scaling of quantum-noise limited sensitivity with power incident on the central beamsplitter for a fixed SRC length of 100 m in ET-HF. For future high-power detector designs with significant SRC lengths, we can consider the trade-off between slightly reduced sensitivity due to quantum noise at high frequency, and increased noise due to thermal effects when compensating for SRC length.

interest for ET-LF (up to ≈ 30 Hz) is negligible; this is shown explicitly in Figure 2.9. Figure 2.6 provides the explanation for this behaviour. For a 1 km SRC, we have $\gamma_s = 1.2$ kHz, $\omega_s = 631$ Hz—the split resonances are still too wide to be individually resolved, and the splitting frequency is much greater than the top end of the frequency range of interest.

2.3 Optimised Filter Cavities for ET

In order to produce frequency-dependent squeezing to improve the quantum-noise limited sensitivity of ET, frequency-independent squeezed light is reflected from one or more filter cavities. This induces a frequency-dependent phase shift in the reflected light. The ET Design Study [35] considered 10 km long filter cavities. Shorter filter cavities are under consideration as a cost saving change to the design, as the vacuum and tunnel infrastructure are one of the main costs of the future observatory. Significantly shortened filter cavities would allow a simplification of the infrastructure design—example lengths being considered are a reduction from 10 km down to 1 km for ET-LF, and 300 m for ET-HF. The performance of filter cavities is determined by their loss per unit length [21], and this reduction in length




Figure 2.9: Effect of increasing L_{src} on ET-LF, without changing any other interferometer parameters. We see that even for $L_{\text{src}} = 1$ km, the reduction in sensitivity is only 2% at 30 Hz, and for a realistic length of 100 m, the effect is on the order of 0.1%.

will lead to a corresponding increase in squeezing loss in the filter cavities. In practice, the optical loss will be determined by the detailed properties of the optical surface and the beam radius [48], the minimum value of which is dependent on cavity length. A simple extrapolation from other experiments would suggest the following optical losses are achievable with current technology and techniques for the different filter cavity lengths [40]: 30 ppm @ 300 m, 40 ppm @ 1 km and 75 ppm @ 10 km. Throughout this section, we do not attempt to predict detailed optical losses, but provide quantum-noise limited sensitivity curves for a range of possible round-trip filter cavity power losses.

2.3.1 ET-HF

For ET-HF in its tuned, broadband configuration, only one filter cavity is required. In this case, an analytical solution for the optimal filter cavity detuning and bandwidth is given in [49, Equations (31, 33, 49, 50 & 53)] as

$$\Delta\omega_{\text{fc}} = \sqrt{1 - \epsilon}\gamma_{\text{fc}}, \quad (2.3)$$

$$\gamma_{\text{fc}} = \sqrt{\frac{2}{(2 - \epsilon)\sqrt{1 - \epsilon}}} \frac{\Omega_{\text{SQL}}}{2}, \quad (2.4)$$

$$\text{where } \epsilon = \frac{4}{2 + \sqrt{2 + 2\sqrt{1 + \left(\frac{2\Omega_{\text{SQL}}}{f_{\text{FSR}}\Lambda_{\text{rt}}^2}\right)^4}}} \quad (2.5)$$

$$\text{and } \Omega_{\text{SQL}} \simeq \frac{t_{\text{sr}}}{1 + r_{\text{sr}}} \frac{8}{c} \sqrt{\frac{P_{\text{arm}}\omega_0}{mT_{\text{itm}}}}. \quad (2.6)$$

Here, f_{FSR} is the free spectral range of the filter cavity, Λ_{rt}^2 is the round-trip power loss in the filter cavity, t_{sr} & r_{sr} are the amplitude transmissivity and reflectivity of the signal recycling mirror, P_{arm} is the circulating power in the arm cavities, ω_0 is the carrier frequency, m is the mass of the test masses, and T_{itm} is the power transmissivity of the input test mass mirror.




Figure 2.10: The effect of filter cavity loss on ET-HF, with $L_{\text{fc}} = 300$ m. For a round-trip loss of 70 ppm, we see an increase in quantum noise of $\simeq 75\%$ at 30 Hz, relative to that of ET-D. However, ET-D is limited by thermal noise around this frequency, so the overall sensitivity is not affected. As filter cavity performance is determined by loss per unit length, the lower loss curves can be used to infer curves for other lengths e.g. a 700 m filter cavity with 70 ppm round-trip loss would give the 30 ppm curve shown.

L_{fc} [m]	Tuning [Hz]	Half-bandwidth [Hz]
300	-29.9520	5.2305

Table 2.1: Optimal filter cavity parameters for ET-HF, with 70 ppm round-trip filter cavity loss.

Filter	Tuning [Hz]	Half-bandwidth [Hz]
FC ₁	25.3574	5.6830
FC ₂	-6.6366	1.4468

Table 2.2: Optimal filter cavity parameters for ET-LF, with 70 ppm round-trip filter cavity loss.

Figure 2.10 shows the effect of loss on ET-HF, with filter cavity length $L_{\text{fc}} = 300$ m, and Table 2.1 gives the optimal filter cavity parameters according to Equations (2.3) and (2.4). The effect of the increased losses is a reduction in the quantum-noise limited sensitivity at low frequencies, especially around 30 Hz. However, in this frequency band the current ET-HF design is entirely limited by thermal noise, so the overall sensitivity is not affected. We expect that upgrades to the interferometers in the long-term infrastructure of ET will reduce this thermal noise. Space for a filter cavity of modest length should therefore already be allocated in the initial infrastructure.

2.3.2 ET-LF

ET-LF operates with a detuned SRM, and thus requires two filter cavities to achieve optimal squeezing. Unlike in the single filter cavity case, no analytical solution exists for multiple lossy filter cavities. A good approximation is provided by [50, Appendix A], however this assumes lossless filter cavities. Thus, we start with this approximation, and then optimise numerically. Figure 2.11 shows the effects of losses on ET-LF, with $L_{\text{fc}} = 1$ km, and Table 2.2 gives optimal filter cavity parameters. We see that a reduction in length leads to a much more stringent requirement for the optical loss to avoid spoiling the sensitivity at low frequencies around 7 Hz.

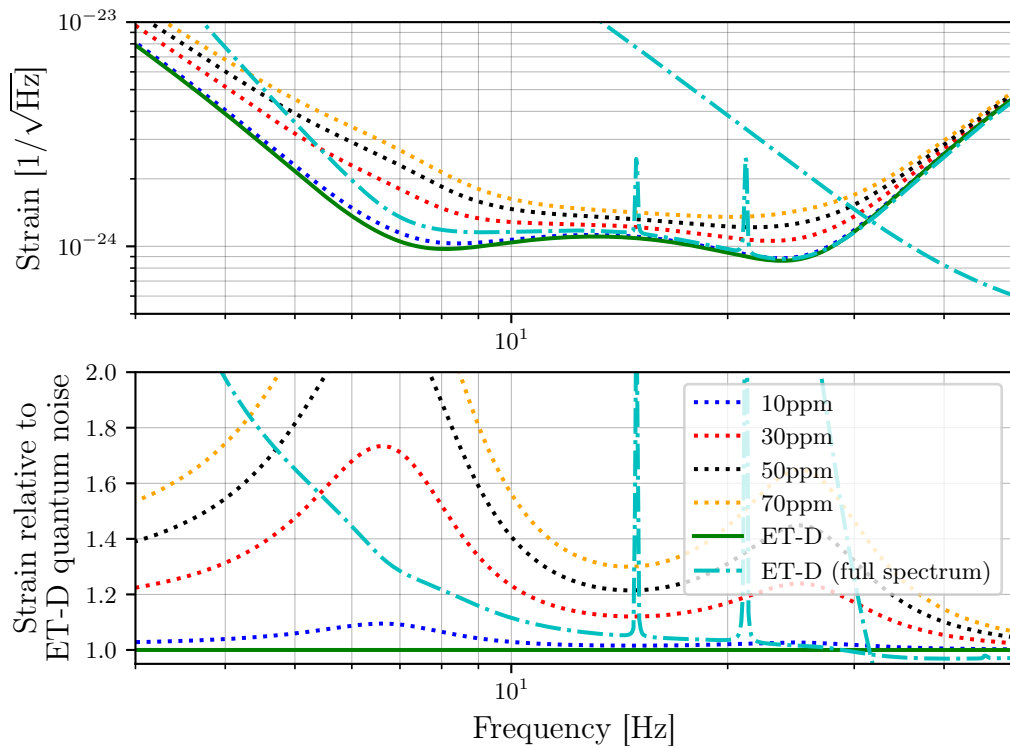


Figure 2.11: The effect of filter cavity loss on ET-LF, with $L_{fc} = 1$ km. For a round-trip loss of 70 ppm, we see an increase in quantum noise of up to 160 % at 7 Hz, relative to that of ET-D (75 ppm loss, $L_{fc} = 10$ km), and a minimum of 30 % over the full spectrum below 30 Hz.




Figure 2.12: Comparison of a detuned ET-LF with 2×1 km filter cavities (solid curves) and a tuned ET-LF with 1×10 km filter cavity (dotted curves). For a round-trip power loss > 30 ppm, the tuned system with one filter cavity performs better across the entire frequency range, especially at high frequencies.

The Virgo detector has already demonstrated round-trip losses of 55 ± 10 ppm in km-scale cavities [51]. These optical losses are dominated by deficiencies in the mirror surface quality, and research is ongoing to identify and mitigate the loss due to light scattering by surface defects. Using the same technology it should currently be possible to realise a 1 km long cavity with round-trip losses of less than 40 ppm. It is reasonable to believe that in the future we can improve the mirror surface quality further, to achieve a round trip loss of 20 ppm, with careful use of state-of-the-art technologies and care regarding polishing, coating, handling and installation.

The main motivation for reducing the filter cavity lengths in ET-LF is the cost of the infrastructure. There is a disincentive to use the main 10 km tunnels for the arm cavities as well as two filter cavities, due to the scaling of excavation cost with tunnel diameter. Given the lack of a low-frequency dip for a 1 km filter cavity with loss $\gtrsim 20$ ppm, it is interesting to compare the previous results to the performance of a tuned SRM for ET-LF with only one filter cavity. A comparison of this tuned ET-LF with a single 10 km long filter cavity, and a detuned ET-LF with two 1 km long filter cavities, is shown in Figure 2.12. For a round-trip power loss > 30 ppm, the tuned system with one filter cavity performs better than the detuned system with two short filter cavities.

Type	Transmissivity	Tuning [degrees]
Two cavities	$FC_{1,\text{in}} = 4.617 \times 10^{-3}$	$FC_{1,\text{end}} = 3.049 \times 10^{-1}$
	$FC_{2,\text{in}} = 1.210 \times 10^{-3}$	$FC_{2,\text{end}} = -7.971 \times 10^{-2}$
Coupled cavity	$FC_{\text{in}} = 5.856 \times 10^{-3}$	$FC_{\text{mid}} = 2.276 \times 10^{-1}$
	$FC_{\text{mid}} = 3.099 \times 10^{-5}$	$FC_{\text{end}} = 2.256 \times 10^{-1}$

Table 2.3: Optimal filter cavity parameters for ET-LF, for both two filter cavities and a single coupled filter cavity. Each individual cavity is 10 km long, with 37.5 ppm loss per optic. Note the low transmissivity required for the middle mirror in the coupled filter cavity. As cavity length decreases, so too does the required value of middle mirror transmissivity.

The design for the ET-LF filter cavity scheme is more complex than for ET-HF, and has to include a careful trade-off between excavation cost, expected optical losses and practical constraints for arranging the vacuum system.

2.3.3 Coupled Filter Cavities

The purpose of using filter cavities in the squeezing path is to replicate the quadrature rotation of the interferometer as seen by the signal light. For a detuned signal-recycled Michelson such as ET-LF, two separate rotations are required. Current plans for ET-LF achieve the desired rotation with two separate filter cavities in series, with each filter cavity producing a single rotation around its resonance. As a coupled cavity exhibits two separate resonances, these two independent filter cavities could potentially be replaced with a coupled filter cavity. To investigate this possibility, a model of a coupled filter cavity was numerically fit to give the same squeezing angle rotation as the two filter cavities, and then further optimised to maximise the quantum-noise limited sensitivity from 5–30 Hz.

Figure 2.13 compares the quantum-noise limited sensitivity of ET-LF for a 20 km coupled filter cavity vs 2×10 km filter cavities, both with and without losses. Optimal parameters are given in Table 2.3. There are a few noteworthy points here. Firstly, from Figure 2.13, we can see that the fit was performed successfully, and as such a coupled filter cavity could in theory be used in place of two independent

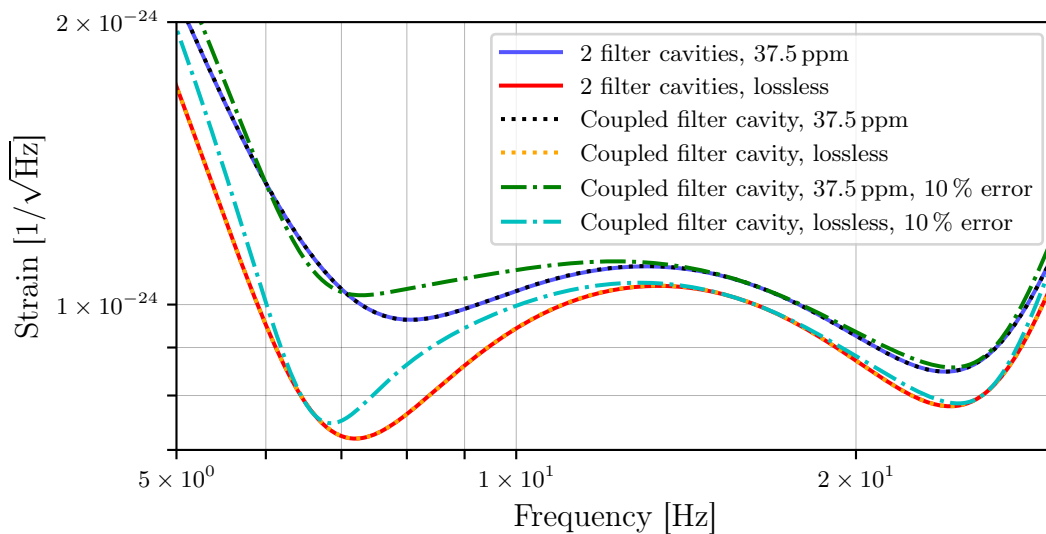


Figure 2.13: Comparison of using two separate 10 km filter cavities vs. a coupled filter cavity totalling 20 km in ET-LF, both with lossless mirrors and 37.5 ppm loss per surface. The parameter with the lowest tolerance is the transmissivity of the middle mirror; the dash-dotted lines show the coupled cavity with a 10% increase in the middle mirror transmissivity from the optimal value, and the other parameters adjusted to compensate. A coupled filter cavity can be used to replicate the effects of two independent filter cavities in ET-LF. Perhaps surprisingly, the performance in the presence of losses is identical for the same loss per surface.

cavities in a detuned, dual-recycled Michelson such as ET-LF. Additionally, we see that the scaling of performance with mirror losses is identical in the coupled and independent cases. We also see that the performance of the filter cavity at low frequencies (~ 7 – 12 Hz) is fairly sensitive to the middle mirror transmissivity.

There are a few motivating factors that make the coupled filter cavity design worth further study. When two individual cavities are used, some extra optics such as Faraday isolators must be introduced to direct the beam from one cavity to the next; this is not needed in a coupled filter cavity. Without these extra Faraday isolators the overall optical loss in the input squeezing path can be reduced [31], to increase the effective squeezing level achievable. Additionally, the same considerations for using a tuned vs. detuned Michelson for ET-LF from Section 2.3.2 apply here: in the case of the coupled cavity scheme, the total filter cavity length is arranged sequentially in one long vacuum system, whereas the scheme with two filter cavities requires a shorter but wider space for two parallel vacuum systems. It should be noted that a




Figure 2.14: Comparison of a detuned ET-LF with a 10 km total length coupled filter cavity (solid curves) and a tuned ET-LF with a single 10 km filter cavity (dotted curves). For all values of per-surface power loss plotted, the coupled filter cavity leads to better average sensitivity from ~ 4 –30 Hz, and a loss of 20 ppm brings the sensitivity in line with that of ET-D.

coupled filter cavity of total length 10 km could provide much better sensitivity than a tuned detector in each of ET-LF’s ‘dips’. This is shown explicitly in Figure 2.14.

In summary, a coupled filter cavity could be used in place of two independent filter cavities. There are two main advantages to this substitution: the lack of a need for an extra Faraday isolator results in lower losses in the squeezing path, and in the case of ET the form factor of the vacuum system could be advantageous. These advantages provide motivation for further study.

2.4 Summary

Practical considerations may motivate the introduction of longer signal recycling cavities in future detectors, such as ET. If the length of the SRC is not accounted for, it can lead to an overestimation of the sensitivity at high frequencies, as shown in Figure 2.4. This can be avoided by considering the SRC-arm system as a coupled cavity, the response of which is described loosely by a split resonance, with separation frequency ω_s and half-bandwidth γ_s , as given in Equation (2.2). We have shown that the change in the response of a detector due to a longer SRC can be counteracted by increasing both T_{itm} & T_{srm} . To maintain arm cavity power and thus sensitivity,

we must also increase input power, either directly or by increasing the PRC gain. Increasing the length of the SRC, while maintaining sensitivity at high frequencies, therefore leads to an increase in the power incident on the central beamsplitter, as shown in Figure 2.7.

In the specific case of ET-HF, we suggest an SRC length of 100 m, which would result in a reduction of the quantum-noise limited sensitivity of only 25 % at 10 kHz compared to an SRC of negligible length. We further show how this loss of sensitivity can be compensated for, at the cost of increased laser power at the beamsplitter. The scaling of this sensitivity with beamsplitter power can be seen in Figure 2.8. For ET-LF, the frequencies at which the decrease in sensitivity occurs are too high to be of consequence. Figure 2.9 shows that for a 100 m SRC, the reduction in sensitivity at 30 Hz is on the order of 0.1 %.

Constraints such as the size and cost of the underground infrastructure provide motivation for reducing the length of the filter cavities used for frequency-dependent squeezing in ET. This has the effect of increasing filter cavity loss per unit length, leading to a reduction in performance. For ET-HF, a much shorter filter cavity of, for example, 300 m has limited consequences, as thermal noise of the main interferometer remains the limiting factor in the frequency range affected by filter cavity losses. In ET-LF, the consequences of reducing the filter cavity lengths from 10 km to 1 km would be more severe, giving up to a factor of 2 reduction in quantum-noise limited sensitivity at 7 Hz for a currently achievable round-trip optical power loss of 40 ppm. However, with expected improvements in optical losses such as an increase in mirror surface quality, a significant reduction of the filter cavity length would be possible. In addition, the use of a coupled filter cavity in place of two independent filter cavities for ET-LF was investigated, and found to perform identically for the same length and per-surface loss. Combined with the fact that coupled cavities would use one less Faraday isolator in the injection path, which further reduces the optical losses, coupled filter cavities should be studied further.

Chapter 3

Output path filter cavities

This chapter is a modified version of my paper ‘Studies of input and output filter cavities for Einstein Telescope’, which is currently awaiting submission for publication.

In the previous chapter, we investigated the properties of filter cavities for frequency dependent squeezing in the context of Einstein Telescope (ET). In that case, only input squeezing was considered, with the filter cavities located in the squeezed light input path. In this chapter, I will discuss other possible arrangements of filter cavities for frequency dependent squeezing, including those with a filter cavity in the interferometer output path. This is again in the context of ET, specifically the low frequency interferometer ET-LF. We will see the potential improvements to sensitivity that could be made for ET and other detectors, and where some problems with these alternate designs lie. Finally, for schemes with two filter cavities, I will present the case for using different lengths for each cavity. This may be desirable when practical factors limit the length of filter cavities.

3.1 Frequency dependent squeezing

Frequency-dependent squeezing is achieved via the use of one or more filter cavities. In their paper, Kimble *et al.* [21] discuss two designs that make use of frequency-dependent squeezing:

1. *Squeezed-Input (input filtering)*: In this scheme, frequency-independent squeezed light is first reflected off one or more filter cavities, before injection into the interferometer at the dark port.
2. *Squeezed-Variational (output filtering)*: Here, frequency-independent squeezed light is directly injected into the interferometer at the dark port, and then the combined output containing the signal and squeezed light is reflected off the filter cavities. A *Variational-Output Interferometer* has the same design, but without squeezing injection.

The motivation for the variational-output setup can be seen by considering what causes increasing radiation pressure noise at low frequencies in an interferometer. Briefly put, the suspended mirrors in the arm cavities act to couple the amplitude and phase quadratures of light, and as such the interferometer squeezes any input noise. The squeezing angle, and thus the optimal readout angle, changes with sideband frequency. It turns out that if the optimal readout angle can be used for all sideband frequencies, radiation pressure noise can be avoided entirely, resulting in a purely shot noise limited interferometer. This can be achieved via the use of filter cavities, similarly to the frequency dependent input squeezing we have already seen. For more detail, see Sections IV.C and V of [21].

Example layouts of the squeezing and output optics for both squeezed-input and squeezed-variational designs are shown in Figure 3.1. Most designs for future detectors are of the squeezed-input variety [10, 23, 52, 53], as this presents little change to the interferometer layout over a design without squeezing. It is therefore




Figure 3.1: Example squeezing and output optic layouts of (a) *Squeezed-Input* and (b) *Squeezed-Variational* interferometers. For a current interferometer using frequency-independent squeezing and DC readout, the changes necessary to add frequency-dependent input squeezing are fairly self-contained. Output filtering requires more significant alterations; DC readout must be exchanged for another scheme such as homodyne detection, as pictured here, and either the geometry of the output must change or extra steering optics have to be introduced.

the ideal candidate for upgrades to current detectors, and investigations into such upgrades are currently underway [23]. The squeezed-variational interferometer has seen little practical consideration by comparison, despite having the potential to perform much better at low frequencies. There are multiple reasons for this. First, it requires a more substantial change to the interferometer design than adding input squeezing, making it unsuitable for upgrades to current detectors. For example, DC readout as currently used cannot work in a variational-readout interferometer, as it fixes the relative phase of the local oscillator to the signal field, which the variational-readout design seeks to make frequency-dependent. Secondly, optical losses in the output path can severely limit the performance of the squeezed-variational case, and we will investigate this further in Section 3.4.

For future detectors focused on low frequencies, such as the yet to be constructed ET-LF, balanced homodyne readout will be implemented. Without the constraints of existing infrastructure, and with future improvements of suspension thermal noise and environmental noise from both the application of cryogenic technology and a shift to an underground location, the study of output-filtered interferometer designs may become more interesting. The lower losses expected to be achievable in future filter cavities, and our stronger understanding of interferometry in general, make it the right time to revisit these designs.

3.2 Motivation

In their review, Danilishin & Khalili [54] provide a good overview of the behaviour of the two designs mentioned in Section 3.1. A combination of the two designs is also discussed; double-filtering with filter cavities in both the squeezed light input path and the interferometer output path. In all cases the optimal quantum noise

power spectral density is derived, giving

$$S_{\text{in}}^h(\Omega) = \frac{2\hbar}{mL^2\Omega^2} \left[\frac{e^{-2r} + \epsilon_d^2}{\mathcal{K}(\Omega)} + e^{-2r}\mathcal{K}(\Omega) \right] \quad (3.1)$$

for the squeezed-input interferometer,

$$S_{\text{out}}^h(\Omega) = \frac{2\hbar}{mL^2\Omega^2} \left[\frac{e^{-2r} + \epsilon_d^2}{\mathcal{K}(\Omega)} + \frac{\epsilon_d^2}{1 + \epsilon_d^2 e^{-2r}} \mathcal{K}(\Omega) \right] \quad (3.2)$$

in the squeezed-variational case, and

$$S_{\text{I/O}}^h(\Omega) = \frac{2\hbar}{mL^2\Omega^2} \left[\frac{e^{-2r} + \epsilon_d^2}{\mathcal{K}(\Omega)} + \frac{\epsilon_d^2}{1 + \epsilon_d^2 e^{2r}} \mathcal{K}(\Omega) \right] \quad (3.3)$$

for combined input-output filtering, where m is the arm cavity mirror mass, L the cavity length, Ω the gravitational wave signal frequency, r the squeezing factor, and ϵ_d^2 the photodetector quantum inefficiency. The term $\mathcal{K}(\Omega)$ is an optomechanical coupling factor introduced in [21], defined as

$$\mathcal{K}(\Omega) = \frac{2(I_0/I_{\text{SQL}})\gamma^4}{\Omega^2(\gamma^2 + \Omega^2)}, \quad (3.4)$$

where I_0 is the circulating arm cavity power, γ the interferometer bandwidth, and

$$I_{\text{SQL}} = \frac{mL^2\gamma^4}{4\omega_0} \quad (3.5)$$

the power required to reach the SQL in a standard interferometer with laser frequency ω_0 .

In Equations (3.1) to (3.3), the first term in brackets is the shot noise of the detector, and the second term is the noise due to back action or radiation pressure. We see that in all three cases, the ideal shot noise performance is identical, and it is only

the radiation pressure term that differs. An observation of note is that for both cases with output filtering, the radiation pressure noise is completely eliminated in a lossless interferometer. In the lossy interferometer case, however, the effects of squeezing differ. Radiation pressure noise is enhanced by squeezing in the squeezed-variational design, whereas it is suppressed for input-output filtering.

It may seem from Equations (3.1) to (3.3) that combined input-output filtering will always perform better than either of the individual filtering schemes; the shot noise performance is identical, and taking the ratio of the radiation pressure terms yields

$$\frac{S_{\text{inRP}}^h}{S_{\text{I/O,RP}}^h} = \frac{e^{-2r}}{\epsilon_d^2/(1 + \epsilon_d^2 e^{2r})} = 1 + \frac{e^{-2r}}{\epsilon_d^2} > 1 \quad (3.6)$$

$$\frac{S_{\text{outRP}}^h}{S_{\text{I/O,RP}}^h} = \frac{\epsilon_d^2/(1 + \epsilon_d^2 e^{-2r})}{\epsilon_d^2/(1 + \epsilon_d^2 e^{2r})} = \frac{1 + \epsilon_d^2 e^{2r}}{1 + \epsilon_d^2 e^{-2r}} \geq 1 \quad (3.7)$$

for $r \geq 0$, i.e. the noise due to radiation pressure will always be lower in the input-output filtering case. However, these equations neglect to account for losses in the filter cavities themselves, which will be non-negligible for future gravitational wave detectors. By including two filter cavities rather than one, the combined input-output filter will necessarily have greater losses. While [54] further discusses the effect of filter cavity losses in the input and output filtering schemes individually, no comparison to the combined input-output filtering design is made. In Section 3.4, we investigate the application of the combined filtering design—including filter cavity losses—to the specific case of ET-LF, and compare the relative performance to the currently proposed detuned SRM setup.




Figure 3.2: Comparison of detuned and tuned signal recycling for ET-LF. Both cases use 1 km filter cavities with 20 ppm round-trip power losses, and 18 kW of power in the arms. The two dips where a detuned SRM improves sensitivity can be seen around 8 Hz & 25 Hz, along with the reduced bandwidth.

3.3 Detuned ET-LF

The currently proposed design for ET-LF consists of a dual-recycled Fabry-Perot Michelson interferometer with 10 km arm cavities, a detuned SRM, and 2×1 km input filter cavities providing frequency-dependent squeezing [10]. The target frequency range is low, at 3–100 Hz.

The motivation for this setup is discussed in the ET Design Report [35]. By detuning the SRM, the quantum-noise limited sensitivity of the interferometer is increased around two frequencies; the effective resonance frequency of the signal recycling cavity, and the resonance frequency of the resulting optical spring in the arm cavities. This gives rise to the double-peaked bucket shape seen in Figure 3.2.

There are two main drawbacks to the use of a detuned signal recycling cavity in ET-LF. Firstly, as can be seen in Figure 3.2, the bandwidth of the interferometer is reduced, leading to a large reduction of sensitivity above 30 Hz compared with the tuned case. Due to ET’s xylophone design, this problem is mostly irrelevant as ET-HF reaches comparable sensitivity around this frequency, as long as the

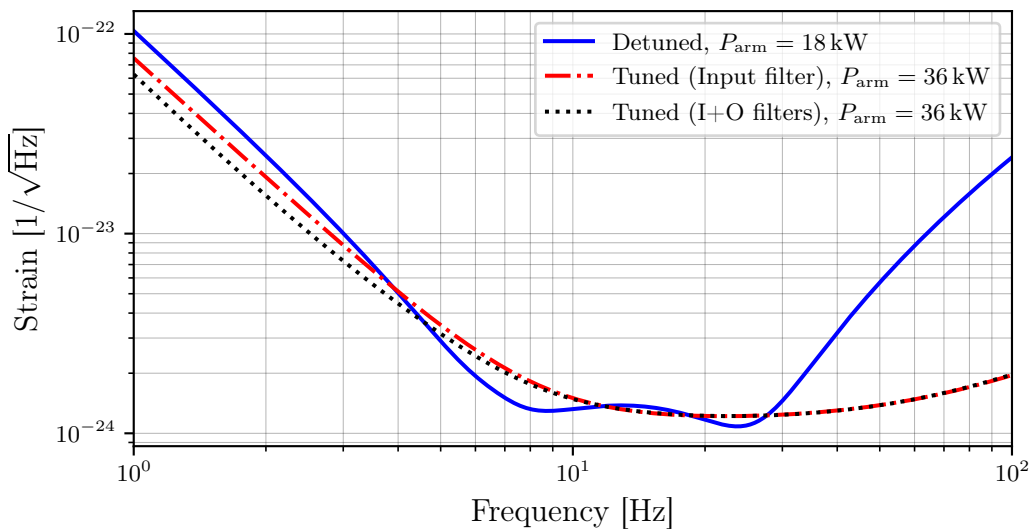


Figure 3.3: Comparison of the combined input-output filtering design for ET-LF with both the current design and tuned input filtering. All cases use 1 km filter cavities with 20 ppm round-trip power losses. While including a factor 2 increase in arm power, the tuned cases perform well across most of the frequency range, especially at high and low frequencies. The combined input-output filter shows a small improvement over the input filtering case at low frequencies less than ~ 7 Hz.

design sensitivity can be achieved. The second issue with detuned signal recycling is a practical one; both GEO600 and LIGO have seen that inclusion of a detuned SRM makes lock acquisition and control of the interferometer significantly more complex [55, 56]. This provides motivation to find an alternative to detuned signal recycling for ET-LF, if a scheme can be found that provides comparable sensitivity within the facility limits.

3.4 Tuned ET-LF with combined I/O filter cavities

Figure 3.3 shows a comparison of two possible tuned SRM designs discussed in Section 3.2 with the currently proposed detuned SRM setup, taking into account filter cavity losses. A photodetector quantum inefficiency of $\epsilon_d^2 = 0.05$ is assumed throughout this paper. Filter cavity parameters were found via numerical optimisation. We see that the combined input-output filtering scheme provides sensitivity comparable with the detuned input-squeezing case, with slightly worse sensitivity in the dips at ~ 4 – 10 Hz & ~ 20 – 28 Hz, but better sensitivity below 4 Hz & above 20 Hz.




Figure 3.4: Sensitivity scaling with arm cavity power for the combined input-output filtering design. T_{srm} is adjusted in each case to keep the radiation-pressure limited performance the same.

One of the potentially limiting factors of the tuned SRM designs is the increased arm power required to match the design sensitivity around 20 Hz. ET-LF makes use of cryogenic test masses to reduce coating thermal noise, placing an upper bound on the amount of arm cavity power allowable while still achieving the desired temperature. A factor two increase to 36 kW is almost sufficient to match the detuned case in the high frequency dip at 25 Hz, but is not enough to match the sensitivity around 7 Hz. Figure 3.4 shows the scaling of sensitivity with arm cavity power in the input-output filtering case, varying the SRM transmissivity to maintain radiation-pressure limited performance.

Conventional wisdom from modern interferometers might suggest that this increase in power is not necessary. The tuned case already out-performs the detuned case at high and low frequencies, so perhaps we could reduce the bandwidth and increase the peak sensitivity of the interferometer by changing the SRM transmissivity. Unfortunately, due to the loss-limited nature of the combined input-output squeezed interferometer, this is not possible. As Figure 3.5 demonstrates, attempting to do so only reduces the bandwidth of the interferometer without increasing peak sensitivity.




Figure 3.5: Effects of varying the bandwidth of the combined input-output filtering design. Arm power is held at 18 kW for all curves. We see that increasing T_{srm} past ~ 0.6 provides very little improvement in peak sensitivity, while still reducing the bandwidth.

For the purposes of detector design, it is useful to see the contribution of each of the losses in the interferometer to the overall quantum noise curve. This is shown in Figure 3.6, from which we can see that filter cavity losses are the limiting factor below ~ 14 Hz. Note that sensitivity is compared—rather than just quantum noise—as losses in the output filter also affect the signal light reaching the detector. Unfortunately, we cannot cleanly separate the effects of losses in the two filter cavities, as they are strongly linked. For example, if there are no losses in the output filter, the input filter provides no benefit anyway.

3.5 Dissimilar filter cavities

Both the original [35] and updated [10] ET design reports propose the use of two filter cavities of equal length for ET-LF. Originally, 10 km cavities in the same tunnels as the arms were suggested, which were then reduced to 1 km in separate auxiliary tunnels as a cost-saving measure. Using the same length for both filter cavities may not be the optimal solution, however. One filter cavity is used to compensate the signal recycling detuning, while the other compensates the radiation pressure induced rotation from the interferometer. As these effects occur at different frequencies, the




Figure 3.6: Effects of different loss sources in the combined input-output filtering design with 18 kW arm power and 10 dB squeezing. The readout loss curve represents a 5% photodiode quantum inefficiency, and the filter cavity loss curve uses 1 km filter cavities with 20 ppm round-trip power losses. We see that below about 14 Hz, filter cavity losses are the limiting factor in the input-output design.

performance of each filter cavity affects different parts of the sensitivity curve. It may therefore be desirable to increase the length of one filter cavity more than the other. In the case of ET-LF, if one filter cavity can fit in the arm tunnels, then perhaps a setup with $1 \times 1 \text{ km} + 1 \times 10 \text{ km}$ filter cavities could be used.

A different argument with the same outcome can be applied to the input-output filtering design. As the output filter is much more sensitive to losses than the input filter, the longest possible output filter should be used if a choice has to be made.

Figure 3.7 shows the effects of increasing the length of just one of the filter cavities in ET-LF to 10 km, for both the detuned and input-output filtering cases. The round-trip power loss in this cavity is also increased to 40 ppm, as the loss depends on the length [48], and extrapolation from current experiments suggests roughly double the loss for a factor 10 increase in length [40]. As expected, we see that the sensitivity of the detuned configuration is increased only in part of the frequency range shown, with significant improvement around 4–10 Hz. The input-output filtering design shows a more modest improvement, but over all frequencies below ~ 10 Hz.

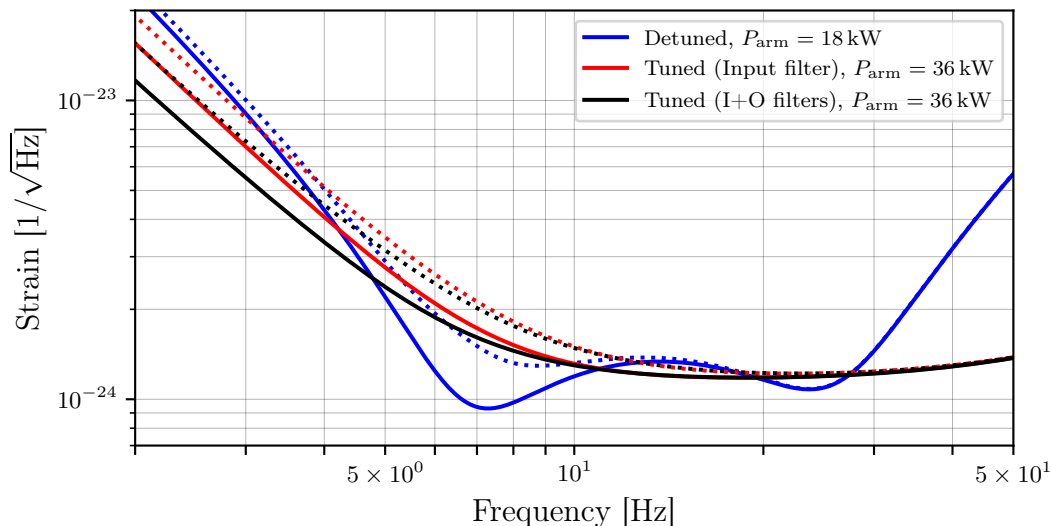


Figure 3.7: Comparison of the detuned, tuned input filtering and combined input-output filter tuned cases with different filter cavities. The dotted curves correspond to 2×1 km filter cavities with 20 ppm round-trip power losses. The solid curves correspond to 1×1 km filter cavity with 20 ppm round-trip loss and 1×10 km filter cavity with 40 ppm round-trip loss. For the tuned input filtering case, only one filter cavity is required. For the detuned case, the filter cavity that gave the largest sensitivity improvement when increased in length was chosen. The tuned cases see an improvement at all frequencies below ~ 20 Hz, with the combined input-output filter benefiting more, while for the detuned case the improvement is limited to ~ 4 –10 Hz.

3.6 Summary

The use of output filter cavities in gravitational wave interferometers has seen relatively little investigation compared to input filters. This is largely due to their degraded performance in the presence of optical losses. Third generation detectors, such as Einstein Telescope, will be able to achieve low enough losses to make designs incorporating output filter cavities feasible. I have shown that a design which combines input and output filtering cavities with realistic losses could be of practical use to future detectors, performing better than the solely input-filtered case for a tuned SRM. These investigations did not take the effects of higher-order spatial modes and backscattering of light from the output into account, however, which could further degrade the performance of the scheme. Additionally, in the specific case of ET-LF, the increased arm cavity power required and the inability to match the performance of a detuned SRM design around 7 Hz make such a design less desirable.

The currently suggested filter cavities for ET-LF are much shorter than originally proposed, at 1 km rather than 10 km. This is so that they can be moved out of the arm tunnels, reducing the main tunnel diameter and thus cost. I have shown that, if a single filter cavity can fit in the arm tunnel, a setup with one long and one short filter cavity provides a significant improvement in performance in both the detuned and input-output filtered tuned cases. Different length filter cavities may be of use in general to future gravitational wave detectors, where the cavity lengths are limited by practical considerations.

Chapter 4

A two-carrier scheme: evading the 3 dB quantum noise penalty of heterodyne readout in gravitational wave detectors

This chapter details simulation work I performed to demonstrate the viability of a completely new readout scheme we proposed in [2]. My contributions were to verify the quantum noise limited sensitivity of the design in the interferometer simulation tool FINESSE, and to derive the requirements on Schnupp asymmetry and carrier separation frequency presented in Section 4.2.1. Figures 4.3, 4.4 and 4.6 are recreations of similar figures published in [2].

An important choice in the design of a modern gravitational wave interferometer is the readout scheme to be used. There are a few major options, each with their own benefits and drawbacks. In this chapter, I will discuss a design that avoids one of the problems with heterodyne readout, which we recently proposed in [2]. This is a completely new scheme, that provides an exciting option for both future detectors

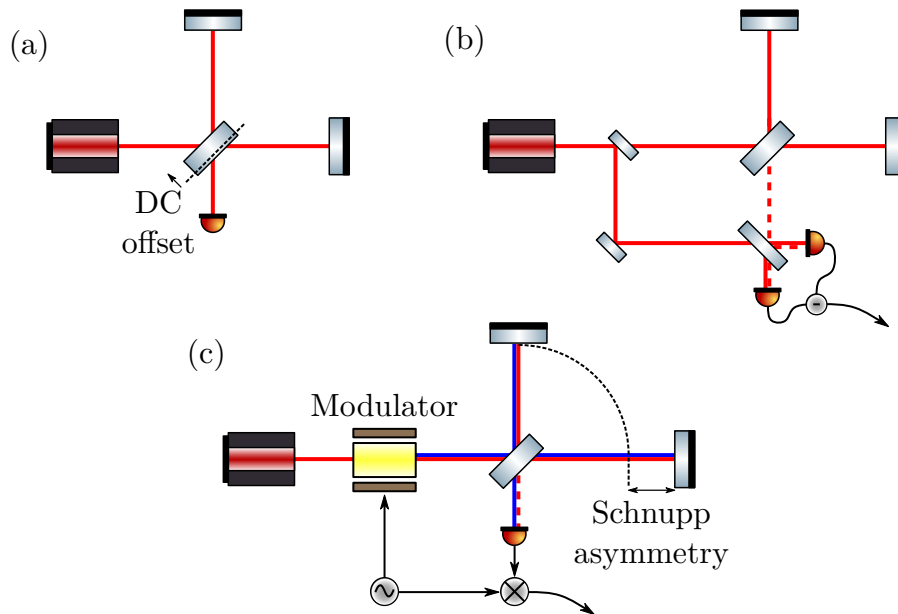


Figure 4.1: Example layouts of three interferometer readout schemes: (a) DC readout, (b) balanced homodyne readout, and (c) heterodyne readout with *frontal* or *Schnupp* modulation. Arm cavities and recycling mirrors have been excluded for compactness.

and upgrades to existing detectors. Starting with a brief overview of existing readout schemes, I will then show the proposed design and the contributions I made by numerically verifying it in the frequency domain modelling tool FINESSE [26].

4.1 Interferometer readout schemes

When measuring the light from a gravitational wave signal in an interferometer, we cannot directly detect the produced sidebands. Instead, we therefore measure the beating of these sidebands with an additional beam called a local oscillator. The beating produces a sinusoidal signal in the power incident on a photodiode, with a frequency equal to the difference between the sideband and local oscillator frequencies. This oscillating power can then be measured. The choice of the local oscillator is the key difference between the major readout schemes available. Example layouts for three readout schemes used in gravitational wave detectors are shown in Figure 4.1.

DC readout is a kind of homodyne readout, meaning that the local oscillator used for readout is at the same frequency as the field from which signal sidebands are generated in the arms. The local oscillator in this case comes from microscopically tuning either the central beamsplitter or one of the arms away from the dark fringe, letting a small amount of carrier light reach the detector. This has the advantage of being simple to implement, but also allows noises present at the interferometer input—such as laser frequency noise—to couple to the output, potentially limiting sensitivity. These noises are discussed in more detail in Chapter 6.

Balanced homodyne is an alternative method, in which the local oscillator light is ‘picked off’ from the main carrier before it enters the interferometer, and guided to the detector via a different path. This removes some of the noise coupling problems present in DC readout, and allows an arbitrary selection of the local oscillator phase, which is required for some designs such as the output path filter cavities discussed in the previous chapter. However, it necessitates the introduction of many more optics than shown in Figure 4.1, and creates another path for noise in the input to couple to the output and limit sensitivity. Additionally, as the local oscillator does not pass through the same optics as the main beam, it must be carefully positioned when they are recombined before detection. If the beams are not aligned, or have different transverse distributions or sizes, there will be excess noise and lost signal [57].

Heterodyne readout is a scheme in which the local oscillator is at a different frequency than the main carrier. This frequency difference is usually at radio frequencies (RF, \sim a few MHz) for the purposes of gravitational wave detectors, far above any signal frequency of interest. In Figure 4.1, the RF field is generated by a modulator in the interferometer input path. This field is then allowed to reach the output by introducing a macroscopic difference in the length of the arms called Schnupp asymmetry. By choosing the length appropriately, the interferometer can be placed on a dark fringe for the carrier field, and a bright fringe for the RF field. The signal is then read out by demodulating the output of a photodiode at the




Figure 4.2: Frequencies present in (a) homodyne readout and (b) heterodyne readout with a single RF field. Ω is the gravitational wave signal frequency, ω_0 the carrier frequency, and ω_m the heterodyne modulation frequency. The trees at the top show the frequencies that will contribute to the output of a demodulated photodiode. Vacuum noise is present at all frequencies. For homodyne readout, the photodiode is demodulated at Ω , and picks up the signal sidebands at $\omega_0 \pm \Omega$, along with the vacuum noise present at those frequencies. For heterodyne readout, the photodiode is demodulated first at ω_m and then Ω . We see that in addition to the signal sidebands and vacuum noise around the carrier frequency ω_0 , extra vacuum noise around $\omega_0 + 2\omega_m$ will also be picked up, reducing sensitivity.

RF modulation frequency. This approach brings a number of advantages. There is no DC offset, reducing the coupling of input noises to the output. Additionally, the demodulation at RF avoids such noises around the local oscillator, as they are typically insignificant at such high frequencies (for instance, laser frequency noise falls off like $1/f$). The drawback is an inherent drop in sensitivity due to the demodulation also picking up vacuum noise around twice the modulation frequency. This is shown in Figure 4.2. The result is a doubling of the quantum noise to signal power ratio, or a factor $\sqrt{2}$ in amplitude, which is known as the 3 dB penalty [58]. This doubling can be seen fairly straightforwardly as a result of measuring the power spectral density (PSD) of uncorrelated noise fields. More information on this is given in Chapter 7, but simply put, as the different noise fields are incoherent, their powers sum in the PSD. For the homodyne case, two noise fields are picked up, whereas in the heterodyne case, four fields are picked up, hence the factor two increase in noise power. In the first generation of gravitational wave detectors, heterodyne readout with two RF fields symmetric around the carrier was used. This can be shown to perform somewhat better, reducing the penalty to a factor 1.5 in power [59].

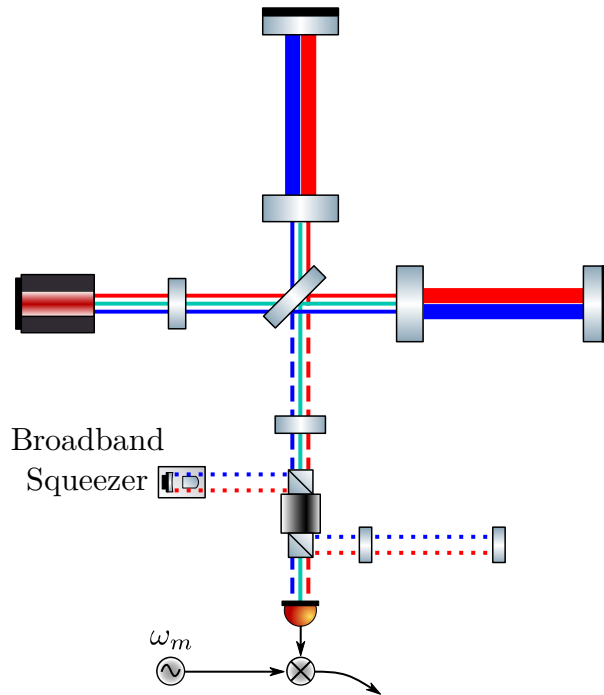


Figure 4.3: Two-carrier scheme proposed in [2]. Both the red field at ω_1 and blue field at ω_2 resonate in the arm cavities and are modulated by an incoming gravitational wave signal, generating signal sidebands. The cyan field at $\omega_L = (\omega_1 + \omega_2)/2$ does not enter the arms, and is used for heterodyne readout of the two signal fields by demodulation at the field separation frequency ω_m . The broadband squeezer generates correlated pairs of sidebands around ω_L , with a bandwidth of at least $\omega_m + \Omega$. The various frequencies present at the output and demodulation scheme are shown in Figure 4.4

4.2 The two-carrier scheme

The question to be answered, then, is how can this 3 dB penalty for heterodyne readout be avoided? The answer is conceptually quite simple; by ensuring that all frequencies that are picked up by a demodulated photodiode contain a coherent signal. This can be achieved by including two carriers in the interferometer, rather than one, leading to the design proposed in [2]. The layout of this two-carrier scheme is shown in Figure 4.3. In this scheme, three fields equally spaced in frequency are injected into the interferometer: two carriers at ω_1 and ω_2 , and a central field used as the local oscillator for heterodyne readout at $\omega_L = (\omega_1 + \omega_2)/2$. As both carriers resonate in the arm cavities, an incoming gravitational wave will generate signal sidebands around each of them. The fields at the output of the interferometer are




Figure 4.4: Frequencies present at the output of the two-carrier scheme. Compared with Figure 4.2, we see that all frequencies included in the black demodulation tree at the top contain signal sidebands. This eliminates the reduction in sensitivity otherwise seen with heterodyne readout. When a broadband squeezer is included, pairs of signal sidebands symmetric around ω_L are entangled, as shown by the grey lines.

then as shown in Figure 4.4. It can be shown that if the same total power in the arm cavities as in other schemes is split equally between each carrier, the 3 dB penalty is evaded [2].

When including squeezing, some additional modifications are also required for the two carrier scheme. In homodyne readout, a squeezer is pumped at $2\omega_0$, generating correlated pairs of sidebands around ω_0 . It has proven difficult to generate squeezed sidebands close to the local oscillator frequency, however, due to effects such as scattering of the local oscillator light into the squeezer [60]. For the two-carrier scheme, [2] shows that we can instead use a broadband squeezer pumped at $2\omega_L$, with a bandwidth covering the range from $\omega_1 - \Omega$ to $\omega_2 + \Omega$. This will then generate squeezed sidebands centered around ω_L , as shown in Figure 4.4. Squeezing at high frequencies avoids the problems caused by scattering of the local oscillator, although any residual light from either carrier present at the output could still scatter. Other benefits of the two-carrier scheme include compatibility with output path filters as described in the previous chapter, which normal heterodyne measurement is incompatible with due to the presence of unsqueezed vacuum noise in the photodiode output.

Due to the benefits discussed above, we propose the two-carrier scheme as a new option for future gravitational wave detectors, or for upgrades to existing detectors

such as aLIGO. This upgrade can be carried out without major changes to the interferometer optics.

In order to verify the feasibility of this new scheme, I performed simulations in FINESSE that successfully demonstrated equivalent quantum noise limited sensitivity to homodyne readout. A plot of this equivalence can be seen later in Figure 4.6. During the course of these simulations, I also discovered strict requirements on the field separation frequency ω_m and the Schnupp asymmetry, which I will detail in Section 4.2.1. The Schnupp asymmetry must be controlled to within a specific fringe (i.e. the mirrors' positions must be correct to within a wavelength of the laser), but this is already achievable at aLIGO.

4.2.1 Resonance conditions

In a Michelson interferometer with only one carrier field, the necessary resonance conditions are fairly straightforward to achieve. The carrier must be resonant in the power recycling cavity (PRC) and the arm cavities, and antiresonant in the signal recycling cavity (SRC). This can be achieved by shifting one of the arm cavities by 90° (i.e. $\lambda/4$) relative to the other. In the two-carrier scheme with multiple carriers and a local oscillator field present, the situation becomes more complex. The solution is perhaps easiest to understand by building up to a full interferometer and considering the necessary phases at each stage.

Let us start with a Michelson interferometer with arm cavities, no recycling cavities and no asymmetries. We must first make the arms resonant for both ω_1 and ω_2 , but not ω_L . The free spectral range (FSR) of the arm cavities must therefore satisfy $2\omega_m = (2n + 1) \cdot \text{FSR}_{\text{arm}}$ for integer n . Next, we try to move the interferometer to a dark fringe, by tuning one of the arms by 90° . A problem immediately arises, however; 90° corresponds to a different physical distance for each of ω_1 and ω_2 , so they won't both be at the dark fringe. This is shown in Figure 4.5. Thankfully, we can co-locate these fringes via careful choice of the Schnupp asymmetry.




Figure 4.5: Differently located dark fringes observed for both carriers in the two-carrier scheme. ω_1 is taken as the reference, placing its dark fringe at exactly 90° . ω_2 is a slightly higher frequency, requiring a smaller displacement of the arm to achieve the same phase change, and shifting the dark fringe.

Schnupp asymmetry

We need to shift one of the arms by some distance d that results in an odd multiple of 90° for both carriers, i.e.

$$\Delta\phi_1 = (2n + 1) \cdot 90^\circ,$$

$$\Delta\phi_2 = (2m + 1) \cdot 90^\circ,$$

$$\therefore \frac{\omega_2}{\omega_1} = \frac{2n + 1}{2m + 1}, \quad (4.1)$$

for integer n & m , where $\Delta\phi_{1,2}$ are the phase picked up by $\omega_{1,2}$ while propagating through the distance d . This distance is then given by

$$d = (2n + 1) \frac{\lambda_1}{4} = (2m + 1) \frac{\lambda_2}{4}, \quad (4.2)$$

where $\lambda_{1,2}$ are the wavelengths of $\omega_{1,2}$. This set of requirements leads to a slightly problematic conclusion: for a finite distance, d , we cannot always place an arbitrary pair of frequencies on a dark fringe at the interferometer output port. The closer

the frequencies are, the greater n & m have to be to accurately approximate the frequency ratio, and therefore the greater d must be.

In the specific case of aLIGO as discussed in [2], we take the usual laser wavelength of $\lambda_1 = 1064$ nm for ω_1 , and have $\omega_m = 2\pi \times 45$ MHz. This gives

$$\omega_1 = 2\pi \times \frac{c}{1064 \text{ nm}} \approx 2\pi \times 282 \text{ THz}, \quad (4.3)$$

$$\omega_2 = \omega_1 \pm 2\pi \times 90 \text{ MHz}, \quad (4.4)$$

giving a ratio

$$\frac{\omega_2}{\omega_1} = \frac{2n+1}{2m+1} \approx 1 + 3.19 \times 10^{-6}. \quad (4.5)$$

We can find a decent first approximation to n by setting $m = n - 1$, giving

$$\begin{aligned} \frac{2n+1}{2n-1} &= \frac{\omega_2}{\omega_1}, \\ \rightarrow n &= \frac{1}{2} \cdot \frac{\frac{\omega_2}{\omega_1} + 1}{\frac{\omega_2}{\omega_1} - 1} \\ &= \frac{1}{2} \cdot \frac{\omega_2 + \omega_1}{\omega_2 - \omega_1} \\ &= \frac{1}{2} \cdot \frac{2\omega_1 + 2\omega_m}{2\omega_m} \\ n &= \frac{\omega_1}{2\omega_m} + \frac{1}{2}, \end{aligned} \quad (4.6)$$

Putting in our values from before, we get

$$\begin{aligned} n &= 3130665.266 \\ &= 3130665 \end{aligned}$$

rounded down to the nearest integer, which corresponds to a length of

$$d = (2n+1) \frac{\lambda_1}{4} \approx 1.666 \text{ m}. \quad (4.7)$$

However, as n had to be rounded down, our phase-offset conditions will not be met exactly. In reality, this is not really an issue, as we do not need to use exactly 90 MHz as the frequency difference between our carriers. We can simply tweak the frequency ω_m to give a modified frequency ω'_m which will work for our desired Schnupp asymmetry, by doing the above calculations in reverse. Starting from Equation (4.6) for both our original ω_m and modified ω'_m ,

$$n = \left\lfloor \frac{\omega_1}{2\omega_m} + \frac{1}{2} \right\rfloor = \frac{\omega_1}{2\omega'_m} + \frac{1}{2}$$

we can rearrange to get

$$\omega'_m = \frac{\omega_1}{2 \left\lfloor \frac{\omega_1}{2\omega_m} + \frac{1}{2} \right\rfloor - 1}, \quad (4.8)$$

which for our original values gives $\omega'_m = 2\pi \times (45 \text{ MHz} + \sim 3.8 \text{ Hz})$.

Aside: Lengths in finesse

Care must be taken when putting this length into FINESSE, due to the distinction made between macroscopic lengths and microscopic tunings. See Section 2.5 of [27] for details on this convention; I will only state it here. When a space is given a length in FINESSE, an assumption is made that the reference carrier light ω_1 accumulates an integer multiple of 360° phase (i.e. the length is an integer multiple of λ_1). The microscopic tuning at a mirror's surface then allows us to achieve an arbitrary phase. However, for light at some other frequency ω_2 , the phase change through a space is calculated from the exact length given, rather than the length rounded to an integer multiple of λ_1 . This assumption simplifies some things, but it makes it possible to create a physically impossible situation: by specifying the macroscopic length of a space to greater precision than the main carrier wavelength λ_1 , we can make any two arbitrary frequencies 'resonate' in a space. The solution to this is to make sure that we only specify macroscopic lengths in FINESSE to the nearest λ_1 , and use a

Field	Reflection from PRM	Propagation through PRC	Reflection from arm	Total
ω_1	180°	0°	180°	360°
ω_L	180°	180°	0°	360°
ω_2	180°	0°	180°	360°

Table 4.1: Collection of phase needed for each field to resonate in the power recycling cavity i.e. have a total round-trip phase of 360° .

mirror's tuning, φ , to add a microscopic offset. The correct FINESSE equivalent of Equation (4.7) is then

$$d = \left\lfloor \frac{n}{2} \right\rfloor \lambda_1, \quad \varphi = ((2n + 1) \bmod 4) \cdot 90^\circ. \quad (4.9)$$

Power recycling

The next step is to add the power recycling cavity, which must be resonant for all fields present. As ω_1 and ω_2 are resonant in the arm cavities, they pick up 180° phase on reflection, whereas ω_L does not. We therefore need to introduce an extra 180° phase for only ω_1 and ω_2 . This can be achieved by requiring that $2\omega_m = (2n + 1) \cdot \text{FSR}_{\text{PRC}}$, to cause ω_L to pick up 180° phase through propagation, then detuning the power recycling mirror (PRM) by 90° , to cause all fields to gain another 180° on reflection. This is easier to see in table form in Table 4.1. The same problem from before still applies here; detuning the PRM by 90° corresponds to a different distance for each field. This is not a serious issue, however, as the large bandwidth of the PRC in aLIGO means that a slightly incorrect tuning does not affect circulating power noticeably.

Signal recycling

To complete the main interferometer, we add in our signal recycling cavity. This must be resonant for ω_L , but not ω_1 or ω_2 . From the phases each field picks up on

Property	Condition	Description
ω'_m	$\frac{\omega_1}{2\lfloor \frac{\omega_1}{2\omega'_m} + \frac{1}{2} \rfloor - 1}$	Modified separation frequency
L_{Schnupp}	$\left\lfloor \frac{1}{2} \cdot \left(\frac{\omega_1}{2\omega'_m} + \frac{1}{2} \right) \right\rfloor \lambda_1$	Schnupp asymmetry
FSR_{arm}	$\frac{2\omega'_m}{2n+1}$	Arm cavity free spectral range
FSR_{PRC}	$\frac{2\omega'_m}{2n+1}$	PRC free spectral range
FSR_{SRC}	$\frac{\omega'_m}{n}$	SRC free spectral range
FSR_{fc}	$\frac{2\omega'_m}{n}$	Filter cavity free spectral range

Table 4.2: Conditions on various parameters throughout the interferometer that are necessary for the two-carrier scheme. In each condition, n is an independent integer.

reflection from the arm cavities, we see that each field needs to pick up no extra phase from propagating through the SRC, i.e. $\omega_m = n \cdot \text{FSR}_{\text{SRC}}$.

Frequency dependent squeezing

Finally, we add frequency dependent input squeezing via the use of a filter cavity. The requirement here is that both sets of squeezed sidebands shown in Figure 4.4 experience the same frequency dependent rotation, i.e. $2\omega_m = n \cdot \text{FSR}_{\text{fc}}$.

4.2.2 Results

After each cavity length has been chosen to satisfy these conditions, and the carrier separation frequency ω_m appropriately modified, we are finally ready to model the two-carrier scheme in FINESSE. A comparison of the various schemes is shown in Figure 4.6, with the expected result that the quantum noise limited sensitivity of the two-carrier scheme is identical to that of homodyne readout. Note that FINESSE does not include radiation pressure calculations from beatings between the two carriers in the arm cavities, which could potentially limit sensitivity. In [2], however, it is shown that for aLIGO the mirror motion induced by these beatings is insignificant, and they are not problematic.




Figure 4.6: Comparison of the quantum noise limited sensitivity of heterodyne and homodyne readout with the two-carrier scheme in FINESSE. As expected, we see that heterodyne readout suffers a relative factor $\sqrt{1.5}$ penalty in shot noise to signal amplitude spectral density ratio, while the two-carrier scheme performs identically to homodyne readout. This performance is kept when frequency-dependent squeezing is added.

This result presents a great opportunity for improving the sensitivity of future detectors. Quantum noise is the limiting source of noise across much of a detector’s frequency range, especially at high frequencies where shot noise dominates. At lower frequencies, technical noises are significant. Some of these noises, such as laser frequency noise, arise in the interferometer input optics, and will couple to the output if there is any offset from the dark fringe.

Heterodyne readout is naturally immune to many technical noises, as the interferometer is kept at a perfect dark fringe for the main carrier light, and was used in the initial generation of gravitational wave detectors [61–63]. However, as can be seen from Figure 4.6, there is a relative reduction in shot noise limited sensitivity of $\sqrt{1.5}$ compared with homodyne readout. Due to the importance of improving quantum noise performance, upgrades to these detectors have involved a switch to DC readout [64], a variant of homodyne readout. This also introduces a direct path for technical noises to couple to the interferometer output, however, potentially limiting sensitivity at low frequencies.

The two-carrier scheme is a variant of heterodyne readout, and as such shares the same benefits of reduced technical noise coupling due to operating at the dark fringe. It also keeps the shot noise limited sensitivity of homodyne readout, shown in Figure 4.6. When frequency dependent squeezing is introduced, identical filter cavity parameters can be used for both homodyne readout and the two-carrier scheme, and the performance is the same. This is shown by the bottom two curves in Figure 4.6. Additionally, the generated squeezing only needs to be good at RF. Both homodyne and conventional heterodyne readout require squeezing at audio frequencies, which has proven difficult to achieve. As an added bonus, the two-carrier scheme is compatible with output path filter cavities as described in Chapter 3. These factors provide strong motivation for consideration of the two-carrier scheme for future detectors and upgrades.

4.3 Summary

The choice of local oscillator for signal readout is a key step in the design of a modern gravitational wave detector. Available schemes—such as DC, heterodyne, and balanced homodyne readout—each have tradeoffs between simplicity, sensitivity, and coupling of input noise sources to the interferometer output. Heterodyne readout in particular results in a reduction of quantum shot noise limited sensitivity of a factor between $\sqrt{2}$ and $\sqrt{1.5}$, known as the 3 dB penalty. We recently proposed a two-carrier scheme [2], a variant of heterodyne readout, which successfully eliminates the 3 dB penalty as well as allowing other potential sensitivity improvements by avoiding technical noises and only requiring good squeezing at radio frequencies.

In this chapter, I detailed the numerical simulations I carried out that verified the performance of the two-carrier scheme, and discussed the implications of this result. These simulations led me to find stringent requirements on the frequencies of each carrier and the Schnupp asymmetry, which must be controlled well enough to select a specific fringe of the interferometer.

Chapter 5

Broadening the bandwidth of gravitational wave detectors

Gravitational wave detectors use a photodetector at the output port of the interferometer to detect light from a gravitational wave signal. The output of the photodetector depends on the arrival rate of incoming photons, which is subject to Poisson statistics. There will therefore always be an error $\propto 1/\sqrt{P}$ [65], where P is the power at the photodetector. This is known as quantum shot noise. The error is independent of the frequency of the incoming light, so we could expect the shot-noise limited sensitivity curve of an interferometer to be a flat line. In reality, we see that the shot noise is indeed flat at lower frequencies, but above a certain frequency the sensitivity starts to decrease, giving the detector a finite bandwidth.

This bandwidth is due to the use of arm and signal-recycling cavities in the detector. In a tuned DRFPMI, all of the various cavities are chosen to be resonant or anti-resonant for the main carrier frequency, i.e. the carrier picks up an integer multiple of 180° phase on each round-trip. The signal sidebands produced when the carrier field is modulated by an incoming gravitational wave signal are therefore not fully resonant in the arm and signal-recycling cavities, as their frequency differs from

the main carrier by the frequency of the signal. The evolution of a light field with frequency ω over time t is given by

$$\vec{E}(t) = E_0 e^{-i\omega t}. \quad (5.1)$$

When modulated by a gravitational wave signal, sidebands \vec{E}_\pm are produced, offset by the signal frequency Ω ,

$$\vec{E}_\pm(t) \propto E_0 e^{-i(\omega \pm \Omega)t} \quad (5.2)$$

$$\propto E_0 e^{-i\omega t} e^{-i(\pm\Omega)t}. \quad (5.3)$$

The sidebands thus pick up a different amount of phase compared to the carrier—the upper sideband \vec{E}_+ picking up more phase per cavity round-trip, and the lower sideband \vec{E}_- picking up less. This is known as positive dispersion, and in this case is linear in Ω . As the signal frequency increases, the sidebands will eventually move out of resonance in the arm or signal recycling cavities, giving the reduction in sensitivity observed in current detectors.

To avoid this reduction in sensitivity, multiple systems have been proposed that exhibit *negative* dispersion, in which the lower signal sideband picks up less phase than the carrier, and the upper signal sideband picks up more. If this negative dispersion follows the correct relation, we can exactly cancel the positive dispersion, and broaden the bandwidth of the detector by restoring a flat shot noise response. The detector then becomes what is known as a white light cavity, responding equally well to any frequency.

One option to achieve negative dispersion is to include a special medium inside a cavity. This has been studied for many years, using atomic transitions to provide the negative dispersion, with limited success. While such white light cavities have been demonstrated in practice [66, 67], they have a number of caveats, such as absorption

of light by the medium, the requirement of a different laser wavelength to that used in gravitational wave detectors, and difficulties in producing the required medium [68]. In addition, a recent study [69] has shown that the quantum noise introduced by the mediums used so far means that they cannot lead to an overall reduction in noise while the cavity is stable. Future investigations need to look to alternative mediums or using feedback schemes to control the cavity while in the unstable regime.

In this chapter, I will briefly investigate an alternative mechanism that seeks to achieve this negative dispersion, known as the optomechanical filter [70]. I will describe the filter and some of the conditions that must be satisfied for it to work. I then use the numerical interferometer simulation software FINESSE to model the filter in a simple setup, without some of the assumptions made in theoretical treatments, and demonstrate the desired negative dispersion. From here, the effects of realistic parameters for optics such as lossy mirrors can be studied.

5.1 The optomechanical filter

The optomechanical filter is another recently proposed method of realising a white-light cavity, which uses the optomechanical coupling between a suspended mirror and multiple light fields circulating within a cavity. In basic operation, it consists of a cavity resonant at the frequency of some probe field ω_0 , and a suspended mirror with mechanical resonance at ω_m , pumped by a laser with frequency $\omega_0 + \omega_m$, such that the pump field is offset from cavity resonance by the mechanical resonance frequency. A modulated probe field such as that by an incoming gravitational wave with signal sidebands at $\omega_0 \pm \Omega$ (where Ω is the gravitational wave frequency) is then incident on the suspended mirror. This beats with the pump field, driving the suspended mirror at $\omega_m \pm \Omega$ and coupling the pump field into the signal field, leading to amplification of the signal field. An example detector configuration including an optomechanical filter is shown in Figure 5.1.




Figure 5.1: Simplified detector layout, with one possible configuration of the optomechanical filter in place. Here, the triangular cavity prevents feedback of the pump laser into the detector, and the signal recycling cavity acts as the filter cavity.

In its originally proposed layout [70], and in the updated transmission-readout setup [71] of which Figure 5.1 is a variant, the filter is required to be dynamically unstable, meaning that a control scheme would be required to implement it in a real detector. Additionally, some approximations are made that imposed extra limitations. The first is that the filter must be in the resolved sideband regime,

$$\Omega \ll \gamma \ll \omega_m, \quad (5.4)$$

where γ is the bandwidth of the filter cavity. Secondly, the single-mode approximation is made,

$$\Omega < \text{FSR}/2, \quad (5.5)$$

where FSR is the free spectral range of the filter cavity L/c . With these approximations, and the filter bandwidth being much larger than the mechanical resonator

bandwidth γ_m , the input-output relation of the filter becomes¹ [70]

$$\hat{a}_{\text{out}}(\Omega) \approx -\exp\left(\frac{2i\Omega}{\gamma_{\text{opt}}}\right) \hat{a}_{\text{in}}(\Omega), \quad (5.6)$$

where γ_{opt} is the negative mechanical damping rate due to optomechanical interaction, and under the extra condition that $\Omega \ll \gamma_{\text{opt}}$. This clearly exhibits linear negative phase dispersion. If γ_{opt} is tuned correctly, we can thus exactly cancel the positive dispersion seen by the signal sidebands within a limited frequency range.

Since the proposal of the transmission-readout setup, further analysis [72] has shown that a modified design with the probe field injected at the interferometer input can reproduce the increase in shot-noise limited bandwidth without relying on the resolved sideband and single-mode approximations. Additionally, by adjusting the frequency of the pump laser to account for the change in the resonance frequency of the mechanical oscillator due to the optical spring effect, the system can be made self-stabilising. These advances increase the optomechanical filter's prospects with regards to inclusion in future detector designs. We therefore need to be able to model the filter in standard interferometer modelling tools such as FINESSE. This will allow investigation of other aspects of the filter, such as the effect on the radiation pressure limited sensitivity of the detector, and provide useful experience for an eventual physical implementation. The next section describes the simulations I performed to verify the proposed quantum noise limited sensitivity enhancements of the filter.

¹For a full treatment of the maths behind Equation (5.6), see Appendix A. The differing signs between Equation (5.6) and [70] are due to the different mathematical conventions used for the Fourier transform and phase change on reflection from or transmission through a mirror.

5.2 Modelling the filter

5.2.1 Initial verification in FINESSE 2

The first step in modelling a full interferometer setup with an optomechanical filter is to check that the desired negative dispersion can be observed in a simpler system. However, modelling this with separate probe and pump lasers as described so far exposed a problem in the current version of FINESSE.

A core part of FINESSE's design is the separate modelling of carrier and signal fields. In the carrier simulation, all optics are static, and only the carrier fields (lasers and RF modulations) are present. The signal simulation is then modelled as small perturbations to this static state, and includes mirror motions, electronics and gravitational wave signal injections. Importantly, the signal sidebands produced by modulation of a carrier field can be thought of as being owned by that carrier field; they are defined relative to that carrier, rather than being independent fields.

In order to improve performance, FINESSE does not model the interactions between carriers and the sidebands of other carriers at a mirror's surface, as it was assumed that carriers would be separated by RF frequencies, while signals and mechanical resonances would be around audio frequencies. The beating frequency involved would then be much larger than any mechanical resonance frequency present. The susceptibility χ of an undamped simple harmonic oscillator with resonance frequency ω_m to a periodic driving force with frequency Ω is given by

$$\chi(\Omega) = \frac{1}{m(\omega_m^2 - \Omega^2)}, \quad (5.7)$$

which becomes very small for $\Omega \gg \omega_m$. This means that the beating can be safely ignored in many simulations, and results in fewer calculations being required at each optical surface. The optomechanical filter effect, however, arises precisely from these




Figure 5.2: The frequencies present in the optomechanical filter setup. The pump field at $\omega_0 + \omega_m$ beats with the signal sidebands, causing the suspended mirror in Figure 5.1 to oscillate. If $\omega_m < \text{FSR}/2$, where FSR is the Free Spectral Range of the filter cavity, the resultant fields around $\omega_0 + 2\omega_m$ will not be resonant, and so will have little effect. In the current version of FINESSE, the $\omega_m \pm \Omega$ beatings are ignored, and negative dispersion cannot be seen with separate probe and pump lasers.

interactions, and so they must be included for it to be observed. An illustration of the frequencies present is shown in Figure 5.2.

While unable to model the filter with separate pump and probe lasers in FINESSE, I was still able to observe negative dispersion in a simple setup with only a pump laser, as shown in Figure 5.3. By modulating the pump laser at a frequency $\omega_m + \Omega$, we effectively generate a single signal sideband at $\omega_0 - \Omega$. As this sideband is generated from the pump field, FINESSE correctly models the interaction, and negative dispersion is seen. The sideband at $\omega_0 + 2\omega_m + \Omega$ is far out of the filter cavity's bandwidth, so will have negligible effect.

This trick will not allow us to observe the sensitivity improvement from the filter in a full interferometer setup however, as there we need the carrier field and a signal injection at Ω . At this point, FINESSE 2 had been pushed to its limit. Instead of extending FINESSE 2 to include the necessary features, I decided that these would be part of FINESSE 3, a complete rewrite of FINESSE. More motivation for this decision and further discussion of FINESSE 3 in general is given in Chapter 8. In the next section, I will explain these features, and demonstrate negative dispersion with separate pump and probe lasers in FINESSE 3.

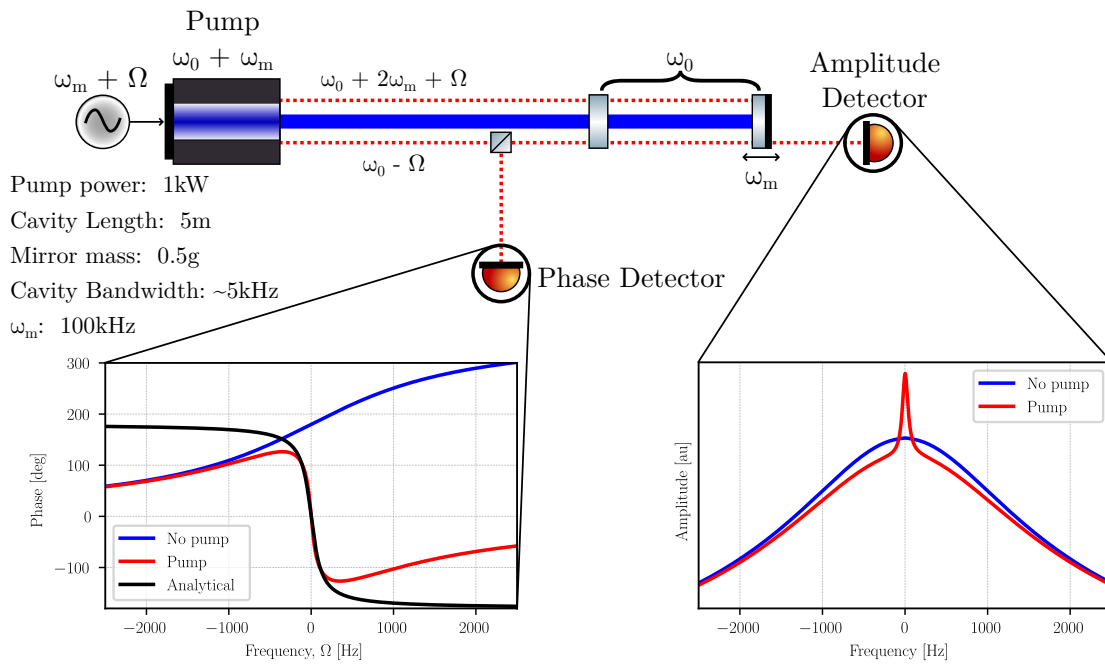


Figure 5.3: Negative dispersion as first observed in FINESSE, originally presented in [73]. This is a toy model, and doesn't include optical losses or control systems. With the pump switched off, we see the normal positive dispersion, as experienced by the signal sidebands in a detector's arm cavities. When the pump is switched on, this changes to linear negative dispersion around $\Omega = 0$ Hz. Also exhibited is *Optomechanically Induced Transparency* (OMIT) [74], as this is also dependent on being in the resolved sideband regime, where the pump laser is offset from cavity resonance by the mechanical oscillator frequency.

5.2.2 Improved modelling in FINESSE 3

FINESSE 3 is a complete rewrite of FINESSE in a mixture of Python [75] and Cython [76]. Here, I will only mention the changes from FINESSE relevant for modelling the optomechanical filter. For a more complete description, see Chapter 8.

Both FINESSE 2 and FINESSE 3 model interferometers as a set of components connected via nodes. In FINESSE 2, only optical components are available. Mechanics can be modelled, but only as they apply to optics. For example, a mirror can be given a mechanical transfer function for its suspension, but it is only used to modify the optical input to output coupling in the final matrix to be solved. This reduces the flexibility of FINESSE 2—to add the ability to model the optomechanical filter, we would either have to modify the existing mirror component to optionally perform the necessary calculations, or create a new component for the filter that re-uses the optical calculations from the normal mirror. Neither of these is desirable, as either the complexity of the mirror or the maintenance burden of the code increases, making future maintenance and expansion more difficult.

This limitation was one of my motivating factors in the development of FINESSE 3, and the new implementation of signal and quantum noise calculations. In this regard, the main difference in FINESSE 3 is the addition of mechanical nodes and components, greatly increasing the flexibility of the simulation. A mirror no longer performs mechanical computations internally, but instead exposes mechanical nodes. Separate mechanical components can then be connected to these nodes, such that the mirror only needs to perform optical calculations, and adding new suspension types etc. need not modify the mirror code. A simplified diagram of this structure is shown in Figure 5.4.

Another key change in FINESSE 3 that allows us to model the optomechanical filter is the ability of components to inform the simulation exactly which beatings between signals and carriers they will perform. The extra beatings required to model

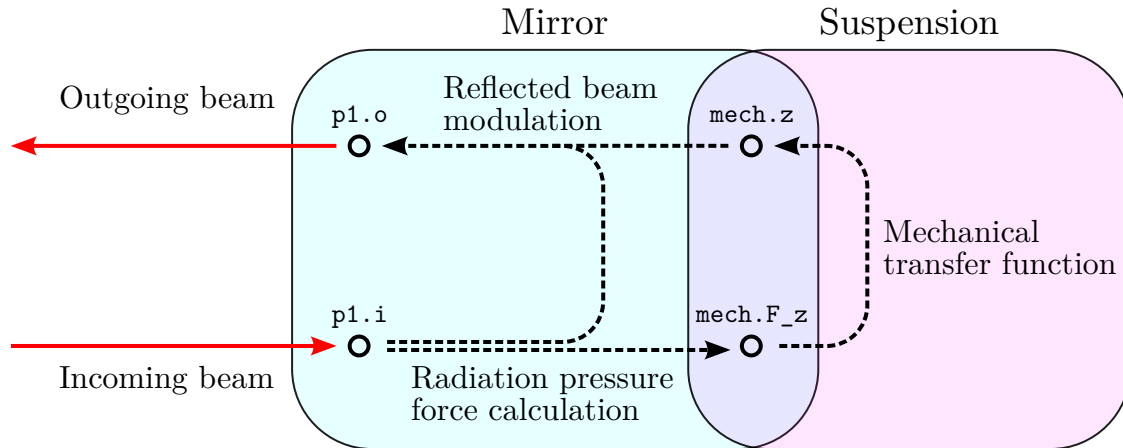


Figure 5.4: Simplified diagram of the couplings involved when a plane-wave beam is reflected from a suspended mirror in FINESSE 3, excluding unused nodes. The mirror and suspension are only responsible for optical and mechanical calculations respectively, interfacing via the `mech.z` and `mech.F_z` nodes. Crucially, the suspension can easily be swapped out for another kind (or some other mechanical system), or removed entirely, without having to modify the mirror.

the filter can then be ignored as before for most components. Instead, we can define a special ‘microresonator’ mechanical component, which informs the simulation that it wants to model these beatings, and then performs the necessary calculations. This can then be connected to a normal mirror component, and we are now ready to model the filter. Figure 5.5 shows a test layout with separate pump and probe lasers, with results given in Figure 5.6. The expected negative dispersion can be seen, and we are now ready to model the optomechanical filter in a full interferometer setup.

5.3 Summary

The optomechanical filter [70] is a proposed technique to reduce quantum noise in gravitational wave interferometers. By pumping a mechanical resonator in the signal recycling cavity with a laser, offset from the main carrier frequency by the mechanical resonance frequency of the resonator, linear negative dispersion can be seen around the carrier frequency. A recent study [72] has shown that this filter can be self-stabilising, removing the need for extra control loops which were one of the biggest potential barriers to a physical implementation.




Figure 5.5: Layout used to model negative dispersion in FINESSE 3, with separate pump and probe lasers. Parameters are as given in Figure 5.3, with results shown in Figure 5.6.




Figure 5.6: Negative dispersion observed in FINESSE 3. The ‘two laser’ curves from the layout in Figure 5.5 match those seen previously in FINESSE 2 using one laser. The optomechanical filter can now be modelled in a full interferometer setup.

In this chapter, I have reported on the first numerical tests of the optomechanical filter. I originally demonstrated negative dispersion with FINESSE 2 in a limited set up, and this work supported the development of the optomechanical filter idea by others. In order to model the filter in a full setup, new features were added in the development of FINESSE 3 to allow modelling of mechanical components and beatings between carriers and sidebands of other carriers. I have then demonstrated the filter performing as expected in a test using these beatings, laying the groundwork for more advanced modelling in the future.

Chapter 6

Noise requirements for ETpathfinder

In this chapter, I will describe modelling work I undertook for ETpathfinder, a prototype interferometer designed to test novel technologies expected to be employed in ET. This is a completely new interferometer, and is currently under construction. Many of the finer details of the optical design are therefore still under discussion. As part of this work, I was tasked with investigating the requirements on laser frequency noise and input mode cleaner (IMC) motion noise. I show how these noise sources are affected by asymmetries in the interferometer, and derive some rough order-of-magnitude estimates for requirements of these noises. Example parameters give these requirements at a signal frequency of 100 Hz as $\sim 0.2 \text{ Hz}/\sqrt{\text{Hz}}$ laser frequency noise, and $\sim 7 \times 10^{-13} \text{ m}/\sqrt{\text{Hz}}$ IMC motion noise.

6.1 ETpathfinder

In order to improve upon the sensitivity of current gravitational wave detectors, Einstein Telescope (ET) will make use of several novel technologies. These include

cryogenic test masses, a change in mirror material from the fused silica used in LIGO and Virgo to silicon, and correspondingly a shift in laser wavelength [36]. Before installation in a billion Euro facility, these technologies must be tested on a smaller and cheaper scale.

ETpathfinder is a detector designed to do exactly that; test many of the technologies required for ET in one facility, on a much smaller scale [29]. With arms only 10 m in length, compared to 4 km at LIGO and 10 km for ET, the aim of ETpathfinder is not to detect gravitational waves. In order to ensure that the desired strain sensitivity can be reached in ET, however, ETpathfinder must make sure to reach comparable order of magnitude displacement sensitivities for all of the usual noise sources present in a gravitational wave detector. Only then can the necessary technologies be fully tested.

To best explore the wide range of potential parameters for ET, ETpathfinder consists of two independent interferometers in one facility, each operating at different laser wavelengths: 1550 nm and 2 μ m. A schematic of the central optics for these two interferometers is shown in Figure 6.1. As the purpose of ETpathfinder is not to detect gravitational waves, the arms do not need to be positioned in the characteristic interferometer L shape. ETpathfinder therefore places both arms of each detector in the same direction, side by side. This allows separate cryogenic systems to be used to cool each interferometer's arm cavity mirrors, meaning that two temperatures of interest, 10 K and 120 K, can be investigated in parallel.

The facility to house ETpathfinder is currently under construction, and the finer details of its optical design are still in flux. It will reach similar displacement sensitivity to that planned for ET at the end of 'Phase 2' of its development [29], around 2028–2030. Before this, ETpathfinder will go through a more exploratory 'Phase 1' with a simpler setup, such as smaller test masses, no recycling cavities, and only 120 K cooling. While the required sensitivity of this initial phase, around

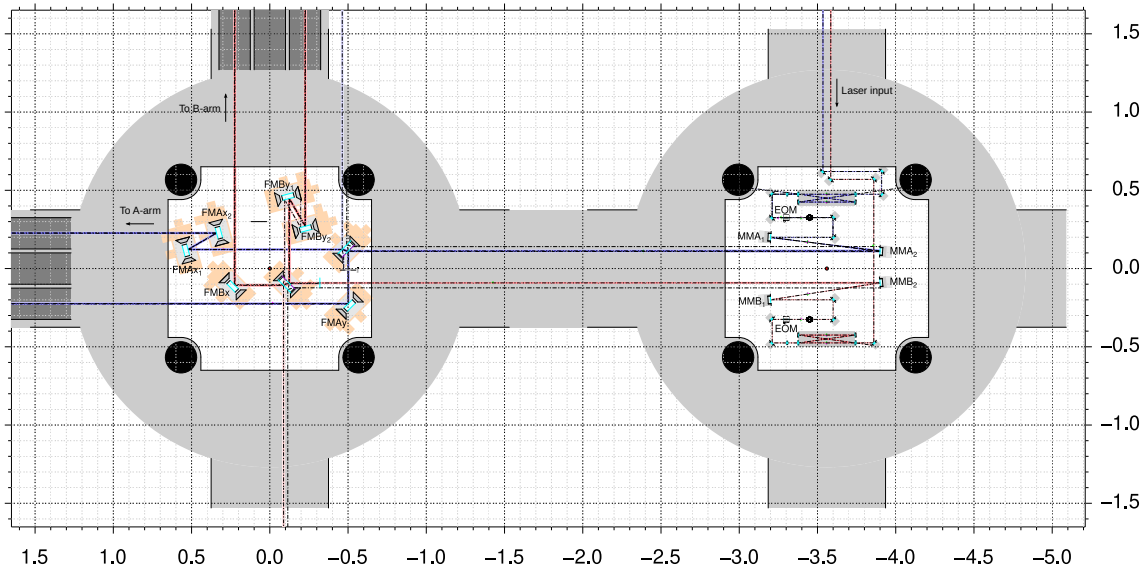


Figure 6.1: Current status of the physical layout of the input and central optics for ETpathfinder. Axes units are distances in metres. The red beams are from the 1550 nm ‘B’ interferometer, and the blue beams are the 2 μm ‘A’ interferometer. Each interferometer has both arm cavities in the same direction. This figure was produced from an OPTOCAD [77] layout made by other members of the ETpathfinder team, and is an updated version of Figure 6.2 from [29].

$1 \times 10^{-18} \text{ m}/\sqrt{\text{Hz}}$, is less than that of the final stages, choices in interferometer design made early on could limit the options available to ETpathfinder later in its lifecycle. We must therefore already understand how the sensitivity of ETpathfinder to various noise sources depends on detector design, and make estimates of requirements on these noise sources.

6.2 Noise sources in interferometer input optics

The basic layout of ETpathfinder investigated in this chapter is shown in Figure 6.2.

Three key parts have been highlighted. These are:

- The IMC — this is used to filter the input beam, by passively suppressing both intensity and frequency noise, and unwanted higher-order spatial modes. It is also used to actively stabilise the laser frequency.




Figure 6.2: Basic ETpathfinder layout used for the simulations in this chapter. Highlighted are the input mode cleaner (IMC), The DC offset (modelled here as a change in beamsplitter tuning), and the Schnupp asymmetry L_{Schnupp} .

- The DC offset — by microscopically tuning the central beamsplitter away from the dark fringe, a small amount of power is transmitted to the output to be used in signal readout.
- The Schnupp asymmetry L_{Schnupp} — a macroscopic difference between the distances from the beamsplitter to each of the arm cavities. This can be chosen to keep the main carrier at (or close to) the dark fringe, while still allowing light from RF sidebands through to the output to be used for control.

Throughout the rest of this thesis, I have largely focused on radiation pressure and shot noise, collectively known as quantum noise. There are a multitude of other noise sources which affect a gravitational wave detector, however. For ETpathfinder, laser frequency noise and IMC motion noise are two noise sources for which requirements were largely unknown, and which have implications for the initial design of the input optics.

Laser frequency noise is the short term jitter in a laser beam’s frequency, arising from many factors that affect the laser cavity such as mechanical vibrations, temperature fluctuations and internal quantum noise [78]. The laser is a closed system, and

we cannot normally individually stabilise the optics within, so this is treated as a single noise. Input mode cleaner motion noise is similar, caused by vibration of any of the mirrors in the input mode cleaner. Each of these noises effectively modulate the laser beam, and can be modelled as creating sidebands around the main laser frequency. If these sidebands reach the output of the interferometer, they will be detected along with the gravitational wave signal, reducing sensitivity.

6.2.1 Noise coupling and suppression

In an ideal Michelson interferometer, the arms are exactly symmetric, and there is no power at the output. This is true for any input beam frequency, and so any noise in the input beam does not couple to the output. Real interferometers contain asymmetries however, both deliberate and unintentional, which couple noise to the interferometer output:

- The DC offset is the deliberate introduction of a coupling from input to output, and so will also couple noise in the input beam. The strength of the coupling is dependent on the displacement of the beamsplitter relative to the wavelength of the input beam, and so is frequency dependent.
- The Schnupp asymmetry causes a phase difference in the light from each arm, chosen to be an integer multiple of 360° for the carrier frequency such that no coupling to the output is introduced. For other frequency fields such as noise sidebands, however, the phase change will be different, and a frequency dependent coupling will be present.
- Differences in the properties of the arm cavity mirrors can cause different reflectivities and signal responses of the arm cavities, changing the amount and phase of reflected light, and therefore coupling input noise to the interferometer output. This coupling is complicated, and depends on other asymmetries.

Modern interferometers therefore include various input noise suppression schemes, of which the IMC is a part. A good overview of the Advanced LIGO suppression scheme is given in [79], which shows how multiple stages of control systems and locking schemes are used to minimise laser frequency noise. From the requirements found in the next section and expected noise spectra for each noise source, we can then determine how effective these control systems must be to reach a target sensitivity.

6.3 Modelling noise requirements

The simulations in this chapter were performed using FINESSE [26] with a model based on a simplified ETpathfinder design with parameters proposed at the time I performed them. This basic layout is shown in Figure 6.2.

In order to accurately state requirements for a given noise source, we must have an overall noise budget and knowledge of all other noise sources present. For an interferometer whose design is still in flux, like ETpathfinder, this is not possible. Despite this, we can still make estimations for noise requirements and determine how these requirements are affected by interferometer parameters.

For a given noise source, the noise-limited displacement sensitivity is given by

$$S_n(f) \left[\frac{\text{m}}{\sqrt{\text{Hz}}} \right] = \frac{N_n(f) |H_n(f)|}{|H_s(f)|} \left[\frac{\frac{\text{U}}{\sqrt{\text{Hz}}} \cdot \frac{\text{W}}{\text{U}}}{\frac{\text{W}}{\text{m}}} \right], \quad (6.1)$$

where $N_n(f)$ is the input noise amplitude spectral density, $H_n(f)$ is the noise transfer function to the output, $H_s(f)$ is the signal transfer function, and U is the unit of the noise source e.g. Hz for laser frequency noise. To obtain an upper bound for the requirement on this noise source, we can ignore all other noise sources, and set this noise-limited sensitivity equal to the target noise budget S_{tot} . This gives us the

maximum allowable noise spectrum,

$$N_n(f) \left[\frac{\text{U}}{\sqrt{\text{Hz}}} \right] = \frac{S_{\text{tot}}(f) |H_s(f)|}{|H_n(f)|} \left[\frac{\frac{\text{m}}{\sqrt{\text{Hz}}} \cdot \frac{\text{W}}{\text{m}}}{\frac{\text{W}}{\text{U}}} \right]. \quad (6.2)$$

Using FINESSE, I calculated scalings for the laser frequency and IMC motion noise requirements with Schnupp asymmetry and DC offset of ETpathfinder. These are shown in Figures 6.3 to 6.6. Where held constant, the Schnupp asymmetry was chosen to be 10 cm and the DC offset was 9.1 nm (equivalent to 2.114° in FINESSE at a wavelength of $\lambda = 1550$ nm); in a real interferometer, these are set by other requirements. A single frequency of 100 Hz was chosen for plotting, but the scaling holds true for all frequencies. We see that for each noise source and asymmetry, the requirement on the noise is inversely proportional to the magnitude of the asymmetry. This can inform later choices for the Schnupp asymmetry and DC offset; they should be as small as possible while still satisfying other requirements, to minimise the effect of noises arising in the input path of the interferometer.




Figure 6.3: Laser frequency noise requirement as a function of Schnupp asymmetry. The requirement on the noise is inversely proportional to the asymmetry. The dashed line at 10 cm shows the value used for the overall noise curves in Figure 6.7.




Figure 6.4: Laser frequency noise requirement as a function of DC offset. The dashed line at $\sim 2^\circ$ shows the value used for the overall noise curves in Figure 6.7, and corresponds to a dark port power of around 4 mW. *N.B. The beamsplitter is modelled with an angle of incidence of 0° , hence the second dark fringe being located at 90° tuning.*




Figure 6.5: IMC motion requirement as a function of Schnupp asymmetry. As for laser frequency noise in Figure 6.3, the requirement on the noise is inversely proportional to the asymmetry, and the dashed line at 10 cm shows the value used for the overall noise curves in Figure 6.7.

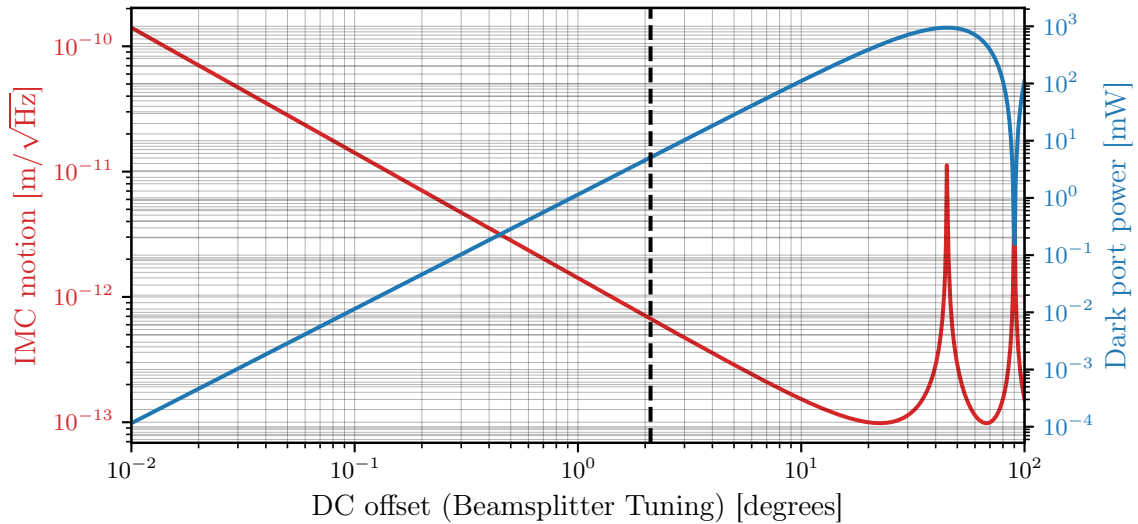


Figure 6.6: IMC motion noise requirement as a function of DC offset. The dashed line at $\sim 2^\circ$ shows the value used for the overall noise curves in Figure 6.7, and corresponds to a dark port power of around 4 mW. *N.B.* The beamsplitter is modelled with an angle of incidence of 0° , hence the second dark fringe being located at 90° tuning.




Figure 6.7: Noise requirements for both laser frequency and IMC motion noise in ET-pathfinder. The Schnupp asymmetry L_{Schnupp} and DC offset were chosen somewhat arbitrarily to give an idea of the orders of magnitude involved; these are set by other requirements in a real interferometer.

Target noise spectra for both laser frequency noise and input mode cleaner motion noise are shown in Figure 6.7, along with the total noise target at the time of simulation. The scaling with frequency of each noise source can be understood as arising from a few factors. Firstly, the coupling of noise in the input to the output via the DC offset and Schnupp asymmetry is proportional to the frequency f . For IMC motion noise, which is pure phase noise, this therefore leads to a $1/f$ scaling in the noise requirement compared to the total noise curve. Laser frequency noise has an additional factor; small amplitude frequency modulation can be understood as phase modulation with a modulation index proportional to $1/f$ [27]. The frequency scalings therefore cancel, leading to a flat overall response.

During the course of these investigations, it became apparent that any imbalance in the finesse of each arm cavity strongly increases the coupling of these noises to the output, thus making the requirements stricter. Such an imbalance could arise from e.g. the etalon effect, in which some light is reflected from the anti-reflective coating on the ITM. This slightly alters the reflectivity and transmissivity of the




Figure 6.8: IMC motion noise requirements vs ITM reflectivity-transmissivity imbalance for the X arm. A positive x-axis value means increased reflectivity, and decreased transmissivity. Loss is kept constant. The blue curve shows the requirement on IMC motion noise, while the red line shows the corresponding arm cavity finesse.

ITM as a whole, and therefore arm cavity finesse. Other potential sources include ice build-up on the ITMs when cooling them to cryogenic temperatures. The effect of such an imbalance is shown for IMC motion noise in Figure 6.8. As we can see, the effect is marked; a few ppm shift in reflectivity and transmissivity leads to an order of magnitude decrease in the allowed input motion cleaner noise. As for other asymmetries, the same scaling applies to laser frequency noise.

6.4 Summary

Einstein Telescope will make use of several novel technologies to reach its target noise budget. Before then, knowledge of and expertise with these technologies must be gained by testing them together on a smaller scale. ETpathfinder aims to do precisely this, eventually reaching comparable displacement sensitivity to Einstein Telescope.

Before reaching its final sensitivity target, ETpathfinder will go through a much

more exploratory phase, with less well defined requirements. Design choices made early on may limit future options for development, however, and so we must already start to investigate the impact and behaviour of potential noise sources.

In this chapter, I have described the simulation work I performed for ETpathfinder. I have given rough estimates for the requirements on laser frequency noise and input mode cleaner motion noise at a signal frequency of 100 Hz as $\sim 0.2 \text{ Hz}/\sqrt{\text{Hz}}$ laser frequency noise, and $\sim 7 \times 10^{-13} \text{ m}/\sqrt{\text{Hz}}$ IMC motion noise. I have then shown how these requirements scale with asymmetries in the interferometer. This will help inform choices for deliberate asymmetries, such as the Schnupp asymmetry and the DC offset, and shows the importance of controlling unintentional asymmetries, such as a finesse imbalance between the two arm cavities.

Chapter 7

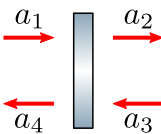
Quantum noise modelling in

FINESSE 3

Throughout much of this thesis, I have shown the results of quantum noise simulations using the frequency domain modelling tool FINESSE [26], and have described some of the extensions that myself and others have implemented in the latest version, FINESSE 3, in order to model newly proposed detector schemes. The time has now come to discuss how these simulations are performed.

I was responsible for a complete reimplementaion of the quantum noise calculations in FINESSE 3. In doing so, I also made generalisations to allow more noise modelling features to be added, and wrote tests that validated some of the code for the first time. Appendix B describes one such set of tests, comparing the radiation pressure calculations performed by FINESSE 3 with analytical expressions for a variety of scenarios.

Much of the core workings and implementation of quantum noise in FINESSE 3 are similar to that seen in previous versions. This has already been discussed in detail in other works [27, 80, 81]. In this chapter, I will first provide a summary of how FINESSE models interferometers in general, and the methods used to calculate

$$\begin{pmatrix} a_2 \\ a_4 \end{pmatrix} = \begin{pmatrix} it & r \\ r & it \end{pmatrix} \begin{pmatrix} a_1 \\ a_3 \end{pmatrix}$$


The diagram shows a vertical blue bar representing a mirror. On the left side, there are two red arrows: one pointing right labeled a_1 and one pointing left labeled a_4 . On the right side, there are two red arrows: one pointing right labeled a_2 and one pointing left labeled a_3 .

Figure 7.1: The coupling matrix for the fields at a mirror with amplitude reflectivity r and transmissivity t [27]. The factor i on transmission is a consequence of the convention used for the phase relation at a reflective surface in FINESSE (see Section 2.4 of [27] for a detailed explanation).

quantum noise. I will then show additions I made when reimplementing these quantum noise calculations in FINESSE 3. For a broader overview of FINESSE 3 and the design choices made, see Chapter 8.

7.1 Frequency domain modelling

There are multiple ways one can go about modelling an optical system, each with various benefits and drawbacks. For example, time domain models can capture complex dynamics of a system, whereas frequency domain models are much more computationally efficient for the subset of behaviour they are able to model.

FINESSE is a frequency domain tool. As such, it makes assumptions about the systems being modelled, the main one being that the system is in a steady state (that is, the power in each of the optical fields present does not change over time). The system can then be described by a set of linear equations, allowing efficient computation. Certain non-linear effects, such as radiation pressure forces, can be linearly approximated around some operating point as small perturbations to this steady state.

To understand how FINESSE models interferometers, consider a single light field incident on each side of a mirror with amplitude reflectivity r and transmissivity t , as shown in Figure 7.1. We can write down the outgoing fields a_2 & a_4 in terms of

the incoming fields a_1 & a_3 , giving

$$a_2 = ita_1 + ra_3, \quad (7.1)$$

$$a_4 = ra_1 + ita_3. \quad (7.2)$$

This set of linear equations can also be expressed in matrix form, as in Figure 7.1. Similar coupling matrices can be written for various other optical components, such as spaces, beamsplitters, modulators and so on. The job of FINESSE is to link all of these coupling matrices together, in order to calculate the fields at some desired output nodes. One of the most flexible ways to do this is to write one equation for every field [27], which for the mirror results in

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ -it & 1 & -r & 0 \\ 0 & 0 & 1 & 0 \\ -r & 0 & -it & 1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix} = \begin{pmatrix} a_1 \\ 0 \\ a_3 \\ 0 \end{pmatrix}. \quad (7.3)$$

This has the basic form

$$\mathbf{M}\vec{a} = \vec{b}, \quad (7.4)$$

where \vec{b} is the ‘right-hand side’ (RHS) vector which contains the known input field amplitudes, and \vec{a} is the ‘solution’ vector which contains the amplitude of the field at all points in the model. \vec{a} can then be found by inverting \mathbf{M} to give

$$\vec{a} = \mathbf{M}^{-1}\vec{b}. \quad (7.5)$$

This is the approach that FINESSE takes. One benefit of this is that construction of the matrix \mathbf{M} is fairly straightforward, even for a complex optical system, as the elements are just the normal component couplings (up to a minus sign). For example, the setup shown in Figure 7.2 is represented by




Figure 7.2: Setup in which a laser beam of amplitude E_0 propagates through a space of length D before impinging on a mirror, with the coupling matrices for each component shown in the same form as in Figure 7.1. $k = 2\pi/\lambda$ is the wavenumber of the beam.

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ -e^{-ikD} & 1 & 0 & 0 & 0 & 0 \\ 0 & -it & 1 & -r & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & -r & 0 & -it & 1 & 0 \\ 0 & 0 & 0 & 0 & -e^{-ikD} & 1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \end{pmatrix} = \begin{pmatrix} E_0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (7.6)$$

Setups involving e.g. multiple optical frequencies and higher-order spatial modes (HOMs) can also be modelled in this way, by treating each combination of (location, frequency, HOM) as a separate field, and including the relevant couplings for each component.

A potential drawback is the size of the matrix \mathbf{M} , whose dimensions grow linearly with the number of fields to be modelled. An interferometer simulation may contain thousands of fields if multiple frequencies and HOMs are included, leading to memory requirements on the order of gigabytes if every element of \mathbf{M} is to be stored. However, the vast majority of these elements are zero, especially for setups with many components. \mathbf{M} is therefore what is known as a sparse matrix. Many efficient data structures and algorithms to solve equations like Equation (7.5) for sparse matrices have already been written; the one used by FINESSE is called KLU [82]. These avoid storing and performing calculations with the zeros in \mathbf{M} , alleviating the problem. Note that this is not applicable to \mathbf{M}^{-1} , which is often dense, and cannot reasonably

be stored in the memory of a modern computer for complex interferometer models. We therefore cannot simply invert \mathbf{M} once and calculate the output to any arbitrary input by changing \vec{b} and performing a matrix multiplication; the original equation must be solved each time.

7.1.1 Simulating signals

A common task when modelling gravitational wave interferometers is to see how a system in a steady state will respond to a sinusoidal perturbation of one of its parameters. This could be, for example, an incoming gravitational wave changing the length of a space, or a mechanical vibration causing a mirror to move. For a small, infinitely long, sinusoidal signal with frequency Ω , a field with frequency ω_0 interacting with such a perturbation will be modulated to create two extra fields at $\omega_0 \pm \Omega$ (see Appendix B). These extra fields are known as *signal sidebands* (or audio sidebands, due to their Hz–kHz range) of the original central frequency or *carrier* field. By injecting a signal at some input point of the model and measuring the corresponding response of a photodiode demodulated at Ω , we can calculate a transfer function from the input to the output.

In theory, we could include these signal injections in our original matrix \mathbf{M} . This brings a number of problems, however. The signals must be small enough that the response is to remain predominantly linear, and including numbers separated by many orders of magnitude in \mathbf{M} is undesirable for numerical accuracy reasons. Additionally, we wish to be able to model some non-linear effects, such as the radiation pressure forces discussed later, which cannot be represented via a single matrix solve. For these reasons, FINESSE splits the matrix into two parts. The original matrix \mathbf{M} as described in the previous section is renamed as the *carrier matrix* \mathbf{M}_c , and a new matrix containing only the signal sidebands is introduced as the *signal matrix* \mathbf{M}_s . The modulations creating the signal sidebands are assumed to be small enough that

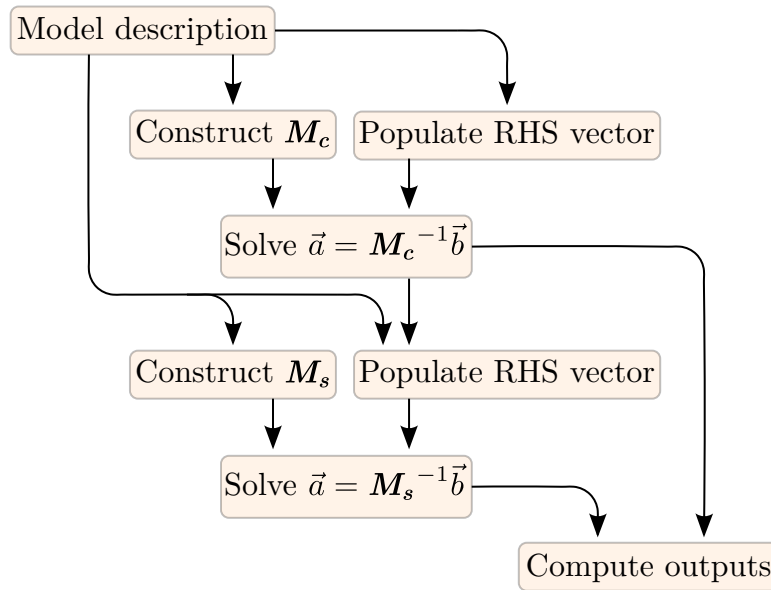


Figure 7.3: Simplified flowchart of a FINESSE simulation, with time flowing roughly from top to bottom. Not pictured here are radiation pressure effects, which necessitate the use of the carrier simulation solution to construct \mathbf{M}_s as explained in Section 7.1.2, or the quantum noise simulation discussed later in this chapter.

any coupling back from the sidebands into the carrier fields is negligible.

FINESSE can therefore perform a simulation as follows. First, we construct \mathbf{M}_c with all of the frequencies present prior to signal injection (the carrier fields) and populate the RHS vector with our known inputs, i.e. the output of any lasers present. This is then solved via Equation (7.5). Next, we construct \mathbf{M}_s , and use our carrier field solution to populate the RHS vector with correctly scaled signal sidebands. Again, this is solved via Equation (7.5). Finally, the solutions to each of these can be used to calculate transfer functions and photodiode outputs. This process is shown as a flow chart in Figure 7.3. Note that the results of the carrier simulation can be re-used for many different signal simulations, allowing us to, for example, sweep the frequency of an input signal and calculate the frequency response of a transfer function, without having to re-solve the carrier matrix for each point.

7.1.2 Radiation pressure

Splitting the simulation in two as previously described allows us to model some non-linear effects such as radiation pressure. The radiation pressure force exerted on a mirror such as that shown in Figure 7.1 is given by

$$F_{\text{rad}} = \frac{P_1 - P_2 - P_3 + P_4}{c}, \quad (7.7)$$

where P_n is the power in the field a_n . If a single field with power P_i is incident on this mirror from the left hand side, the radiation pressure force will then be

$$F_{\text{rad}} = \frac{1 + R - T}{c} P_i, \quad (7.8)$$

where R and T are the power reflectivity and transmissivity of the mirror. Consider an incoming carrier field of frequency ω_0 with signal sidebands at $\omega_0 \pm \Omega$,

$$E_i = E_0 e^{i\omega_0 t} + E_+ e^{i(\omega_0 + \Omega)t} + E_- e^{i(\omega_0 - \Omega)t}, \quad (7.9)$$

where E_0 and E_{\pm} are the complex amplitudes of the carrier and upper and lower sidebands respectively. In this complex notation, we define the power in the field as

$$P_i = |E_i|^2 = E_i E_i^*, \quad (7.10)$$

which gives

$$\begin{aligned} P_i = & |E_0|^2 + |E_+|^2 + |E_-|^2 \\ & + (E_0^* E_+ + E_0 E_-^*) e^{i\Omega t} \\ & + (E_0 E_+^* + E_0^* E_-) e^{-i\Omega t} \\ & + E_+ E_-^* e^{i2\Omega t} + E_+^* E_- e^{-i2\Omega t}. \end{aligned} \quad (7.11)$$

The first line represents the static component of the force; in FINESSE, this is assumed to be balanced out by a control system to maintain the steady state, and so can be ignored. The last line represents beatings between the two sidebands. The amplitudes of the sidebands are assumed to be much smaller than that of the carrier, such that these terms will be insignificant compared to the terms containing E_0 and can also be neglected. This leaves us with

$$P_i = (E_0^* E_+ + E_0 E_-^*) e^{i\Omega t} + \text{c.c.}, \quad (7.12)$$

where c.c. means the complex conjugate of the right hand side. This represents the beating of the carrier with its sidebands, which will induce a sinusoidal force on the mirror with frequency Ω and magnitude dependent on the amplitude of each field,

$$F_{\text{rad}} = \frac{1 + R - T}{c} (E_0^* E_+ + E_0 E_-^*) e^{i\Omega t} + \text{c.c.} \quad (7.13)$$

In Appendix B.1.1, I show that for a suspended mirror, a similar sinusoidal force will cause sinusoidal motion of the mirror, which in turn will modulate the reflected carrier field, coupling it with its sidebands. Importantly, the magnitude of the motion, and therefore the coupling between the carrier and its sidebands, is dependent on the power in the fields. We therefore could not model this effect with one matrix containing the carrier and signal fields together, as we would have to already know the solution for the incoming fields to construct the matrix!

By assuming the amplitudes of the signal sidebands are small and performing the simulation in two parts, however, the problem becomes solvable. Briefly put, in the signal simulation we already know the amplitude of the carrier field. This is allowed because we treat the signal sidebands as small, such that any coupling between them and the carrier field does not change the amplitude of the carrier significantly. The radiation pressure force given in Equation (7.13) is then linear with respect to the sideband amplitude. The response of the mirror, and thus the coupling from the

carrier field into the sidebands, is therefore also linear in sideband amplitude. Finally, this implies that the coupling from incoming to reflected sidebands is independent of sideband amplitude, depending only on the amplitude of the carrier field, and we are able to model radiation pressure effects. For a more detailed mathematical treatment of this reasoning, along with verification that the implementation in FINESSE behaves as expected, see Appendix B.

We have now seen in broad terms how FINESSE performs simulations of coherent fields, and possess the necessary background to continue. The rest of this chapter is focused on how to model incoherent fields, specifically looking at quantum noise.

7.2 Quantum noise spectra

A coherent state, such as that produced by a laser, can be described by two quadratures; its amplitude a and phase ϕ . Together these form a conjugate pair, and so cannot be precisely measured at the same time [83]. This uncertainty can be thought to arise from the spontaneous creation and annihilation of photons incoherently at all frequencies, which leads to a background of normally distributed noise. In the absence of effects such as squeezing, the variances in amplitude σ_a^2 and phase σ_ϕ^2 of this noise are equal and at the minimum allowed by the Heisenberg uncertainty principle, i.e.

$$\sigma_a \sigma_\phi = \frac{\hbar\omega}{2}, \quad (7.14)$$

where ω is the light frequency. This is known as vacuum noise. The quantum noise we observe in a laser is then the addition of a coherent carrier with this background, as shown in Figure 7.4.

In Section 6.1 of [27], Bond et. al. describe how we can think of this quantum noise background in a more natural way for the purposes of gravitational wave interferometry. They start with a single carrier field with amplitude a_0 and frequency




Figure 7.4: Vacuum noise fluctuations (a) in the absence of any carrier and (b) superimposed on a coherent laser field with complex amplitude a_0 and corresponding phase $\phi = \arg(a_0)$. The presence of the vacuum fluctuations causes uncertainty in the amplitude and phase of the laser field, with variances of σ_a^2 & σ_ϕ^2 respectively.

ω_0 with a continuum of noise,

$$E(t) = \frac{a_0}{2} e^{i\omega_0 t} + \frac{1}{2} \int_0^\infty q(\omega) e^{i\omega t} d\omega + \text{c.c.} \quad (7.15)$$

where $q(\omega)$ represents the quantum fluctuations of the field, and *c.c.* represents the complex conjugate part required to make $E(t)$ real. In the notation used by FINESSE and elsewhere in this thesis, in which only the positive half of the frequency spectrum is considered, this becomes

$$E(t) = a_0 e^{i\omega_0 t} + \int_0^\infty q(\omega) e^{i\omega t} d\omega. \quad (7.16)$$

By transforming ω to be in reference to the carrier frequency ω_0 , considering only those noise fields at frequencies within a gravitational wave detector's bandwidth B (and thus only the fields that will affect sensitivity), and assuming that B (typically a few kHz) is much less than ω_0 (typically a few THz), it is shown that we can thus treat the quantum noise as sidebands of the carrier described by

$$E(t) = e^{i\omega_0 t} \left[a_0 + \int_{-B}^B q(\omega_0 + \Omega) e^{i\Omega t} d\Omega \right], \quad (7.17)$$

where

$$q(\omega) = n_a(\omega) + in_\phi(\omega) \quad (7.18)$$

with $n_{a,\phi}(\omega)$ being normally distributed probability density functions with mean $\mu_{a,\phi} = 0$ and variance $\sigma_{a,\phi}^2$.

7.3 Noise spectral density

When talking about modelling quantum noise, the measurable quantity we are really interested in is the power spectral density (PSD) of noise in the output photocurrent of a photodiode. In general, the PSD quantifies the frequency distribution of power present in a signal. As quantum noise is stochastic, the amplitude and phase of a sinusoid at some frequency are ill-defined. Additionally, from Equation (7.18) we see that the expectation value of a quantum noise field $q(\omega)$ is [27]

$$\langle q(\omega) \rangle = \langle \mu_a \rangle + i \langle \mu_\phi \rangle = 0. \quad (7.19)$$

The PSD, on the other hand, is well-defined—for a frequency-domain signal $x(\omega)$, the single-sided (positive frequency spectrum) PSD can be written in terms of the covariance,

$$S_{xx}(\omega)\delta(\omega - \omega') = 2\langle x(\omega)x^*(\omega') \rangle. \quad (7.20)$$

In the case of vacuum noise, the fields at different frequencies are incoherent, and the variance is limited by the uncertainty principle, giving

$$\langle q(\omega)q^*(\omega') \rangle = \frac{\hbar\omega}{2}\delta(\omega - \omega'), \quad (7.21)$$

$$\langle q(\omega)q(\omega') \rangle = 0. \quad (7.22)$$

The photocurrent I of a photodiode is proportional to the power P of the incident

field, which is given by

$$P(t) = E(t)E^*(t). \quad (7.23)$$

For a single carrier with quantum noise sidebands, we can insert Equation (7.17) to obtain

$$\begin{aligned} P(t) = & |a_0|^2 \\ & + a_0 \int_{-B}^B q^*(\omega_0 + \Omega) e^{-i\Omega t} d\Omega \\ & + a_0^* \int_{-B}^B q(\omega_0 + \Omega) e^{i\Omega t} d\Omega \\ & + O(q^2). \end{aligned} \quad (7.24)$$

Here the first line is simply the DC power of the carrier, the next two represent the beating of the carrier with its quantum noise sidebands, and the last line represents beatings between noise sidebands, which are assumed to be negligible.

The positive frequency photocurrent spectrum within the interferometer bandwidth is then simply the Fourier transform [27],

$$I(\Omega) = a_0 q^*(\omega_0 - \Omega) + a_0^* q(\omega_0 + \Omega), \quad (7.25)$$

where the proportionality constant between I and P has been ignored. Applying Equation (7.20), and with the notation

$$\begin{aligned} q_{\pm} & \equiv q(\omega_0 \pm \Omega) \\ q'_{\pm} & \equiv q(\omega_0 \pm \Omega'), \end{aligned}$$

we obtain

$$\begin{aligned} S_I(\Omega)\delta(\Omega - \Omega') &= 2a_0a_0^* \left(\langle q_+q_+^* \rangle + \langle q_-^*q_- \rangle \right) + 2a_0^2 \langle q_-^*q_+^* \rangle + 2a_0^{*2} \langle q_+q_- \rangle \\ &= 2P_0 \left(\langle q_+q_+^* \rangle + \langle q_-^*q_- \rangle \right). \end{aligned} \quad (7.26)$$

From here we can plug in the covariance for vacuum noise from Equation (7.21) to get

$$S_I(\Omega) = 2P_0\hbar\omega_0. \quad (7.27)$$

For the purpose of understanding how quantum noise is modelled in FINESSE, the form of Equation (7.26) is important. The noise PSD is calculated from some combination quantum noise field covariances, scaled by the power in the carrier field. While the end result here is simple, cases involving multiple carriers, photodiode demodulations, and squeezed quantum noise fields become significantly more complex—see [81] for derivations of the PSD in these scenarios. The building blocks remain similar, however. As long as we know the amplitudes of any carrier fields and the covariance of all quantum noise fields that will contribute to the output of a photodiode, we can calculate the noise PSD.

One way to express the covariance of a finite set of quantum noise fields q_1, q_2, \dots, q_N is via the variance-covariance matrix,

$$\mathbf{V} = \begin{bmatrix} \langle q_1q_1^* \rangle & \langle q_1q_2^* \rangle & \cdots & \langle q_1q_N^* \rangle \\ \langle q_2q_1^* \rangle & \langle q_2q_2^* \rangle & & \\ \vdots & & \ddots & \\ \langle q_Nq_1^* \rangle & & & \langle q_Nq_N^* \rangle \end{bmatrix}. \quad (7.28)$$

This contains the covariance of every pair of quantum noise fields, with single field variances along the diagonal. For pure vacuum noise, all of the off-diagonal elements will be zero, and the matrix is trivial. This is not the case in the presence of effects

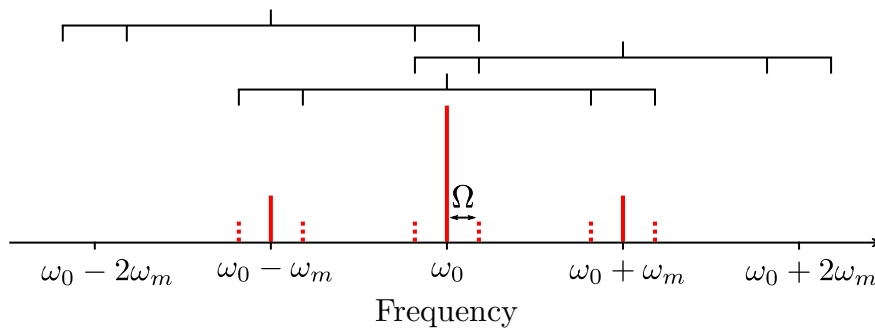


Figure 7.5: Demodulation of three carriers with signal sidebands. In this setup, the central carrier at ω_0 has been modulated by both a modulator with frequency ω_m , and a signal with frequency Ω . The three trees at the top show which frequencies each carrier beats with when demodulating twice, first by ω_m and then by Ω . The central tree includes only frequencies present in the simulation, but the two outer trees contain contributions from pure vacuum noise at frequencies not usually modelled.

such as squeezing, however, which causes correlations between fields. In Section 7.6, we will see how the variance-covariance matrix is used to calculate quantum noise PSDs in FINESSE.

7.4 Photodiode demodulation

As previously mentioned, calculating the quantum noise PSD in the output of a demodulated photodiode is significantly more complex than in the non-demodulated case. Consider the setup shown in Figure 7.5, in which three carrier fields with signal sidebands are incident on a photodiode demodulated at the carrier separation frequency ω_m and the signal frequency Ω . The trees at the top show which sidebands each carrier will beat with. From these, a couple of the complicating factors are apparent.

Firstly, the quantum noise sidebands around $\omega_0 \pm 2\omega_m$ will contribute to the photodiode output, but there are no signal fields present at these frequencies. This suggests that we need to include these extra frequencies in the quantum noise simulation, reducing performance. A workaround can be found, however, by assuming that the separation between carriers ω_m is much greater than the resonance frequencies

of any mechanical oscillators in the model. As the response of a mechanical oscillator falls off like $1/\omega^2$ above resonance, the radiation pressure effects from these extra quantum noise fields beating with any carriers will be negligible, and thus so will any squeezing effects on them. We can therefore treat these as pure vacuum noise, and account for them when calculating the photodiode output without having to propagate them through the model. Note that this assumption is explicitly invalid in the case of the optomechanical filter described in Chapter 5, but in that case the necessary frequencies are already included in the signal simulation.

The other complication shown in Figure 7.5 is that the quantum noise fields at $\omega_0 \pm \Omega$ are included in the demodulation trees of both carriers at $\omega_0 \pm \omega_m$. Special care must be taken to ensure that the contributions from these fields are accounted for correctly, as they will be coherent between the two carriers [81].

In [81], an algorithm is derived by which the various contributions of each quantum noise field to the output PSD of a demodulated photodiode can be calculated. This is the algorithm used by FINESSE and is well summarised in Appendix D of [80]. The rest of this section (Section 7.4) is a replication of this summary for completeness, although I switch from linear to angular frequency for coherence with the rest of this chapter:

Consider N_C carrier fields with frequencies $\omega_{C,i}$, where $i \in \{0, 1, \dots, N_C - 1\}$, incident on a photodiode performing N_D demodulations at frequencies $\omega_{D,j}$ and phases $\phi_{D,j}$, where $j \in \{0, 1, \dots, N_D - 1\}$. We construct a demodulation frequency matrix \mathcal{F} with dimensions $N_C \times 2^{N_D}$ and a demodulation phase vector $\vec{\Phi}_D$ of length 2^{N_D} . \mathcal{F} contains all of the frequencies that will be mixed due to demodulation (e.g. those present in the demodulation trees in Figure 7.5), and $\vec{\Phi}_D$ lists the total

demodulation phase of each mix.

$$\beta_{ij} = (-1)^{\lfloor i/2^{N_D-j-1} \rfloor} \quad (7.29)$$

$$\vec{\Phi}_{D,v} = \sum_i^{N_D-1} \beta_{vi} \phi_{D,i} \quad (7.30)$$

$$\mathcal{F}_{uv} = \omega_{C,u} + \sum_i^{N_D-1} \beta_{vi} \omega_{D,i} \quad (7.31)$$

$$u \in \{0, 1, \dots, N_C - 1\}$$

$$v \in \{0, 1, \dots, 2^{N_D-1}\}$$

To help visualise this, Equation (7.32) shows the frequency matrix \mathcal{F} for the demodulation trees seen in Figure 7.5, with frequencies taken relative to the reference frequency ω_0 :

$$\mathcal{F} = \begin{pmatrix} -\omega_m - \Omega & -\omega_m + \Omega & \omega_m - \Omega & \omega_m + \Omega \\ -2\omega_m - \Omega & -2\omega_m + \Omega & -\Omega & \Omega \\ -\Omega & \Omega & 2\omega_m - \Omega & 2\omega_m + \Omega \end{pmatrix}. \quad (7.32)$$

Each row lists the frequencies each carrier field is mixed with. As in Figure 7.5, we see that some frequencies—those at $\pm\Omega$ —are mixed with multiple carrier fields. The contribution of the unique frequencies in \mathcal{F} to the noise PSD is given by

$$S_I^U(\mathcal{F}_{ij}) = \frac{\hbar(\omega_0 + \mathcal{F}_{ij})}{4^{N_D}} |a_{C,i}|^2, \quad (7.33)$$

where $a_{C,i}$ is the complex amplitude of the i th carrier. The set of N_s indices $\{(n, d)\}$ of the frequencies that appear multiple times in \mathcal{F} must be grouped together as $\{(n_k, d_k)\}$, where $k \in \{0, 1, \dots, N_s - 1\}$. The contribution of the shared frequencies

in \mathcal{F} to the noise PSD is then given by

$$S_I^S(\mathcal{F}_{\{(n,d)\}}) = \frac{\hbar(\omega_0 + \mathcal{F}_{ij})}{4^{N_D}} \left| \sum_{k=0}^{N_s-1} a_{C,n_k}^* e^{i\vec{\Phi}_{D,d_k}} \right|^2, \quad (7.34)$$

where the carriers are added coherently. The total noise PSD is then the sum of these terms for a particular \mathcal{F} where the set of unique frequency indices is $\{(m_k, e_k)\}$ for $k \in \{0, 1, \dots, 2^{N_D} - N_s - 1\}$,

$$S_I^T(\mathcal{F}) = S_I^S(\mathcal{F}_{\{(n,d)\}}) + S_I^U(\mathcal{F}_{\{(m,e)\}}). \quad (7.35)$$

For pure vacuum noise, this is all that is needed to calculate the quantum noise PSD. When squeezing is involved, however, we need to propagate the quantum noise fields through our model to the photodiode.

7.5 Propagation of noise sidebands in FINESSE

For the purpose of modelling interferometers in FINESSE, the semi-classical approach discussed in Section 7.2 is extremely useful, as these stochastic quantum noise sidebands will propagate through a linear system via the same transfer function as any signal sidebands [27]. In practical terms, this means that the matrix \mathbf{M}_s used to calculate the signal response of any model in FINESSE can be reused for the quantum noise simulation, with a minor change which will be discussed in Section 7.7. This slightly modified matrix shall be referred to as the quantum matrix \mathbf{M}_q .

There is a large caveat, however. The inclusion of all input fields in one right-hand side vector for the signal simulation relies on them being coherent, meaning that fields can be linearly combined. For the quantum noise simulation, this is not the case. Each source of noise is incoherent with all others, and so if multiple noise sources were present at once, any couplings between them would have to be performed

incoherently. The incoherent mixing of fields is not a linear process, and so cannot be represented in the interferometer matrix.

The simplest solution to this problem would be to only include one quantum noise source at a time in the RHS vector and repeatedly solve the interferometer matrix for each one individually. This is less than ideal from a performance standpoint, however, due to the large number of quantum noise sources present; quantum noise is introduced incoherently at each open port or lossy component, at each frequency modelled, and for each higher-order mode. For a realistic interferometer model, this can amount to thousands of quantum noise sources, and so treating each one individually can be prohibitively expensive. An alternative solution is therefore highly desirable.

7.6 The MISO method

Compared to the large number of quantum noise sources in a typical interferometer model, there are normally very few points at which we want to measure this noise—often just at the interferometer output. As previously mentioned, we cannot find the quantum noise at all outputs by solving the usual interferometer quantum matrix \mathbf{M}_q once for all quantum noise inputs present, as they are not coherent. We also wish to avoid having to solve \mathbf{M}_q repeatedly for each noise source, due to their large number in a typical model. Instead, FINESSE uses the multiple-input single-output (MISO) method, described in [80]. Here, Brown derives a method by which we can calculate the contribution of each noise input to a single output, in a way which is efficient to compute with a modern sparse matrix solver. This is given by

$$\mathbf{V}_o \vec{s} = \mathbf{M}_q^{-1} \mathbf{V}_i \mathbf{M}_q^{-1 \dagger} \vec{s}, \quad (7.36)$$

where $\mathbf{V}_{\{o,i\}}$ are respectively the output and input variance-covariance matrices of the noise sidebands, and \vec{s} is known as the selection vector. This selection vector is output-dependent, and serves to select out only those elements of \mathbf{V}_o which contribute to the associated output. Specifically, for each quantum noise sideband that will be detected by a demodulated photodiode, the relevant index in \vec{s} contains the amplitude of the carrier field corresponding to that sideband. The rest of \vec{s} is filled with zeros. The total noise PSD for a photodiode is then ideally just

$$S_I^T = \frac{\vec{s}^\dagger \mathbf{V}_o \vec{s}}{4^{N_D}}. \quad (7.37)$$

This is not quite correct due to the exclusion of some of the necessary quantum noise fields from the simulation, for reasons described in Section 7.4. As discussed there, however, these missing frequencies can be assumed to only contain vacuum noise, and so their contributions can be added back in manually. We therefore add to Equation (7.37) the contributions from Equations (7.33) and (7.34) for the unique indices $\{(m', e')\}$ and shared indices $\{(n', d')\}$ in \mathcal{F} that do not appear in \mathbf{M}_q (and thus \mathbf{V}_o) [80],

$$S_I^T = \frac{\vec{s}^\dagger \mathbf{V}_o \vec{s}}{4^{N_D}} + S_I^S(\mathcal{F}_{\{(n', d')\}}) + S_I^U(\mathcal{F}_{\{(m', e')\}}). \quad (7.38)$$

7.7 The quantum matrix M_q

Ideally, the matrix describing how quantum noise sidebands propagate through the model, \mathbf{M}_q , would be exactly the same as that for signal sidebands, \mathbf{M}_s . A complication arises when modulators are involved, however. Figure 7.5 shows how the fields and beatings in an example setup with a modulator could look. A single-order modulator will couple any fields which differ in frequency by the modulation frequency ω_m . In this case, we see that the modulator will couple sidebands from around $\omega_0 \pm 2\omega_m$ with those around $\omega_0 \pm \omega_m$. For the purposes of the signal simulation,

this is not an issue; there are no signal sidebands at these outermost frequencies. In the quantum noise simulation however, this presents a problem. There is a background of vacuum noise at all frequencies, and so in general we would have to include these frequencies in the simulation. Doing so would greatly slow down the simulation however. For a modulator producing sidebands of order N , we would have to introduce $8N$ extra frequencies, compared to $4N$ for the signal simulation i.e. the number of fields would almost double.

In practice, this can be solved when only vacuum noise is present at the input to a modulator. Consider a field at a single frequency; a modulator will incoherently couple equal amounts of vacuum noise out of and into this field, effectively leaving it unchanged. Therefore, if only pure vacuum noise is present at a modulator, it can be correctly modelled by ‘removing’ it (in practice, setting its coupling matrix to the identity matrix). This is nearly always valid when modelling gravitational wave detectors, in which modulators often feature as part of the input optics only and do not interact with squeezed fields. The modulator is therefore modelled in this way in both the older and new versions of FINESSE.

7.8 Generalisation to other sources of noise

Previous versions of FINESSE provided the ability to calculate quantum noise as so far described. Much of the same logic applies equally well to other classical sources of incoherent noise, such as mirror thermal motion and laser frequency noise. Therefore, one of the additions I made when implementing noise calculations in FINESSE 3 was the ability to model these classical noise sources.

To see how this can be done, it is worth revisiting some of the maths we have seen so far and considering which parts are specific to quantum noise, and which are

more general. Starting from Equations (7.17) and (7.18), we had

$$E(t) = e^{i\omega_0 t} \left[a_0 + \int_{-B}^B q(\omega_0 + \Omega) e^{i\Omega t} d\Omega \right] \quad (7.17)$$

where

$$q(\omega) = n_a(\omega) + in_\phi(\omega), \quad (7.18)$$

with $n_{a,\phi}(\omega)$ being normally distributed probability density functions with mean $\mu_{a,\phi} = 0$ and variance $\sigma_{a,\phi}^2$. This form for $q(\omega)$ is not quantum noise specific; it describes any continuum of noise sidebands, with an arbitrary frequency dependence. For this case of a single carrier, we then reached a description for the noise PSD at some sideband frequency Ω in Equation (7.26),

$$S_I(\Omega)\delta(\Omega - \Omega') = 2P_0 \left(\langle q_+ q_+^* \rangle + \langle q_-^* q_- \rangle \right). \quad (7.26)$$

The difference here in the case of classical noise is that we do not know the covariance, as there is no default vacuum noise case; it always depends on the source of the noise and how it has propagated to our detector. The MISO method took this propagation into account in Equation (7.36),

$$\mathbf{V}_o \vec{s} = \mathbf{M}_q^{-1} \mathbf{V}_i \mathbf{M}_q^{-1\dagger} \vec{s}. \quad (7.36)$$

Here, there were two parts specific to quantum noise: the input variance-covariance matrix \mathbf{V}_i ; and \mathbf{M}_q , which differed from the signal matrix \mathbf{M}_s only by the effective removal of modulators to correctly preserve vacuum noise. For classical noises, \mathbf{V}_i can simply be replaced by the correct input variance-covariance matrix containing the noises that the user wishes to model. The modulator, however, is problematic—we cannot perform the same trick of removing it, as that relied on the noise being ‘white’ i.e. frequency-independent.

Another similar problem appears when we try to calculate the classical noise PSD via Equation (7.38),

$$S_I^T = \frac{\vec{s}^\dagger \mathbf{V}_o \vec{s}}{4^{N_D}} + S_I^S(\mathcal{F}_{\{(n',d')\}}) + S_I^U(\mathcal{F}_{\{(m',e')\}}). \quad (7.38)$$

The extra terms containing the noise contributions from frequencies not modelled, $S_I^S(\mathcal{F}_{\{(n',d')\}})$ and $S_I^U(\mathcal{F}_{\{(m',e')\}})$, are calculated in Equations (7.33) and (7.34) under the assumption that only vacuum noise will be present at these frequencies. Again, this is not the case for some arbitrary user-defined classical noise.

In implementing classical noise calculations in FINESSE 3, I therefore made the assumption that all classical noise sources fall off with increasing frequency, such that the generated sidebands are insignificant at the frequencies any carrier fields are separated by. This assumption is valid for many noises, such as laser frequency noise, which falls off like $1/f$, and seismic motion of mirrors, which couples to the laser beam through a suspension chain whose mechanical susceptibility often falls off like $1/f^2$ or more. It may not be true for all noises, however, and the user must be aware of this condition. Future work in this area should involve investigating alternatives to this assumption.

With this assumption made, the problems involved with modelling classical noise are alleviated. Looking once again at Figure 7.5, we see that the modulator can behave as normal, as we assume that any classical noise sidebands around $\omega_0 \pm 2\omega_m$ are insignificant. The normal signal matrix \mathbf{M}_s can therefore be used. We can also see that these unmodelled frequencies do not need to be accounted for when performing photodiode demodulation, as the contribution from any classical noise sidebands at these frequencies is insignificant. FINESSE 3 therefore computes classical




Figure 7.6: Example layout injecting multiple classical noises into one model simultaneously. Here, independent thermal noise is injected at each mirror along with the same seismic noise. This results in semi-coherent sidebands being generated at each mirror.

noise PSDs with a simplified version of the quantum noise calculations:

$$\mathbf{V}_o \vec{s} = \mathbf{M}_s^{-1} \mathbf{V}_i \mathbf{M}_s^{-1\dagger} \vec{s}, \quad (7.39)$$

$$S_I^T = \frac{\vec{s}^\dagger \mathbf{V}_o \vec{s}}{4^{N_D}}. \quad (7.40)$$

The addition of electrical and mechanical components to FINESSE 3 as described in Sections 5.2.2 and 8.3.2 also allows some optimisation here. In contrast to quantum noise, classical noises are not inherently optical effects; they start as electrical or mechanical noises that ultimately couple to the laser beam through e.g. the motion of a mirror, changing field strength of an electro-optic modulator etc. Many classical noises can enter the system at the same node in the simulation. For example, both thermal motion of a mirror and vibration caused by seismic motion will couple to the laser beam via the mirror's surface. If classical noise calculations were implemented in previous versions of FINESSE, injecting multiple noises at the same node in this way would either mean we have to take the incoherent sum of them prior to injection, at which point they become inseparable, or we have to perform a separate simulation for each type of noise. In FINESSE 3, however, these noises can easily be added in the matrix by some 'mixer' component; a simple electrical or mechanical component that adds multiple inputs together and provides an output. This is shown diagrammatically in Figure 7.6, which also shows how correlated noises can be

modelled by propagating the same noise source to multiple mechanical or electrical nodes. In this way, every classical noise can be included in one variance-covariance matrix, and we only have to solve Equation (7.39) once for each output.

Often, however, we want to investigate how only one type of noise, or some subset of the noises present, affects our sensitivity. This could be done by only inputting the noises we wish to model, but this quickly becomes cumbersome. For example, if we wanted to plot a noise budget for an interferometer, we would have to re-run the simulation once for each noise source, and again for the total noise. To this end, another addition I made to the noise calculations in FINESSE 3 is the ability for users to specify which noise sources they wish a given detector to observe. This is achieved via a slight modification of the MISO method, giving

$$\mathbf{V}_o \vec{s} = \mathbf{M}_s^{-1} \mathbf{V}_i \vec{s}_{\text{source}} \mathbf{M}_s^{-1\dagger} \vec{s}, \quad (7.41)$$

where \vec{s}_{source} is known as the source selection vector. The form of \vec{s}_{source} is simple; it contains ones for all of the input nodes that we wish to detect noise arriving from, and zero elsewhere. Different combinations of classical noises can then be detected in one simulation by including multiple detectors, each with different source selection vectors. This also works for the quantum noise calculation, allowing the user to see how subsets of the quantum noises present propagate to the output. Note that this is largely a user convenience feature, rather than a performance one; Equation (7.41) must still be solved for each detector present, and therefore for each combination of noises detected.

7.9 Summary

FINESSE is a frequency domain interferometer simulation tool, which I have used throughout this thesis to model quantum noise in gravitational wave detectors. I was

responsible for the reimplementations of the quantum noise calculations in the latest version, FINESSE 3. These were based on the methods used in the previous version, and I added tests and checks to confirm that the FINESSE 3 reimplementations correctly replicate the existing capabilities and match analytical expressions. The MISO method [80] is used to perform quantum noise simulations quickly and efficiently, by computing the propagation through a model of all sources of quantum noise to a single output simultaneously. I then extended these calculations to include classical sources of noise, and to allow the selection of arbitrary subsets of the noises present on a per-detector basis. These extensions allow entire noise budgets to be produced in a single simulation run, increasing the power and flexibility of FINESSE.

In this chapter, I have given a mathematical description of how the different levels of carrier, signal and noise calculations are performed in FINESSE, providing the necessary background to understand the quantum noise implementation and the additions I made.

Chapter 8

FINESSE 3

In the previous chapter, we took a focused look at the core of FINESSE from a mostly mathematical point of view. In this chapter, I will provide a broader and more practically-minded description of FINESSE 3, the latest version of FINESSE, for which source code [84] and documentation [85] are available. We will see how the interferometer simulation scene looked in general before work on FINESSE 3 started, and the motivations behind undertaking a complete rewrite of the FINESSE code. The architecture of FINESSE 3 will be discussed, along with the reasoning behind some key design choices. Finally, we will see how work on FINESSE 3 has opened the doors to more advanced modelling techniques in the future, potentially allowing faster and more complete simulations.

I was part of the core development team for FINESSE 3, and work on it accounted for roughly half the time I spent doing this PhD over the past four years. The core team consists of 5 people, with a number of other contributors. To give a sense of scale for the project, the FINESSE 3 git repository contains over 3500 commits, with the source amounting to around 62,000 lines of code (including documentation strings), 21,000 lines of tests, and 13,000 lines of manually written documentation. Starting from scratch in 2017, FINESSE 3 has since surpassed feature parity with

FINESSE 2. The first alpha release was made in April 2021, and it has already started to be used for serious modelling work, some of which I will describe in this chapter.

My contributions to FINESSE 3 are widespread throughout the code, with some particularly notable highlights. I was responsible for the complete reimplemention of the quantum noise calculations in FINESSE 3, and extended it to be able to model classical noise sources in a similar way, as described in the previous chapter. I also wrote tests to validate the quantum noise and other parts of the code—tests used to validate radiation pressure calculations can be found in Appendix B. I made significant contributions to early versions of the KatScript parser described later in this chapter, which underwent many iterations before reaching its current form. Smaller contributions I made include the syntax highlighters and Jupyter kernel discussed in Section 8.3.4.

8.1 State of the field

A key tradeoff when simulating any physical system is speed versus modelling capability. More accurate and powerful simulations may take too long to run to be practical for complex models, whereas faster and higher-level simulations may be unable to represent certain aspects of a system. For this reason, multiple tools to simulate gravitational wave interferometers have been written, covering a wide range of use cases. The general-purpose tools can be broadly split up into two categories:

- Time domain models simulate the evolution of a system over time. This allows them to be extremely powerful, as they can model dynamic behaviour and non-linear effects directly. They can also be relatively conceptually simple, as they do not have to make many of the assumptions required by other methods. The price, however, is speed. The wide range of length/time scales present in a single interferometer greatly reduce the efficiency of time domain models, and

they have seen comparatively little investigation in recent years. An example is LIGO End to End [86].

- Frequency domain models assume the system is in a steady state. This allows many optimisations, making them much faster than time domain simulations, and thus more practical for many modelling tasks. They obviously cannot capture dynamic effects, however, and can only model a small number of non-linear effects such as radiation pressure with extra assumptions. Many relevant examples exist, such as FINESSE [26], Optickle [87], OSCAR [88] & SIS [89].

Outside of these categories, other tools also exist to fill specific niches. These include `pygwinc` [9], a higher-level tool to plot noise budgets for various interferometer configurations, and `OPTOCAD` [77], a ray-tracing tool for detailed investigation of Gaussian beams in optical setups.

While a couple of these tools were used peripherally, all of the simulation results presented in this thesis came from FINESSE. During the course of these simulations, limitations of the previous version were run up against. Although it would have been possible to extend the existing code base, it was decided that the time was right to perform a full rewrite of FINESSE. To understand why, it is useful to briefly see the history of FINESSE and its state before the rewrite.

8.2 Motivation for redesigning FINESSE

FINESSE 2 takes the form of a single executable, called `kat`, with source code written in C. Work on FINESSE began over 20 years ago, motivated by the slow speed and inflexibility of existing simulations [90]. Since then, its popularity and feature list have steadily grown, making it one of the most powerful and flexible tools for modelling gravitational wave interferometers available today. Over this time, however, it has

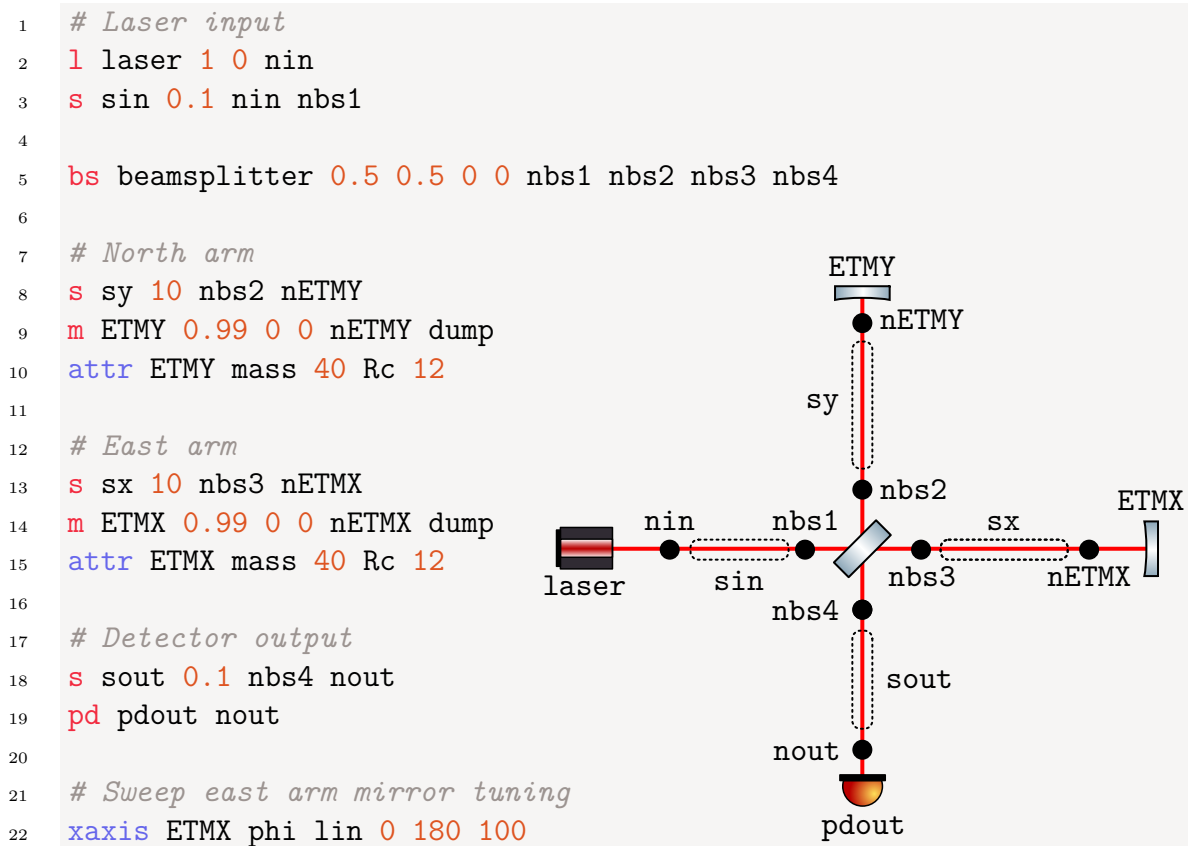
also started to show its age, from both a development and user interface standpoint.

The choice of C as the source language for FINESSE is suitable from a performance point of view. Compared to many other languages, C has very little overhead when performing common operations, and modern compilers are able to produce highly optimised code. For development purposes, however, it comes with some drawbacks. One of the main ones is familiarity; anecdotally, many people in the field of astrophysics are not particularly familiar with C, instead preferring to work in higher level languages, particularly Python. This presents a barrier to entry for anyone wishing to modify the FINESSE source code, reducing the number of willing contributors.

Another issue is caused simply by the age of the codebase, but somewhat exasperated by C. While it is possible to write object-oriented or modular code in C, the language is not particularly accommodating in doing so, and care must be taken to ensure that separate components are kept separate, and that boundaries are well-defined. This is done to some extent in FINESSE 2, but as more and more additions have been made over the years the code has become increasingly brittle and complicated. As a result, while many of the additions made in the rewrite of FINESSE could have been made to the existing codebase, the desire to do so was low among existing developers.

8.2.1 Using FINESSE 2

Traditionally, the input to FINESSE consisted of a single ‘kat file’, a textual description of the optical setup to be modelled and the simulation to run on that model. An example layout and kat file are shown in Listing 8.1. Fully understanding this file is not necessary for now, and I will only give a short outline—more detail can be found in the FINESSE 2 manual [90]. Here, lines 1–19 define the model, and line 22 the simulation to be run. Each line starting with a red keyword defines a component or



Listing 8.1: Kat file and corresponding diagram for a simple Michelson setup with curved, suspended mirrors. Lines starting with a red keyword define components and the spaces between them. These are connected via abstract ‘nodes’. In the last line, the tuning of the east arm mirror is swept from 0–180°.

space with a given name, numerical attributes, and the names of any ‘nodes’ the component is attached to. Connections between components are defined by attaching them to the same node. The order of parameters and the number of nodes is fixed for each component type. For example, the syntax for defining a laser is given by

```
l name power frequency [phase] node
```

where square brackets denote an optional parameter. This terseness has proven popular with those well acquainted with the syntax, however it can be confusing and opaque for those new to FINESSE.

The fixed nature of component definitions leads to one of the quirks of FINESSE 2 syntax, namely the `attr` command, from which the evolution of FINESSE is visible. Certain components, such as mirrors, have enough attributes that listing them all on one line without labels would be uncomfortable to read. For this and other reasons, as more attributes were added over time, they were added to a separate `attr` command which takes named parameters to apply to a component. Users must therefore know both the normal syntax definition of a component, as well as any extra parameters available to be set via `attr`, which is less than ideal. Other such historical quirks and inconsistencies in the FINESSE 2 kat file format also exist, leading to a slight incoherence in the syntax overall.

The single input file nature of FINESSE also came to be a hindrance over time. Often, we wish to perform multiple simulations on a single model. This is not possible with a single kat file; either multiple files must be made, or the same file changed inbetween simulations. Neither choice is desirable when done manually, and so multiple wrappers for FINESSE have been developed to make complex simulations much easier. The most popular of these by far is PYKAT [91].

8.2.2 PYKAT

PYKAT [91] is a set of tools around and interface to FINESSE, written in Python [92]. At its core, it parses kat files into an object-oriented representation. Simulations are run by regenerating a kat file from this representation, feeding it to the kat binary, and storing the results. While sounding simple, this middle step greatly increases the power and flexibility of FINESSE. The whole Python ecosystem becomes available to the user, for purposes such as calculating model parameters and plotting results. The tasks of running complex or multiple simulations on a model are simplified, as the user can programmatically set component properties, make copies of models, and call different simulation routines. This flexibility and ease of use has made PYKAT the de-facto standard way to interact with FINESSE. Indeed, all of the FINESSE 2 results presented in this thesis were generated and plotted with PYKAT and Matplotlib [93], a Python plotting library.

An example use of PYKAT can be seen in Listing 8.2, showing how a model can be manipulated from Python, and multiple simulations can be run on the same model. A few key parts of PYKAT’s design are visible here. The kat file is parsed into a `kat` object, which contains all of the components defined in the file as separate objects. Each component object then contains all of that objects attributes, which can be modified directly as seen on line 37. Also noteworthy is the overall structure of the code. A single kat file contains the initial state of the model we wish to investigate in a compact form. This file is parsed to give the ‘base’ `kat` object, from which we create multiple copies via `base.deepcopy()`. Separate simulations are then run on these copies, to avoid changing the original model. This flow of base-deepcopy-edit-run is ubiquitous in PYKAT usage.

While PYKAT goes a long way to improving the usability of FINESSE, its separation from the main `kat` binary causes a number of small issues. New components and features added to FINESSE also need to be added to PYKAT to be fully usable,

```

1  import pykat
2
3  base = pykat.finesse.kat()
4  base.parse("""
5  # Laser input
6  l laser 1 0 nin
7  s sin 0.1 nin nbs1
8
9  bs beamsplitter 0.5 0.5 0 0 nbs1 nbs2 nbs3 nbs4
10
11 # North arm
12 s sy 10 nbs2 nETMY
13 m ETMY 0.99 0 0 nETMY dump
14 attr ETMY mass 40 Rc 12
15
16 # East arm
17 s sx 10 nbs3 nETMX
18 m ETMX 0.99 0 0 nETMX dump
19 attr ETMX mass 40 Rc 12
20
21 # Detector output
22 s sout 0.1 nbs4 nout
23 pd pdout nout
24 """)
25
26 # Make a copy of our model to run a simulation on
27 kat = base.deepcopy()
28 kat.parse("xaxis ETMY phi lin 0 180 100")
29
30 # Call Finesse to generate the output
31 out = kat.run()
32 out.plot()
33
34 # Change a property and run a different simulation
35 # on a copy of the same model
36 kat = base.deepcopy()
37 kat.ETMY.R = 0.98
38 kat.parse("xaxis beamsplitter phi lin 0 180 100")
39 out = kat.run()
40 out.plot()

```

Listing 8.2: Example usage of PYKAT to run multiple simulations on the file from Listing 8.1.

increasing the maintenance burden. The `kat` binary can be distributed separately from `PYKAT`, and although modern package distribution tools such as Anaconda [94] reduce the problem, version discrepancies and file location issues still arise in practice. Ideally, there would be only one codebase for developers to maintain and users to install. This desire to unify `PYKAT` and `FINESSE` was one of the motivating factors in the development of `FINESSE 3`. I joined the `PYKAT` team towards the end of its feature completion, and some of my work on it contributed to the early state of `FINESSE 3`.

8.3 FINESSE 3 design

`FINESSE 3` is a complete rewrite of `FINESSE` and `PYKAT` in Python. The key aims were to take the lessons learned from these to create a tool that is as modular and extensible as possible, and to add some extra functionality needed for more complex simulations such as seen in Chapter 5. The opportunity was also taken to clean up the various quirks and pieces of cruft that have developed as a result of `FINESSE`'s long history.

8.3.1 Choice of language

A few motivating factors have already been given for choosing Python as the language to write `FINESSE 3` in; the existing familiarity of astrophysicists with Python, and the large Python ecosystem which can be leveraged for both development and simulation purposes. Other factors include the layout of packages in Python, which is well suited for writing modular code with a clear structure.

The main drawback of Python is its performance; Python is much slower than C and other compiled languages, with relatively large overheads on most operations such as function calls and accessing class members. A large portion of this overhead

is due to type checking as a result of Python’s dynamic typing. Thankfully, Python has good interoperability with C, and it is fairly straightforward to call C code from Python. For this reason, many Python packages and tools make use of C internally, especially for core mathematical routines. A notable example is NumPy [95], which describes itself as ‘the fundamental package for scientific computing with Python’. In the development of FINESSE 3, it quickly became clear that at least the core matrix solving code would have to be compiled in some form. Rather than directly writing C code, however, we decided to use Cython [76].

Cython is a Python to C compiler, that allows users to benefit from the ease of writing Python code and the speed of C. While able to compile normal Python code, Cython is most effective when some extra type annotations are added, informing the compiler about the types of variables and functions. This allows it to remove many of the type checks normally performed by Python, as well as other sources of overhead. The benefit can be clearly seen from Listings 8.3 and 8.4, where many lines of generated C are avoided by adding type annotations. Also visible is the similarity of the type annotated code to the unannotated code, though more complex examples necessitate greater changes. Importantly, the type annotated function in Listing 8.4 can be called from regular Python code in exactly the same way as any other function, though performance will be better if making the call from Cythonised code. This allows us to easily mix Cython and Python code in the FINESSE 3 codebase. New components and features can be written in regular Python, lowering the barrier to entry for new contributors, and performance critical sections later translated into Cython by developers more familiar with it. In this way, FINESSE 3 has almost the best of both worlds from Python and C.

```

1 def add(x, y):
2     return x + y
3     # __Pyx_XDECREF(__pyx_r);
4     # __pyx_t_1 = PyNumber_Add(__pyx_v_x, __pyx_v_y);
5     # if (unlikely(!__pyx_t_1)) __PYX_ERR(0, 2, __pyx_L1_error)
6     # __Pyx_GOTREF(__pyx_t_1);
7     # __pyx_r = __pyx_t_1;
8     # __pyx_t_1 = 0;
9     # goto __pyx_L0;
10
11 add(1, 2)
12 # __Pyx_GetModuleGlobalName(__pyx_t_1, __pyx_n_s_add);
13 # if (unlikely(!__pyx_t_1)) __PYX_ERR(0, 4, __pyx_L1_error)
14 # __Pyx_GOTREF(__pyx_t_1);
15 # __pyx_t_2 = __Pyx_PyObject_Call(__pyx_t_1, __pyx_tuple__3, NULL);
16 # if (unlikely(!__pyx_t_2)) __PYX_ERR(0, 4, __pyx_L1_error)
17 # __Pyx_GOTREF(__pyx_t_2);
18 # __Pyx_DECREF(__pyx_t_1); __pyx_t_1 = 0;
19 # __Pyx_DECREF(__pyx_t_2); __pyx_t_2 = 0;

```

Listing 8.3: Example of running Cython on unannotated code. The comments below lines 2 & 11 show the generated C code for each. Note that a single line of Python expands to multiple lines of C, mostly made up of error checking and reference counting.

```

1 cpdef int add(int x, int y):
2     return x + y
3     # __pyx_r = (__pyx_v_x + __pyx_v_y);
4     # goto __pyx_L0;
5
6 add(1, 2)
7 # (void)(__pyx_f_5test2_add(1, 2));

```

Listing 8.4: Example of running Cython on fully type annotated code. The comments below lines 2 & 6 show the generated C code for each. Compared to Listing 8.3, the reference counting and error checking has been eliminated. In more complex setups, this can lead to multiple order of magnitude drops in execution time.

8.3.2 Structure

Much of the high-level structure of FINESSE 3 is directly inspired by that of PYKAT. At the centre is the `Model` object, analogous to the `kat` object from PYKAT. Components and detectors are created and added to the `Model`, either via parsed `KatScript` (the FINESSE 3 equivalent to `kat` files) or directly from Python. Simulations are then performed on the model, and outputs calculated. In this section, I will provide a more detailed view by considering the various steps involved in performing a simulation in order. Along the way, we will see many of the improvements made relative to FINESSE 2.

Parsing `KatScript`

As in previous versions, the majority of simulations begin with a description of the setup to be modelled in a bespoke language. While performing a complete rewrite of FINESSE, the opportunity was also taken to make major changes to the input file format. String handling is particularly cumbersome in C, and standard features in Python meant that a much more advanced parser could be written while keeping complexity reasonable.

The large existing userbase for FINESSE could not be ignored however; users familiar with previous versions (myself included) would rather not have to completely relearn how to write input files for FINESSE to use the new version. To this end, and due to the impracticality of designing a new language overnight, FINESSE 3 started with and still includes a ‘legacy’ parser. This makes a best-effort attempt to parse the old style syntax, allowing most older `kat` files to be used with FINESSE 3 with little modification. The newer syntax then branched off from this legacy parser, meaning the initial form of `KatScript` was very similar to that present in older `kat` files. Over the course of FINESSE 3 development, this has slowly evolved into its current form. Listing 8.5 shows a direct translation of the file from Listing 8.1, and while recognisably similar, demonstrates some of the larger changes made.

```

1  # Laser input
2  l laser P=1
3  s sin laser.p1 beamsplitter.p1 L=0.1
4
5  bs beamsplitter R=0.5 T=0.5
6
7  # North arm
8  s sy beamsplitter.p2 ETMY.p1 L=10
9  m ETMY R=0.99 T=0 Rc=12
10 free_mass ETMY_sus ETMY.mech mass=40
11
12 # East arm
13 s sx beamsplitter.p3 ETMX.p1 L=10
14 m ETMX R=0.99 T=0 Rc=12
15 free_mass ETMX_sus ETMX.mech mass=40
16
17 # Detector output
18 s sout beamsplitter.p4 output.p1 L=0.1
19 nothing output
20 pd pdout sout.p2.o
21
22 # Sweep east arm mirror tuning
23 xaxis(ETMX.phi, lin, 0, 180, 100)

```

Listing 8.5: FINESSE 3 KatScript translation of the file shown in Listing 8.1. Long form names such as `space` are available for each component, but short form names as seen in previous versions of FINESSE are often used for common components.

<pre> 1 m m1 0.9 0.1 0 n1 n2 2 s cav 1 n2 n3 3 m m2 0.45 0.1 0 n3 n4 4 5 set R m1 R 6 func R2 = \$R / 2 7 put m2 R \$R2 </pre>	<pre> 1 m m1 R=0.9 L=0 2 s cav m1.p2 m2.p1 L=1 3 m m2 R=&m1.R/2 L=0 </pre>
--	---

Listing 8.6: Comparison of defining one parameter relative to another in FINESSE 2 (left) and FINESSE 3 (right). FINESSE 3 allows mathematical expressions and references to other parameters directly in parameter definitions, greatly simplifying the task.

A key change is the removal of the explicitly named abstract nodes through which components connect. Instead, components can own ‘ports’, which themselves own nodes. For example, a mirror has three ports; `p1` and `p2`, which represent the two sides of the optical surface, and `mech`, which represents any mechanical couplings the mirror may have. An optical port then has two nodes, `i` & `o`, representing the two directions light can be travelling at that port. Spaces and detectors then connect directly to a component’s ports and nodes respectively. This removes the need to give names to all the nodes in the model, which can be a source of confusion for larger files.

Other visible changes include the addition of keyword arguments alongside positional ones for component definitions, in a style similar to that of Python, and a more function-like definition of the simulation to be run. Another large improvement can be seen in Listing 8.6, which compares the methods used in FINESSE 2 and FINESSE 3 to define one component parameter relative to another, such that changes in the original parameter will automatically be reflected in the derived one. In FINESSE 2, this was achieved via multiple commands; `set` to create a variable `R` tracking the value of `m1.R`, `func` to calculate some value based on this variable, and `put` to use this calculation for the value of `m2.R`. FINESSE 3 by contrast allows mathematical expressions directly in parameter definitions, and introduces the `&` operator to reference component parameters, simplifying the code considerably.

More improvements have been made to KatScript than mentioned here, and together they make it a less cumbersome and much more expressive language than before. This both eases the task of developing and maintaining KatScript files, and reduces the learning burden for new users.

Components, detectors and workspaces

The next step is to build the model from the parsed KatScript. Components and detectors, collectively called elements, are superficially similar to those found in PYKAT, comprising of a set of parameters which can be modified by the user. Components also contain ports and nodes, which are connected together via spaces in the model to form a directed graph. Detectors act as probes, and do not form part of the graph, instead just observing a particular node or set of nodes. These elements then are largely just informational, only describing the relevant object. Many of the various mathematical routines associated with an element are instead implemented as part of a separate *workspace* object.

A workspace is an element-type-specific object responsible for any data structures and calculations an element needs only during a simulation. For example, during the standard matrix simulation as described in the previous chapter, a workspace could hold information about which elements of the interferometer matrix its parent component needs to fill, and calculate the coupling coefficients. By separating this information into its own object, the higher-level element object is kept simple, and can be reused for other simulations by swapping out the workspace. This separation also provides a natural place for the Cython-Python boundary; the workspace needs to perform multiple calculations and interact with the core matrix filling code, so should be written in Cython, whereas the element itself can be written in Python for simplicity.

One of the major additions to FINESSE 3 can also be found here. Previously,

FINESSE could only model optical components, and mechanics only as they applied to these components. In FINESSE 3 however, both mechanical and electrical components are also available. I described this in Section 5.2.2, where the addition of mechanical components simplified the task of modelling the optomechanical filter. This is also apparent from a comparison of Listings 8.1 and 8.5. In FINESSE 2, `mass` is an attribute of a mirror, whereas in FINESSE 3 a `free_mass` component is connected to a mirror's `mech` port, as shown before in Figure 5.4. Crucially, the `free_mass` can be swapped out for any other suspension or series of mechanics without modifying the mirror, as was done for the optomechanical filter. Electrical components are a completely new addition to FINESSE 3, and allow more advanced modelling control systems, though there is still much more work to be done in this area.

Simulations and actions

Now that we have a model with all of the components connected, it's time to do something with it! FINESSE 3 is designed to support a range of simulation types, though for now only one is implemented—the matrix simulation described in Chapter 7. For this, the model must first be 'built': sparse matrices constructed, workspaces created, and locations in the various matrices assigned to each workspace. At this point, the structure of the model is fixed, and elements cannot be added or removed without unbuilding the model, as doing so would require reconstructing the various data structures anyway.

To actually perform modelling tasks with the built model, FINESSE 3 introduces the concept of an *action*. In FINESSE 2, only one class of actions was available. These were the `xaxis`, `x2axis` and `x3axis` commands, which performed 1–3 dimensional parameter sweeps respectively. In addition to these, FINESSE 3 provides many more advanced actions for tasks such as numerical optimisation of some model parameter, measuring multiple transfer functions at once, and so on. It is straightforward for a

```
1 import finesse
2
3 model = finesse.Model()
4 model.parse("""
5     l l1 P=1
6     s s1 l1.p1 m1.p1
7     m m1 R=0.99 L=0
8     series(
9         xaxis(l1.P, lin, 0, 1, 100),
10        xaxis(m1.R, lin, 0, 1, 100),
11    )
12 """)
13
14 out = model.run()
```

Listing 8.7: Minimal example of chaining actions in FINESSE 3. The `series` action takes a list of actions, in this case two `xaxis` actions, and performs them sequentially on the same built model. See Section 8.4 for a more complete example.

user to create their own action for a particular modelling task, and multiple actions can be strung together to perform a series of modelling tasks in one go. We will see a more detailed example of this later in Section 8.4, though Listing 8.7 provides a brief glimpse.

The chaining of actions leads to one of the performance improvements of FINESSE 3 over its predecessors. To perform two separate actions on the same model in PYKAT, as shown in Listing 8.2, two separate calls to FINESSE 2 and an extra copy of the model were required. This results in FINESSE 2 reconstructing all of the sparse matrices and other data structures twice, even though the structure of the model hasn't changed. FINESSE 3, on the other hand, can use the same built simulation for both actions. As matrix construction and other setup is one of the slowest parts of the simulation, FINESSE 3 therefore avoids a large amount of overhead, and can be much faster than FINESSE 2 for long series of actions.

8.3.3 Documentation

Good documentation of software is critical, especially for new users, and FINESSE is no exception. FINESSE 2 provides a few sources of documentation: the manual [90] provides comprehensive coverage of the mathematical background and example kat files, the syntax reference [96] gives a handy overview of the available commands, and modelling tutorials using FINESSE are available from [97].

FINESSE 3 provides all its documentation in the main project source repository, with the pages built in HTML format available online [85]. The documentation is built with Sphinx [98], which allows mixing automatically generated documentation with manually created pages. The FINESSE 3 API pages, which document each of the classes and methods available, are automatically generated from descriptive comments in the source code. This allows the API documentation to be automatically updated whenever new features are added. Other documentation, such as mathematical background, examples, and tutorials, are manually written.

By automatically keeping the documentation up to date, and providing all of it in a single place, the FINESSE 3 documentation aims to make learning about and using FINESSE easier than ever before.

8.3.4 External tooling

In addition to good documentation, other quality of life improvements increase people's willingness to use a tool. As part of the development of FINESSE 3, I wrote a few such improvements, namely syntax highlighters and a kernel for Jupyter Notebooks [99].

Syntax highlighting

Syntax highlighting is omnipresent in the world of programming, making code easier to read, write and understand, by providing hints about its meaning. How syntax highlighting is achieved depends on the environment, as different environments use different tools. As FINESSE uses its own language, KatScript, I had to write four separate highlighters to cover all the places that FINESSE is used by the core development team. These places are:

- Pygments [100] — this is a Python-based syntax highlighter, used for both the syntax highlighting seen throughout this chapter and in the FINESSE 3 documentation. The KatScript Pygments highlighter is part of the core FINESSE 3 package, and leverages the parser to function, making it the most powerful of the highlighters.
- Jupyter Notebooks — these are a web browser based development and presentation environment used commonly throughout the interferometer modelling community, and described in more detail in the next section. They use a JavaScript tool to perform highlighting, and so I wrote a JavaScript extension [101] using regular expressions (regex) to parse the KatScript code.
- Visual Studio Code [102] — this is a popular editor used by many. It also uses a JavaScript syntax highlighter, though a different one to Jupyter Notebooks, necessitating the creation of another extension [103]. This extension again uses regex to parse KatScript code.
- Vim [104] — self described as ‘the ubiquitous editor’. It uses its own language for syntax highlighting, meaning another (still regex based) extension was written [105].

Together, these cover the majority of places KatScript code is likely to be seen, and present a consistent, user-friendly syntax highlighting experience.

Jupyter Notebooks

Jupyter Notebooks [99] are a combined development and presentation environment used throughout the interferometer modelling community and elsewhere. They are popular for their ability to capture an entire workflow, including code, prose and figures, and present that work to others. An example can be seen in Figure 8.1, showing a modified excerpt from a FINESSE workshop I helped run at IUCAA in Pune, India. For this workshop, myself and others prepared teaching material, in the form of Jupyter Notebooks, to guide students through programming tasks. Other workshops run by the interferometer simulation community also make extensive use of Notebooks. Notebooks are also used in parts of the FINESSE 3 documentation for examples and tutorials, and I also perform most of my simulation work in notebooks.

Jupyter Notebooks are not just limited to running Python, however. Each notebook runs a ‘kernel’; a backend responsible for running the code in the notebook and returning results. In Figure 8.1, the current kernel is visible in the top-right, stating that Python 3 is running. In some situations, such as when giving tutorials, it would be nice to be able to directly run KatScript without needing the usual Python boilerplate code around it. I therefore wrote a kernel for KatScript [106], of which an example is shown in Figure 8.2. While still somewhat experimental, this kernel may help new users learn how to model interferometers with FINESSE, especially those unfamiliar with Python, by removing the need for extra Python wrapper code.

8.4 Using FINESSE 3: AdVirgo+ modelling

When modelling a full interferometer, there is some setup that first needs to be done. This includes making sure that all of the optical components are at the correct microscopic tunings, the correct amount of light power is present at the output, the error signals used for control are appropriate and so on. Various setup scripts have

03_Introduction.ipynb
Python 3 (ipykernel)

Simple Finesse Example

Here we construct a simple Finesse model of a laser beam shining on a mirror. In this simulation we just vary the input laser power and see how much is reflected and transmitted. To measure the laser power we put two photodiodes on reflection and transmission at the mirror.

```
[5]: model = finesse.Model()

model.parse("""
#kat
l laser P=1
s s1 laser.p1 m1.p1
m m1 R=0.9 L=0

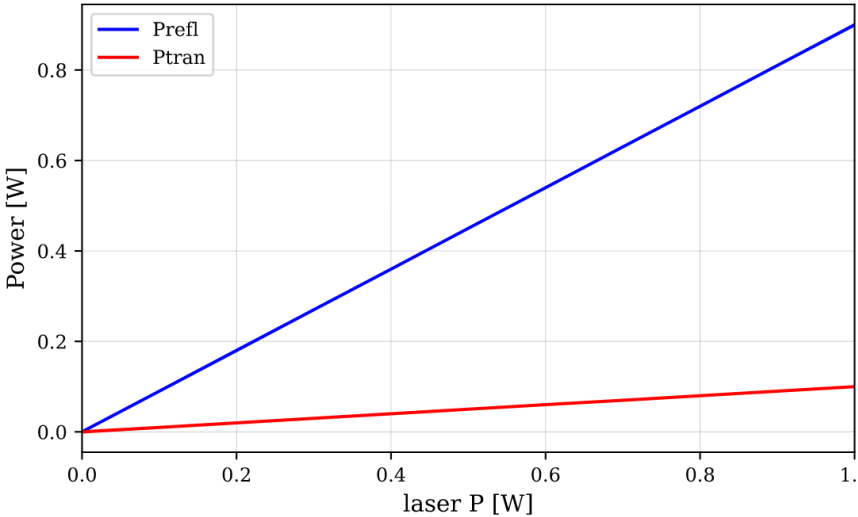
pd Prefl m1.p1.o
pd Ptran m1.p2.o

xaxis(laser.P, lin, 0, 1, 10)
""")

out = model.run()
```

The data can be plotted in multiple ways. If you quickly want to see the results, we suggest using the builtin plotting command for the output:

```
[6]: # Plotting the results of an output can be done using the builtin command
fig = out.plot()
```



This will produce a plot using all the detectors enabled in the model. It should be noted that

Figure 8.1: Example of a Jupyter Notebook, presenting prose, code, and results together. Also present is syntax highlighting of KatScript within a Python string, via the use of a special `#kat` comment.




Figure 8.2: KatScript kernel in a Jupyter Notebook, running the same example shown in Figure 8.1. KatScript is run directly, and simulation results shown automatically, removing the need for Python wrapper code.

therefore been produced for different interferometers, to automate the task of getting the model to the correct operating point. An aLIGO script is included as part of PYKAT.

One of the first practical modelling tasks I used FINESSE 3 for was producing a setup script for AdVirgo+, a future upgrade to the AdVirgo interferometer. This was the first use in anger of the action system, and during the process I realised the potential they had to greatly simplify simulations. A few problems were also encountered along the way. For example, I did not write the `series` action, and was therefore not fully familiar with its working. It proved surprisingly confusing in practice to work out whether changes to the model made by one action in a series would be reset before the next action (they won't), largely because `xaxis` *does* reset the parameter it sweeps after completion. Confusion points such as this often only become apparent through usage for a real task, after which they can be clarified and documented.

As actions hadn't seen much real world use at the time, there also weren't very many of them included in FINESSE—mostly the `xaxis` and a few for specific tasks. I therefore went about writing more actions to carry out the various steps in setting up an interferometer. It soon became clear, however, that actions could be treated as a sort of mini language, and by providing a suite of small, general purpose actions, they could be chained together to perform advanced tasks. I thus made an effort to create actions that could be widely applicable, and an excerpt of the AdVirgo+ setup script using some of these is shown in Listing 8.8. The new actions shown are:

- `Change` — This takes a dictionary of parameters and values, and simply sets each parameter to the specified value. If the `relative` option is set to true, as in line 18 of Listing 8.8, the value specifies an amount to add to the current value of the parameter.
- `MakeTransparent` — A shorthand for a group of `Change` actions, this takes

```

1 Series(
2   # Switch off the modulators for pretuning
3   temporary(Change({"eom1.midx": 0, "eom2.midx": 0}))(
4     # 'remove' SR and PR mirrors by making them transparent
5     temporary(MakeTransparent(["SR", "SRAR"]))(
6       temporary(MakeTransparent(["PR", "PRAR"]))(
7         # Maximise arm power
8         Maximize("B7_DC", "NE_z.DC", bounds=[-90, 90]),
9         Maximize("B8_DC", "WE_z.DC", bounds=[-90, 90]),
10        # Minimise dark fringe power
11        Minimize("B1_DC", "MICH.DC", bounds=[-45, 45])
12      ),
13      # Maximise PRC power
14      Maximize("CAR_AMP_BS", "PRCL.DC", bounds=[-90, 90])
15    ),
16    # Maximise SRC power, then offset by 90°
17    Maximize("B1_DC", "SRCL.DC", bounds=[-90, 90]),
18    Change({"SRCL.DC": -90}, relative=True),
19  )
20 ).run(ifo)

```

Listing 8.8: Example usage of actions in FINESSE 3 to ‘pretune’ an interferometer i.e. make sure all mirrors have the correct tuning. Strings refer to components and parameters of the model (`ifo`) to which the `Series` is applied e.g. `"eom1.midx"` refers to `ifo.eom1.midx`. Most actions are small and straightforward, though `temporary` warrants explanation. The first action passed to it is run, then all the following actions are run, and finally any changes made by the first action are undone. For example, in the above code, the `PR` and `PRAR` mirrors are made transparent on line 6, then have their properties restored on line 12. The other actions are explained in the main body text.

a list of mirrors and sets them to be fully transparent.

- `Maximize` — This takes two positional parameters. The first is the name of a detector in the model, and the second is the name of a parameter. The action then numerically optimises the output of the detector to be at a maximum by varying the stated parameter. The `bounds` option specifies minimum and maximum allowed values for the parameter.
- `Minimize` — The opposite of `Maximize`.
- `temporary` — This is not an action, but a function returning an action. The action passed to it is run, then all the actions passed to the returned action are run, and finally any changes made by the first action are undone. For example, in the above code, the `PR` and `PRAR` mirrors are made transparent on line 6, then have their properties restored on line 12.

This is extremely concise and clear compared to older setup scripts in PYKAT, and even with very little familiarity with FINESSE modelling in general, the overall flow should be roughly understandable.

8.5 The future of FINESSE

By designing FINESSE 3 to be modular and extensible, the path has been made clear for future upgrades to FINESSE, and I will discuss a few potential upgrades here. Most of these have already seen some investigation.

8.5.1 Coupling matrix combination

In Chapter 7, we saw that FINESSE builds a large sparse matrix to represent the model, and computes field amplitudes at every point. This was the chosen method

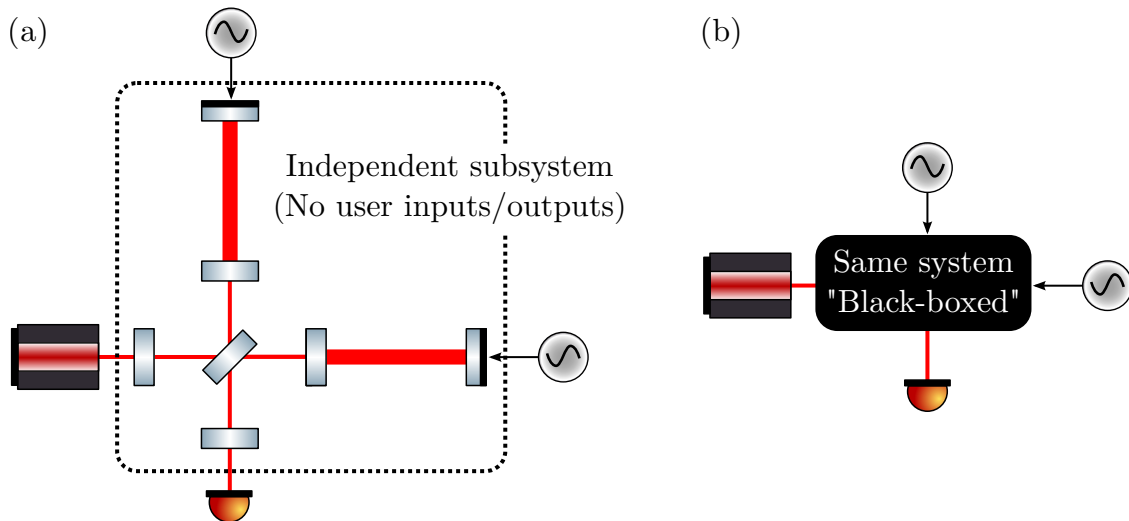


Figure 8.3: Conversion of multiple components (a) into a single ‘black box’ component (b), with a single dense coupling matrix. Only those nodes at which an input or output to the model is present are required.

due to its flexibility and the relative ease of constructing such a matrix. In the majority of cases, however, we only need to know field amplitudes at points where inputs and detectors are connected to the system. By representing the nodes in the simulation as a graph, as FINESSE 3 does, and breaking the graph at those that require input or output, it is possible to identify groups of components that do not need to be treated separately. The individual coupling matrices of these components can then be combined into a single dense matrix, reducing the size of the sparse matrix fed to the matrix solver, and thus reducing the number of calculations required. This is depicted graphically in Figure 8.3. If this ‘black-boxing’ can be done smartly and reliably, it could greatly speed up simulations, allowing more complicated setups to be modelled in a reasonable time frame.

8.5.2 Time domain simulation

While FINESSE 3 is a frequency domain tool, much of the structure is equally applicable to time domain simulations as well, for example KatScript, the `Model` object, and components and detectors. The only parts of FINESSE 3 that are

fully frequency domain specific are the matrix-based simulation and the existing workspace classes. These are designed to be easily switchable, and so there is interest in developing time domain counterparts within FINESSE, especially due to the difficulty of using and extending current time domain simulations such as End to End [86].

By being able to perform both time and frequency domain simulations on the same model, FINESSE 3 would be an extremely useful tool for an even wider range of simulation tasks than currently possible. There is also the potential for unique optimisations. As mentioned in Section 8.1, the wide range of length and time scales in an interferometer can greatly reduce the speed of time domain simulations, as the same temporal resolution is used for the whole model. One idea to get around this is to approximate ‘small’ (i.e. very short dynamic timescale) or uninteresting subsystems as having instantaneous response, and only using time domain simulations where necessary. FINESSE could then calculate the instantaneous response of these subsystems via a frequency domain simulation as usual, and use the resulting transfer functions in the time domain simulation. This could potentially make much larger and more complex time domain simulations possible.

8.5.3 Polarising optics

Currently, no version of FINESSE can model polarised light fields and polarising optics. These are not necessary to model current Michelson interferometers, but ideas for future designs such as certain Sagnac-based velocimeters [107–109] require their use. The addition of polarising optics to FINESSE 3 would be a widespread change, as all components would need to be updated with the correct coupling calculations for each polarisation. Investigations into implementing polarisation in FINESSE 3 are currently ongoing.

8.6 Summary

Numerical modelling tools are widely used in the interferometer simulation community, to investigate both new layouts and behaviours of existing setups. As more advanced techniques and models are proposed, the limitations of existing tools are often run up against. Extensions and new tools are therefore required to progress.

FINESSE 3 is a complete rewrite of one of the most flexible and powerful tools currently in use. Over the past four years, my colleagues and I have developed FINESSE 3 to surpass its predecessors, in speed, flexibility, capability and ease of use. This represents a serious and concerted effort to create the best interferometer modelling tool available, culminating in the first alpha release which was completed in April 2021. Since this release, development has not slowed. In this chapter, I have shown why a complete rewrite of FINESSE was performed, and highlighted some of the major additions made to FINESSE 3. I have also shown how FINESSE 3 has already been successfully used for modelling tasks, and how the new features greatly simplify and speed up the process. Finally, we have seen how FINESSE 3 has been designed to allow even more features to be added in future, to keep its position as one of the best tools of its class for years to come.

Chapter 9

Summary and conclusions

Quantum noise is a result of the inherent uncertainty in the amplitude and phase of the laser beam used in interferometric gravitational wave detectors. For both current and future detectors, quantum noise is a sensitivity limiting source of noise, especially at the high end of the gravitational wave signal frequency spectrum. The reduction of quantum noise is therefore a key area of research and development. In this thesis, I have described my work exploring methods to combat quantum noise, focusing mainly on applications for future detectors such as ET. Within this area, I have worked on a wide range of tasks. Some were more conservative, such as the correction of an overlooked parameter in the ET conceptual design, and the investigation of several alternative designs with a mind to feasibility and cost. Others have been more novel or advanced, like the study of a completely new detection scheme for future interferometers, and the optomechanical filter. Finally, I contributed to the design of a prototype interferometer currently under construction. This has constantly involved the use of the numerical modelling tool FINESSE, for which the development of a successor also constituted significant part of my work.

When investigating a specific part of something as complex as a modern gravitational wave interferometer, we frequently make a number of assumptions and

simplifications. These help focus our efforts, and make the task of designing such an instrument manageable. As new developments are made, however, we must ensure that our original assumptions remain valid. For ET, an assumption had previously been made that the length of the SRC played an insignificant role in the quantum noise limited sensitivity. This was valid for current detectors, such as aLIGO with a 40 m SRC. ET's originally proposed SRC length was 300 m, however, and in Chapter 2 I showed that such a length significantly affected the quantum noise performance of the detector. I then presented options for mitigating the negative effect on the performance. Finally, I explored how the sensitivity was affected by a proposed reduction of the length of the filter cavities used for frequency dependent squeezing at ET. This work was published in: *Philip Jones et al. 'Implications of the quantum noise target for the Einstein Telescope infrastructure design', Phys. Rev. D 101 (8 Apr. 2020) [1]*.

Chapter 3 continued the investigation into filter cavities for frequency dependent squeezing at ET, this time looking at the use of cavities in the interferometer output path. These are not a particularly new idea, but have seen comparatively little investigation due to their extreme sensitivity to lossy optics. In the past it was frequently assumed that achievable losses were too great to make these designs viable. I showed that for parameters expected to be achievable at the completion of ET in 10 years, designs including an output filter are both viable and can be a worthwhile upgrade. I also explored more options for the lengths of the filter cavities in ET, carrying on from my previous work. This work is being prepared for publication, under the title *'Studies of input and output filter cavities for Einstein Telescope'*.

The next two chapters shifted to consider more novel quantum noise reduction techniques. In Chapter 4, I described a recently proposed detector design known as the two-carrier scheme, which seeks to avoid the quantum noise penalty experienced by heterodyne readout while retaining and improving on the benefits. I detailed the simulation work I carried out to verify the performance of the design, and

derived some of the particular requirements on interferometer parameters. This work contributed to the publication of *Teng Zhang et al. ‘Two-Carrier Scheme: Evading the 3 dB Quantum Penalty of Heterodyne Readout in Gravitational-Wave Detectors’*, *Phys. Rev. Lett.* 126 (22 June 2021) [2].

In Chapter 5, I worked on further developing a novel idea known as the opto-mechanical filter, which seeks to broaden the bandwidth of gravitational wave interferometers similarly to a white light cavity. I performed the first numerical simulations that verified the key behaviour exhibited by this filter, namely linear negative dispersion. This sets the stage for more detailed study in future. Studies conducted elsewhere have also shown that the filter can be made self-stabilising, improving the feasibility of implementing such a device in an interferometer.

My simulation work has not focused solely on quantum noise in future detectors, however. ET requires many novel technologies, which must first be tested on a smaller scale. ETpathfinder is a prototype interferometer designed for this testing. Currently under construction, the finer details of its optical layout are still being developed. I undertook the task of finding requirements for laser frequency noise and input motion cleaner noise in ETpathfinder, and provided rough estimates along with their scalings with asymmetries in the detector. These findings have informed decisions regarding the parameters of ETpathfinder.

All of this work was made possible by the interferometer simulation tool FINESSE. Due to its power, flexibility, and freely accessible open source nature, FINESSE is one of the most widely used tools throughout the interferometer simulation community. I have made crucial contributions to the latest version of FINESSE, including performing the complete reimplementations of the quantum noise calculations used throughout this thesis. Chapter 7 provided a mathematical description of these calculations, as well as the extensions I made to allow other sources of noise to be modelled, and Chapter 8 demonstrated the power of FINESSE 3 for simulation work. Having already

surpassed feature parity with its predecessor, the first alpha release of FINESSE 3 was made earlier this year. Development is still well underway, and FINESSE 3 is poised to see even more widespread success than ever before, driving the field of interferometer simulation forward.

Bibliography

- [1] Philip Jones et al. ‘Implications of the quantum noise target for the Einstein Telescope infrastructure design’. In: *Phys. Rev. D* 101 (8 Apr. 2020), p. 082002. DOI: 10.1103/PhysRevD.101.082002.
- [2] Teng Zhang et al. ‘Two-Carrier Scheme: Evading the 3 dB Quantum Penalty of Heterodyne Readout in Gravitational-Wave Detectors’. In: *Phys. Rev. Lett.* 126 (22 June 2021), p. 221301. DOI: 10.1103/PhysRevLett.126.221301.
- [3] B. P. Abbott et al. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Phys. Rev. Lett.* 116 (6 Feb. 2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102.
- [4] B. S. Sathyaprakash and Bernard F. Schutz. ‘Physics, Astrophysics and Cosmology with Gravitational Waves’. In: *Living Reviews in Relativity* 12.1 (Mar. 2009). ISSN: 1433-8351. DOI: 10.12942/lrr-2009-2.
- [5] J. Aasi et al. ‘Advanced LIGO’. In: *Class. Quant. Grav.* 32 (2015), p. 074001. DOI: 10.1088/0264-9381/32/7/074001. arXiv: 1411.4547 [gr-qc].
- [6] F. Acernese et al. ‘Advanced Virgo: a second-generation interferometric gravitational wave detector’. In: *Class. Quant. Grav.* 32.2 (2015), p. 024001. DOI: 10.1088/0264-9381/32/2/024001. arXiv: 1408.3978 [gr-qc].
- [7] T. Akutsu et al. *Overview of KAGRA: Detector design and construction history*. 2020. arXiv: 2005.05574 [physics.ins-det].
- [8] K L Dooley et al. ‘GEO 600 and the GEO-HF upgrade program: successes and challenges’. In: *Classical and Quantum Gravity* 33.7 (Mar. 2016), p. 075009. DOI: 10.1088/0264-9381/33/7/075009.
- [9] *Python Gravitational Wave Interferometer Noise Calculator*. URL: <https://git.ligo.org/gwinc/pygwinc> (visited on 06/09/2021).
- [10] ET Steering Committee Editorial Team. *Einstein Telescope design report update 2020*. Einstein Telescope Collaboration, Sept. 2020.
- [11] F Matichard et al. ‘Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance’. In: *Classical and Quantum Gravity* 32.18 (Aug. 2015), p. 185003. ISSN: 1361-6382. DOI: 10.1088/0264-9381/32/18/185003.
- [12] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford Master Series in Physics. Oxford University Press, 2007. ISBN: 978-0-19-857074-5.

- [13] Carlton M. Caves. ‘Quantum-mechanical noise in an interferometer’. In: *Phys. Rev. D* 23 (8 Apr. 1981), pp. 1693–1708. DOI: 10.1103/PhysRevD.23.1693.
- [14] E. Oelker et al. ‘Squeezed light for advanced gravitational wave detectors and beyond’. In: *Opt. Express* 22.17 (Aug. 2014), pp. 21106–21121. DOI: 10.1364/OE.22.021106.
- [15] J. Aasi et al. ‘Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light’. In: *Nature Photonics* 7 (July 2013), p. 613.
- [16] Roman Schnabel. ‘Squeezed states of light and their applications in laser interferometers’. In: *Phys. Rept.* 684 (2017), pp. 1–51. DOI: 10.1016/j.physrep.2017.04.001. arXiv: 1611.03986 [quant-ph].
- [17] L.-A. Wu, M. Xiao and H. J. Kimble. ‘Squeezed states of light from an optical parametric oscillator’. In: *Journal of the Optical Society of America B Optical Physics* 4 (Oct. 1987), pp. 1465–1475. DOI: 10.1364/JOSAB.4.001465.
- [18] Christopher Gerry and Peter Knight. *Introductory Quantum Optics*. Cambridge University Press, 2004. DOI: 10.1017/CB09780511791239.
- [19] A. Gatti, T. Corti and E. Brambilla. ‘Squeezing and Einstein-Podolsky-Rosen correlation in the mirrorless optical parametric oscillator’. In: *Phys. Rev. A* 96.1, 013820 (July 2017), p. 013820. DOI: 10.1103/PhysRevA.96.013820. arXiv: 1704.09010 [quant-ph].
- [20] Moritz Mehmet and Henning Vahlbruch. ‘The Squeezed Light Source for the Advanced Virgo Detector in the Observation Run O3’. In: *Galaxies* 8.4 (2020). ISSN: 2075-4434. DOI: 10.3390/galaxies8040079.
- [21] H. J. Kimble et al. ‘Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics’. In: *Phys. Rev. D* 65 (2 Dec. 2001), p. 022002. DOI: 10.1103/PhysRevD.65.022002.
- [22] P. Kwee et al. ‘Decoherence and degradation of squeezed states in quantum filter cavities’. In: *Phys. Rev. D* 90 (6 Sept. 2014), p. 062006. DOI: 10.1103/PhysRevD.90.062006.
- [23] L. McCuller et al. ‘Frequency-Dependent Squeezing for Advanced LIGO’. In: *Phys. Rev. Lett.* 124 (17 Apr. 2020), p. 171102. DOI: 10.1103/PhysRevLett.124.171102.
- [24] Eric Oelker et al. ‘Audio-Band Frequency-Dependent Squeezing for Gravitational-Wave Detectors’. In: *Physical Review Letters* 116.4 (Jan. 2016). ISSN: 1079-7114. DOI: 10.1103/physrevlett.116.041102.
- [25] Yuhang Zhao et al. ‘Frequency-Dependent Squeezed Vacuum Source for Broadband Quantum Noise Reduction in Advanced Gravitational-Wave Detectors’. In: *Physical Review Letters* 124.17 (Apr. 2020). ISSN: 1079-7114. DOI: 10.1103/physrevlett.124.171101.
- [26] Daniel David Brown and Andreas Freise. *FINESSE: frequency domain interferometer simulation software*. <http://www.gwoptics.org/finesse>. 2020.

- [27] Charlotte Bond et al. ‘Interferometer techniques for gravitational-wave detection’. In: *Living Reviews in Relativity* 19.1 (Feb. 2017), p. 3. ISSN: 1433-8351. DOI: 10.1007/s41114-016-0002-8.
- [28] David Reitze et al. *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*. 2019. arXiv: 1907.04833 [astro-ph.IM].
- [29] The ETpathfinder Team. *ETpathfinder Design Report*. Jan. 2020.
- [30] H. Grote et al. ‘First Long-Term Application of Squeezed States of Light in a Gravitational-Wave Observatory’. In: *Phys. Rev. Lett.* 110 (18 May 2013), p. 181101. DOI: 10.1103/PhysRevLett.110.181101.
- [31] M. Tse et al. ‘Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy’. In: *Phys. Rev. Lett.* 123 (23 Dec. 2019), p. 231107. DOI: 10.1103/PhysRevLett.123.231107.
- [32] F. Acernese et al. ‘Increasing the Astrophysical Reach of the Advanced Virgo Detector via the Application of Squeezed Vacuum States of Light’. In: *Phys. Rev. Lett.* 123 (23 Dec. 2019), p. 231108. DOI: 10.1103/PhysRevLett.123.231108.
- [33] M Punturo et al. ‘The third generation of gravitational wave observatories and their science reach’. In: *Classical and Quantum Gravity* 27.8 (2010), p. 084007.
- [34] S Hild et al. ‘Sensitivity studies for third-generation gravitational wave observatories’. In: *Classical and Quantum Gravity* 28.9 (2011), p. 094013.
- [35] The ET Science Team. *Einstein gravitational wave Telescope conceptual design*. European Commission, June 2011.
- [36] The ET Steering Committee. *Design Report Update 2020 for the Einstein Telescope*. European Commission, Nov. 2020.
- [37] Stefan Hild et al. ‘A Xylophone Configuration for a third Generation Gravitational Wave Detector’. In: *Class. Quant. Grav.* 27 (2010), p. 015003. DOI: 10.1088/0264-9381/27/1/015003. arXiv: 0906.2655 [gr-qc].
- [38] S Hild et al. ‘Demonstration and comparison of tuned and detuned signal recycling in a large-scale gravitational wave detector’. In: *Classical and Quantum Gravity* 24.6 (Mar. 2007), pp. 1513–1523. DOI: 10.1088/0264-9381/24/6/009.
- [39] Jun Mizuno. ‘Comparison of optical configurations for laser-interferometric gravitational-wave detectors’. PhD thesis. Hannover U., 1995.
- [40] M. Evans et al. ‘Realistic filter cavities for advanced gravitational wave detectors’. In: *Phys. Rev. D* 88 (2 July 2013), p. 022002. DOI: 10.1103/PhysRevD.88.022002.
- [41] W. Winkler et al. ‘Heating by optical absorption and the performance of interferometric gravitational-wave detectors’. In: *Phys. Rev. A* 44 (11 Dec. 1991), pp. 7022–7036. DOI: 10.1103/PhysRevA.44.7022.
- [42] Yu. Levin. ‘Internal thermal noise in the LIGO test masses: A direct approach’. In: *Phys. Rev. D* 57 (2 Jan. 1998), pp. 659–663. DOI: 10.1103/PhysRevD.57.659.

- [43] A Freise et al. ‘Frequency-domain interferometer simulation with higher-order spatial modes’. In: *Classical and Quantum Gravity* 21.5 (Feb. 2004), S1067–S1074. ISSN: 1361-6382. DOI: 10.1088/0264-9381/21/5/102.
- [44] A C Green et al. ‘The influence of dual-recycling on parametric instabilities at Advanced LIGO’. In: *Classical and Quantum Gravity* 34.20 (2017), p. 205004. DOI: 10.1088/1361-6382/aa8af8.
- [45] Alessandra Buonanno and Yan-bei Chen. ‘Scaling law in signal recycled laser interferometer gravitational-wave detectors’. In: *Phys. Rev. D* 67 (2003), p. 062002. DOI: 10.1103/PhysRevD.67.062002. arXiv: gr-qc/0208048 [gr-qc].
- [46] André Thüring, Harald Lück and Karsten Danzmann. ‘Analysis of a four-mirror-cavity enhanced Michelson interferometer’. In: *Phys. Rev. E* 72 (6 Dec. 2005), p. 066615. DOI: 10.1103/PhysRevE.72.066615.
- [47] Denis Martynov et al. ‘Exploring the sensitivity of gravitational wave detectors to neutron star physics’. In: *Phys. Rev.* (2019). arXiv: 1901.03885 [astro-ph.IM].
- [48] T. Isogai et al. ‘Loss in long-storage-time optical cavities’. In: *Opt. Express* 21 (2013), pp. 30114–30125. DOI: 10.1364/OE.21.030114. arXiv: 1310.1820 [physics.optics].
- [49] P. Kwee et al. ‘Decoherence and degradation of squeezed states in quantum filter cavities’. In: *Phys. Rev. D* 90.6 (2014), p. 062006. DOI: 10.1103/PhysRevD.90.062006. arXiv: 1704.03531 [physics.optics].
- [50] Patricia Purdue and Yan-bei Chen. ‘Practical speed meter designs for QND gravitational wave interferometers’. In: *Phys. Rev. D* 66 (2002), p. 122004. DOI: 10.1103/PhysRevD.66.122004. arXiv: gr-qc/0208049 [gr-qc].
- [51] Jérôme Degallaix et al. ‘Large and extremely low loss: the unique challenges of gravitational wave mirrors’. In: *J. Opt. Soc. Am. A* 36.11 (Nov. 2019), pp. C85–C94. DOI: 10.1364/JOSAA.36.000C85.
- [52] Lisa Barsotti et al. ‘The A+ design curve’. In: *LIGO Document T1800042* (Dec. 2018).
- [53] Sheila Dwyer et al. ‘Gravitational wave detector with cosmological reach’. In: *Phys. Rev. D* 91 (8 Apr. 2015), p. 082001. DOI: 10.1103/PhysRevD.91.082001.
- [54] Stefan L. Danilishin and Farid Ya. Khalili. ‘Quantum Measurement Theory in Gravitational-Wave Detectors’. In: *Living Reviews in Relativity* 15.1 (Apr. 2012), p. 5. ISSN: 1433-8351. DOI: 10.12942/lrr-2012-5.
- [55] Hartmut Grote. ‘Making it Work: Second Generation Interferometry in GEO600!’ PhD thesis. Universität Hannover, 2003.
- [56] Osamu Miyakawa et al. ‘Measurement of optical response of a detuned resonant sideband extraction interferometer’. In: *Phys. Rev. D* 74 (2006), p. 022001. DOI: 10.1103/PhysRevD.74.022001. arXiv: gr-qc/0604078 [gr-qc].

- [57] Teng Zhang et al. ‘Effects of static and dynamic higher-order optical modes in balanced homodyne readout for future gravitational waves detectors’. In: *Phys. Rev. D* 95 (6 Mar. 2017), p. 062001. DOI: 10.1103/PhysRevD.95.062001.
- [58] S. Personick. ‘An Image Band Interpretation of Optical Heterodyne Noise’. In: *Bell System Technical Journal* 50 (1971), pp. 213–216.
- [59] T. M. Niebauer et al. ‘Nonstationary shot noise and its effect on the sensitivity of interferometers’. In: *Phys. Rev. A* 43 (9 May 1991), pp. 5022–5029. DOI: 10.1103/PhysRevA.43.5022.
- [60] Kirk McKenzie et al. ‘Squeezing in the Audio Gravitational-Wave Detection Band’. In: *Physical Review Letters* 93.16 (Oct. 2004). ISSN: 1079-7114. DOI: 10.1103/physrevlett.93.161105.
- [61] B P Abbott et al. ‘LIGO: the Laser Interferometer Gravitational-Wave Observatory’. In: *Reports on Progress in Physics* 72.7 (June 2009), p. 076901. ISSN: 1361-6633. DOI: 10.1088/0034-4885/72/7/076901.
- [62] T Accadia, F Acernese et al. ‘Virgo: a laser interferometer to detect gravitational waves’. In: *Journal of Instrumentation* 7.03 (Mar. 2012), P03012–P03012. DOI: 10.1088/1748-0221/7/03/p03012.
- [63] H. Grote. ‘The GEO 600 status’. In: *Class. Quant. Grav.* 27 (2010). Ed. by Zsuzsa Marka and Szabolcs Marka, p. 084003. DOI: 10.1088/0264-9381/27/8/084003.
- [64] S Hild et al. ‘DC-readout of a signal-recycled gravitational wave detector’. In: *Classical and Quantum Gravity* 26.5 (Feb. 2009), p. 055012. ISSN: 1361-6382. DOI: 10.1088/0264-9381/26/5/055012.
- [65] Sascha Husa. ‘Michele Maggiore: Gravitational waves. Volume 1: theory and experiments’. In: *General Relativity and Gravitation* 41.7 (July 2009), pp. 1667–1669. ISSN: 1572-9532. DOI: 10.1007/s10714-009-0762-5.
- [66] G. S. Pati et al. ‘Demonstration of a Tunable-Bandwidth White-Light Interferometer Using Anomalous Dispersion in Atomic Vapor’. In: *Phys. Rev. Lett.* 99 (13 Sept. 2007), p. 133601. DOI: 10.1103/PhysRevLett.99.133601.
- [67] H. N. Yum et al. ‘Demonstration of White Light Cavity Effect Using Stimulated Brillouin Scattering in a Fiber Loop’. In: *Journal of Lightwave Technology* 31 (Dec. 2013), pp. 3865–3872. DOI: 10.1109/JLT.2013.2288326. arXiv: 1307.5272 [physics.optics].
- [68] Minchuan Zhou and Selim M. Shahriar. ‘Optomechanical resonator as a negative dispersion medium for enhancing the sensitivity bandwidth in a gravitational-wave detector’. In: *Phys. Rev. D* 98.2 (2018), p. 022003. DOI: 10.1103/PhysRevD.98.022003. arXiv: 1610.07255 [gr-qc].
- [69] Yiqiu Ma et al. ‘Quantum noise of a white-light cavity using a double-pumped gain medium’. In: *Phys. Rev. A* 92 (2 Aug. 2015), p. 023807. DOI: 10.1103/PhysRevA.92.023807.
- [70] Haixing Miao et al. ‘Enhancing the Bandwidth of Gravitational-Wave Detectors with Unstable Optomechanical Filters’. In: *Physical Review Letters* 115.21 (Nov. 2015). ISSN: 1079-7114. DOI: 10.1103/physrevlett.115.211104.

- [71] Joe Bentley et al. ‘Converting the signal-recycling cavity into an unstable optomechanical filter to enhance the detection bandwidth of gravitational-wave detectors’. In: *Physical Review D* 99.10 (May 2019). ISSN: 2470-0029. DOI: 10.1103/physrevd.99.102001.
- [72] Xiang Li et al. ‘Enhancing interferometer sensitivity without sacrificing bandwidth and stability: Beyond single-mode and resolved-sideband approximation’. In: *Physical Review D* 103.12 (June 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.103.122001.
- [73] Philip Jones et al. *Negative Dispersion with Optomechanical Systems in Interferometer Simulations*. Mar. 2018. LIGO DCC: LIGO-G1800388.
- [74] S. Weis et al. ‘Optomechanically Induced Transparency’. In: *Science* 330 (Dec. 2010), p. 1520. DOI: 10.1126/science.1195596. arXiv: 1007.0565 [quant-ph].
- [75] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1441412697.
- [76] S. Behnel et al. ‘Cython: The Best of Both Worlds’. In: *Computing in Science Engineering* 13.2 (Mar. 2011), pp. 31–39. ISSN: 1521-9615. DOI: 10.1109/MCSE.2010.118.
- [77] Roland Schilling. *OPTOCAD: Tracing Gaussian TEM₀₀ beams through an optical set-up*. URL: <http://www2.mpg.de/~ros/optocad.html> (visited on 06/09/2021).
- [78] Anthony E. Siegman. *Lasers*. University Science Books, 1986. ISBN: 9780935702118.
- [79] Craig Cahillane, Georgia Mansell and Daniel Sigg. *Laser Frequency Noise in Next Generation Gravitational-Wave Detectors*. 2021. arXiv: 2107.14349 [physics.ins-det].
- [80] D. D. Brown. ‘Interaction of light and mirrors: Advanced techniques for modelling future gravitational wave detectors’. PhD thesis. University of Birmingham, 2015.
- [81] Jan Harms, Paul Cochrane and Andreas Freise. ‘Quantum-noise power spectrum of fields with discrete classical components’. In: *Phys. Rev.* A76 (2007), p. 023803. DOI: 10.1103/PhysRevA.76.023803. arXiv: quant-ph/0703119 [QUANT-PH].
- [82] Timothy A. Davis and Ekanathan Palamadai Natarajan. ‘Algorithm 907: KLU, A Direct Sparse Solver for Circuit Simulation Problems’. In: *ACM Trans. Math. Softw.* 37.3 (Sept. 2010). ISSN: 0098-3500. DOI: 10.1145/1824801.1824814.
- [83] Carlton M. Caves and Bonny L. Schumaker. ‘New formalism for two-photon quantum optics. I. Quadrature phases and squeezed states’. In: *Phys. Rev. A* 31 (5 May 1985), pp. 3068–3092. DOI: 10.1103/PhysRevA.31.3068.
- [84] *FINESSE 3 Source Repository*. URL: <https://git.ligo.org/finesse/finesse3> (visited on 10/09/2021).
- [85] *FINESSE 3 Documentation*. URL: <https://finesse.docs.ligo.org/finesse3> (visited on 10/09/2021).

- [86] *LIGO End to End Simulation software*. Aug. 2015. URL: <https://labcit.ligo.caltech.edu/~e2e/> (visited on 06/09/2021).
- [87] *Optickle*. URL: <https://github.com/Optickle/Optickle> (visited on 06/09/2021).
- [88] Jérôme Degallaix. ‘OSCAR a Matlab based optical FFT code’. In: *Journal of Physics: Conference Series* 228 (May 2010), p. 012021. DOI: 10.1088/1742-6596/228/1/012021.
- [89] Hiro Yamamoto and Richard Day. *SIS: Static Interferometer Simulation*. URL: <https://git.ligo.org/IF0sim/SIS> (visited on 06/09/2021).
- [90] A. Freise, D. Brown and C. Bond. *Finesse 2.0 Manual*. Available from <http://www.gwoptics.org/finesse/>. May 2014.
- [91] Daniel D. Brown et al. ‘Pykat: Python package for modelling precision optical interferometers’. In: *SoftwareX* 12 (2020), p. 100613. ISSN: 2352-7110. DOI: 10.1016/j.softx.2020.100613.
- [92] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1441412697.
- [93] J. D. Hunter. ‘Matplotlib: A 2D graphics environment’. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.
- [94] *Anaconda Software Distribution*. Version Vers. 2-2.4.0. 2020.
- [95] Charles R. Harris et al. ‘Array programming with NumPy’. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- [96] *FINESSE Syntax Reference*. 30th Apr. 2016. URL: <http://www.gwoptics.org/finesse/reference> (visited on 10/09/2021).
- [97] *Learn Laser Interferometry*. URL: <http://www.gwoptics.org/learn> (visited on 10/09/2021).
- [98] *Sphinx Python Documentation Generator*. URL: <https://www.sphinx-doc.org> (visited on 10/09/2021).
- [99] Thomas Kluyver et al. ‘Jupyter Notebooks—a publishing format for reproducible computational workflows’. In: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*. Ed. by Fernando Loizides and Birgit Schmidt. IOS Press, 2016, pp. 87–90.
- [100] *Pygments Python syntax highlighter*. URL: <https://pygments.org> (visited on 10/09/2021).
- [101] *KatScript JupyterLab highlighter*. URL: <https://git.ligo.org/finesse/katscript-jupyterlab> (visited on 10/09/2021).
- [102] *Visual Studio Code*. URL: <https://code.visualstudio.com> (visited on 10/09/2021).
- [103] *KatScript VSCode highlighter*. URL: <https://git.ligo.org/finesse/katscript-vscode> (visited on 10/09/2021).
- [104] *Vim - the ubiquitous text editor*. URL: <https://www.vim.org> (visited on 10/09/2021).

-
- [105] *KatScript Vim highlighter*. URL: <https://git.ligo.org/finesse/katscript-vim> (visited on 10/09/2021).
- [106] *KatScript Jupyter Kernel*. URL: <https://git.ligo.org/finesse/katscript-jupyter-kernel> (visited on 10/09/2021).
- [107] Yan-bei Chen. ‘Sagnac interferometer as a speed meter type, quantum non-demolition gravitational wave detector’. In: *Phys. Rev. D* 67 (2003), p. 122004. DOI: 10.1103/PhysRevD.67.122004. arXiv: gr-qc/0208051 [gr-qc].
- [108] N. V. Voronchev, Sh. L. Danilishin and F. Ya. Khalili. ‘A Sagnac interferometer as a gravitational-wave third-generation detector’. In: *Moscow University Physics Bulletin* 69.6 (Nov. 2014), pp. 519–528. ISSN: 1934-8460. DOI: 10.3103/S0027134914060198.
- [109] S H Huttner et al. ‘Candidates for a possible third-generation gravitational wave detector: comparison of ring-Sagnac and sloshing-Sagnac speedmeter interferometers’. In: *Classical and Quantum Gravity* 34.2 (2017), p. 024001.

Appendix A

Optomechanical filter derivation

In this appendix, I will go through the derivation of the negative dispersion relation stated in Equation (5.6). This derivation was carried out in parallel with a colleague, however we use different conventions for minus signs on the Fourier transform and the phase relations at a mirror. I have included this derivation to make it abundantly clear where the minus sign in Equation (5.6) has come from relative to [70].

A.1 Notation

The Fourier transform convention used is

$$F(\omega) = \int_{-\infty}^{\infty} dt e^{-i\omega t} f(t), \quad f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega e^{i\omega t} F(\omega). \quad (\text{A.1})$$

Throughout this appendix, for a given field \hat{a} the notation

$$\vec{a} \equiv \begin{bmatrix} \hat{a}_1 \\ \hat{a}_2 \end{bmatrix}$$

will be used, where \hat{a}_1 and \hat{a}_2 are the amplitude and phase quadrature components respectively. Operators will generally be in the Heisenberg picture, and may be in

the time or frequency domain based on context.

The annihilation operators for photons in the upper and lower sidebands $\pm\Omega$ about the laser frequency ω_0 are defined as

$$\hat{a}_+ \equiv \hat{a}_{\omega_0+\Omega}, \quad \hat{a}_- \equiv \hat{a}_{\omega_0-\Omega}. \quad (\text{A.2})$$

The amplitude and phase quadratures are then defined through

$$\hat{a}_1 = \frac{\hat{a}_+ + \hat{a}_-^\dagger}{\sqrt{2}}, \quad \hat{a}_2 = \frac{\hat{a}_+ - \hat{a}_-^\dagger}{i\sqrt{2}}, \quad (\text{A.3})$$

and conversely,

$$\hat{a}_+ = \frac{\hat{a}_1 + i\hat{a}_2}{\sqrt{2}}, \quad \hat{a}_- = \frac{\hat{a}_1^\dagger + i\hat{a}_2^\dagger}{\sqrt{2}}. \quad (\text{A.4})$$

If we have an arbitrary input-output relation,


$$\vec{\beta} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \vec{\alpha}, \quad (\text{A.5})$$

then the upper output sideband is given by

$$\hat{\beta}_+ = \frac{1}{2} \left[((A + D) + i(C - B))\hat{\alpha}_+ + ((A - D) + i(C + B))\hat{\alpha}_-^\dagger \right] \equiv T_+\hat{\alpha}_+ + T_-^*\hat{\alpha}_-^\dagger, \quad (\text{A.6})$$

and the lower output sideband by

$$\hat{\beta}_- = \frac{1}{2} \left[((A - D)^\dagger + i(C + B)^\dagger)\hat{\alpha}_+^\dagger + ((A + D)^\dagger + i(C - B)^\dagger)\hat{\alpha}_- \right] \equiv T_+^*\hat{\alpha}_+^\dagger + T_-\hat{\alpha}_-. \quad (\text{A.7})$$



A.2 Free Space

The input optical field can be written as

$$\hat{E}_{\text{in}}(t) = \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t) \right) \cos \omega_0 t - \hat{a}_2(t) \sin \omega_0 t \right]. \quad (\text{A.8})$$

Assuming $R = 1$, the output field is simply time-delayed due to travel across the system,

$$\hat{E}_{\text{out}}(t) = \hat{E}_{\text{in}}(t - 2\tau - 2\hat{x}/c),$$

where $\tau = L/c$ is the travel time across the space before the mirror (L is the *tuned* length), and \hat{x} is the displacement of the end mirror. We can write the output field as

$$\begin{aligned} \hat{E}_{\text{out}}(t) &\equiv \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left(\hat{b}_1(t) \cos \omega_0 t - \hat{b}_2(t) \sin \omega_0 t \right) \\ &= \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t - 2\tau - 2\hat{x}/c) \right) \cos \omega_0(t - 2\tau - 2\hat{x}/c) \right. \\ &\quad \left. - \hat{a}_2(t - 2\tau - 2\hat{x}/c) \sin \omega_0(t - 2\tau - 2\hat{x}/c) \right]. \end{aligned}$$

Since L is the tuned length, $\omega_0 L/c = \omega_0 \tau = n\pi$ where n is an integer,

$$\begin{aligned} \hat{E}_{\text{out}}(t) &= \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t - 2\tau - 2\hat{x}/c) \right) \cos \omega_0(t - 2\hat{x}/c) \right. \\ &\quad \left. - \hat{a}_2(t - 2\tau - 2\hat{x}/c) \sin \omega_0(t - 2\hat{x}/c) \right]. \end{aligned}$$

Using some trigonometry, and then assuming that the mirror displacement \hat{x} is much smaller than the macroscopic length L so that the mirror displacement can be dropped from the quadratures,

$$\hat{E}_{\text{out}}(t) \approx \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t-2\tau) \right) \left(\cos\omega_0 t \cos\frac{2\omega_0\hat{x}}{c} + \sin\omega_0 t \sin\frac{2\omega_0\hat{x}}{c} \right) - \hat{a}_2(t-2\tau) \left(\sin\omega_0 t \cos\frac{2\omega_0\hat{x}}{c} - \cos\omega_0 t \sin\frac{2\omega_0\hat{x}}{c} \right) \right].$$

Now assuming $\frac{\omega_0\hat{x}}{c} \ll 1$ and expanding the trig terms to first order,

$$\hat{E}_{\text{out}}(t) \approx \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t-2\tau) \right) \left(\cos\omega_0 t + \sin\omega_0 t \frac{2\omega_0\hat{x}}{c} \right) - \hat{a}_2(t-2\tau) \left(\sin\omega_0 t - \cos\omega_0 t \frac{2\omega_0\hat{x}}{c} \right) \right].$$

Grouping coefficients,

$$\hat{E}_{\text{out}}(t) \approx \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}} \left[\cos\omega_0 t \left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t-2\tau) + \frac{2\omega_0\hat{x}}{c} \hat{a}_2(t-2\tau) \right) - \sin\omega_0 t \left(\hat{a}_2(t-2\tau) - \frac{2\omega_0\hat{x}}{c} \left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t-2\tau) \right) \right) \right].$$

Equating the coefficients of the trig functions we have

$$\begin{aligned} \hat{b}_1(t) &= \hat{a}_1(t-2\tau) - \frac{2\omega_0\hat{x}(t-\tau)}{c} \hat{a}_2(t-2\tau), \\ \hat{b}_2(t) &= \hat{a}_2(t-2\tau) - \frac{2\omega_0\hat{x}(t-\tau)}{c} \left(\sqrt{\frac{2I_0}{\hbar\omega_0}} + \hat{a}_1(t-2\tau) \right) \end{aligned}$$

If we also assume that the quadratures terms are small such that any cross-terms of $\hat{a}_{1,2}$ and $\frac{\omega_0 \hat{x}}{c}$ are negligible,

$$\begin{aligned}\hat{b}_1(t) &= \hat{a}_1(t - 2\tau), \\ \hat{b}_2(t) &= \hat{a}_2(t - 2\tau) - \frac{2\omega_0}{c} \sqrt{\frac{2I_0}{\hbar\omega_0}} \hat{x}(t - \tau).\end{aligned}\tag{A.9}$$

To transform this into the frequency domain, we first take the Fourier transform of eq. (A.9),

$$\begin{aligned}\hat{b}_1(\Omega) &= e^{-2i\Omega\tau} \hat{a}_1(\Omega), \\ \hat{b}_2(\Omega) &= e^{-2i\Omega\tau} \hat{a}_2(\Omega) - \frac{2\omega_0}{c} e^{-i\Omega\tau} \sqrt{\frac{2I_0}{\hbar\omega_0}} \hat{x}(\Omega).\end{aligned}\tag{A.10}$$

Now,

$$m\ddot{x}(t) = \hat{F}_{rp}(t) + \frac{1}{2}mL\ddot{h}(t),\tag{A.11}$$

and we have

$$\hat{F}_{rp}(t) = \frac{2I_0}{c} \left[1 + \sqrt{\frac{2\hbar\omega_0}{I_0}} \hat{a}_1(t - \tau) \right].$$

Putting this into eq. (A.11) and taking the Fourier transform, we find

$$-\Omega^2 m \hat{x}(\Omega) = \frac{2I_0}{c} \sqrt{\frac{2\hbar\omega_0}{I_0}} e^{-i\Omega\tau} \hat{a}_1(\Omega) - \frac{1}{2} m L \Omega^2 \hat{h}(\Omega).$$

If we then put this into eq. (A.10), we obtain

$$\begin{aligned}
\hat{b}_2 &= e^{-2i\Omega\tau}\hat{a}_2(\Omega) + \frac{1}{\Omega^2 m} \frac{2\omega_0}{c} e^{-i\Omega\tau} \sqrt{\frac{2I_0}{\hbar\omega_0}} \left(\frac{2I_0}{c} \sqrt{\frac{2\hbar\omega_0}{I_0}} e^{-i\Omega\tau}\hat{a}_1(\Omega) - \frac{1}{2}mL\Omega^2\hat{h}(\Omega) \right) \\
&= e^{-2i\Omega\tau}\hat{a}_2(\Omega) + \frac{8\omega_0 I_0}{mc^2\Omega^2} e^{-2i\Omega\tau}\hat{a}_1(\Omega) - \sqrt{\frac{2I_0\omega_0 L^2}{\hbar c^2}} e^{-i\Omega\tau}\hat{h}(\Omega) \\
&= e^{-2i\Omega\tau}\hat{a}_2(\Omega) + \kappa e^{-2i\Omega\tau}\hat{a}_1(\Omega) - \sqrt{2\kappa} e^{-i\Omega\tau} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}},
\end{aligned}$$

where


$$\kappa = \frac{8I_0\omega_0}{mc^2\Omega^2}, \quad h_{\text{SQL}} = \sqrt{\frac{8\hbar}{m\Omega^2 L^2}}.$$

We can then write the relation between $\hat{b}(\Omega)$ and $\hat{a}(\Omega)$ cleanly in matrix form,

$$\vec{\hat{b}}(\Omega) = e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \vec{\hat{a}}(\Omega) - e^{-i\Omega\tau} \sqrt{2\kappa} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \frac{h(\Omega)}{h_{\text{SQL}}} \quad (\text{A.12})$$

Ignoring the gravitational wave strain, we can then use the relations from equations (A.6) and (A.7) to obtain

$$\begin{aligned}
\hat{b}_+ &= \frac{e^{-2i\Omega\tau}}{2} ((2 + i\kappa)\hat{a}_+ + i\kappa\hat{a}_-^\dagger) \\
\hat{b}_- &= \frac{e^{2i\Omega\tau}}{2} ((2 + i\kappa)\hat{a}_- + i\kappa\hat{a}_+^\dagger).
\end{aligned}$$



A.3 Tuned Cavity

Relating the fields before (\vec{a}, \vec{b}) and after $(\vec{\alpha}, \vec{\beta})$ the mirror, we have

$$\vec{a} = \sqrt{R}\vec{b} + i\sqrt{T}\vec{\alpha} \quad (\text{A.13})$$

$$\vec{\beta} = i\sqrt{T}\vec{b} + \sqrt{R}\vec{\alpha}. \quad (\text{A.14})$$

Substituting Equation (A.13) into our free space relation from Equation (A.12) gives

$$\vec{b}(\Omega) = e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \left(\sqrt{R}\vec{b}(\Omega) + i\sqrt{T}\vec{\alpha}(\Omega) \right) - \sqrt{2\kappa}e^{-i\Omega\tau} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}}.$$

We can then obtain an expression for $\vec{b}(\Omega)$ in terms of $\vec{\alpha}(\Omega)$,

$$\underbrace{\begin{bmatrix} 1 - \sqrt{R}e^{-2i\Omega\tau} & 0 \\ -\sqrt{R}\kappa e^{-2i\Omega\tau} & 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix}}_M \vec{b}(\Omega) = i\sqrt{T}e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \vec{\alpha} - \sqrt{2\kappa}e^{-i\Omega\tau} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}}.$$

Now we can multiply from the left by \mathbf{M}^{-1} ,

$$\begin{aligned} \mathbf{M}^{-1} &= \frac{1}{(1 - \sqrt{R}e^{-2i\Omega\tau})^2} \begin{bmatrix} 1 - \sqrt{R}e^{-2i\Omega\tau} & 0 \\ \sqrt{R}\kappa e^{-2i\Omega\tau} & 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix}, \\ \therefore \vec{b}(\Omega) &= \frac{i\sqrt{T}e^{-2i\Omega\tau}}{(1 - \sqrt{R}e^{-2i\Omega\tau})^2} \begin{bmatrix} 1 - \sqrt{R}e^{-2i\Omega\tau} & 0 \\ \kappa & 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix} \vec{\alpha} \\ &\quad - \frac{\sqrt{2\kappa}e^{-i\Omega\tau}}{(1 - \sqrt{R}e^{-2i\Omega\tau})^2} \begin{bmatrix} 0 \\ 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}}. \end{aligned}$$

We then substitute in Equation (A.14) to find $\vec{\beta}(\Omega)$ in terms of $\vec{\alpha}(\Omega)$,

$$\begin{aligned} \vec{\beta} &= \left(\frac{-Te^{-2i\Omega\tau}}{(1 - \sqrt{R}e^{-2i\Omega\tau})^2} \begin{bmatrix} 1 - \sqrt{R}e^{-2i\Omega\tau} & 0 \\ \kappa & 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix} + I\sqrt{R} \right) \vec{\alpha} \\ &\quad - \frac{i\sqrt{2\kappa T}e^{-i\Omega\tau}}{(1 - \sqrt{R}e^{-2i\Omega\tau})^2} \begin{bmatrix} 0 \\ 1 - \sqrt{R}e^{-2i\Omega\tau} \end{bmatrix} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}}. \end{aligned} \quad (\text{A.15})$$

First, let us consider the $\hat{\beta}_1$ terms,

$$\hat{\beta}_1 = \left(\frac{-Te^{-2i\Omega\tau}}{1 - \sqrt{R}e^{-2i\Omega\tau}} + \sqrt{R} \right) \hat{\alpha}_1. \quad (\text{A.16})$$

Looking at the term in brackets and assuming a lossless mirror $T + R = 1$, we have

$$\begin{aligned} (\dots) &= \frac{-Te^{-2i\Omega\tau}}{1 - e^{-2i\Omega\tau}\sqrt{1-T}} + \sqrt{1-T} \\ &= \frac{Te^{-2i\Omega\tau} + \sqrt{1-T} - e^{-2i\Omega\tau}(1-T)}{1 - e^{-2i\Omega\tau}\sqrt{1-T}} \\ &= \frac{\sqrt{1-T} - e^{-2i\Omega\tau}}{1 - \sqrt{1-T}e^{-2i\Omega\tau}}. \end{aligned}$$

We then assume

$$T \ll 1 \implies \sqrt{1-T} \approx 1 - \frac{T}{2}$$

$$\Omega\tau \ll 1 \implies e^{-2i\Omega\tau} \approx 1 - 2i\Omega\tau.$$

These are valid assumptions for LIGO, which uses highly reflective, low-loss mirrors, and investigates a frequency of $\Omega \approx 10^3$ Hz with $\tau \approx 10^{-5}$ s. This gives us

$$\begin{aligned} (\dots) &\approx \frac{1 - \frac{T}{2} - 1 + 2i\Omega\tau}{1 - 1 + \underbrace{iT\Omega\tau}_{\approx 0} + \frac{T}{2} + 2i\Omega\tau} \\ &\approx -\frac{\frac{T}{2} - 2i\Omega\tau}{\frac{T}{2} + 2i\Omega\tau}. \end{aligned}$$

Then, defining $\tan \phi = \frac{\Omega}{\gamma}$, with the cavity bandwidth $\gamma = \frac{T}{4\tau} \equiv \frac{cT}{4L}$,

$$\begin{aligned} (\dots) &= -\frac{1 - i \tan \phi}{1 + i \tan \phi} = -\frac{\cos \phi - i \sin \phi}{\cos \phi + i \sin \phi} = -\frac{e^{-i\phi}}{e^{i\phi}} = -e^{-2i\phi}, \\ \therefore \hat{\beta}_1 &= -e^{-2i\phi} \hat{\alpha}_1. \end{aligned} \tag{A.17}$$

Next we shall consider the $\hat{\beta}_2$ terms of Equation (A.15),

$$\begin{aligned} \hat{\beta}_2 &= \frac{-\kappa T e^{-2i\Omega\tau}}{(1 - \sqrt{R} e^{-2i\Omega\tau})^2} \hat{\alpha}_1 \\ &+ \left(\frac{-T e^{-2i\Omega\tau}}{1 - \sqrt{R} e^{-2i\Omega\tau}} + \sqrt{R} \right) \hat{\alpha}_2 \\ &- \frac{i\sqrt{2\kappa T} e^{-i\Omega\tau}}{1 - \sqrt{R} e^{-2i\Omega\tau}} \frac{h}{h_{\text{SQL}}}. \end{aligned}$$

The coefficient of $\hat{\alpha}_2$ is the same as the bracketed term in eq. A.16 and so the same results apply.

For the coefficient of $\hat{\alpha}_1$, we follow the same process as before; assume T &

$\Omega\tau \ll 1$, expand their terms to leading order, and throw away cross terms between the two quantities. This gives

$$\begin{aligned}
 (\dots) &= \frac{-\kappa T e^{-2i\Omega\tau}}{\left(1 - \sqrt{1 - T} e^{-2i\Omega\tau}\right)^2} \\
 &\approx \frac{-\kappa T (1 - 2i\Omega\tau)}{\left(1 - \left(1 - \frac{T}{2}\right) (1 - 2i\Omega\tau)\right)^2} \\
 &\approx \frac{-\kappa T}{\left(\frac{T}{2} + 2i\Omega\tau\right)^2}
 \end{aligned}$$

Next, we multiply the top and bottom by a factor of $T/2 - 2i\Omega\tau$ and rearrange to find

$$\begin{aligned}
 (\dots) &= \frac{-\kappa T \left(\frac{T}{2} - 2i\Omega\tau\right)}{\left(\frac{T}{2} + 2i\Omega\tau\right)^2 \left(\frac{T}{2} - 2i\Omega\tau\right)} \\
 &= \frac{\frac{T}{2} - 2i\Omega\tau}{\frac{T}{2} + 2i\Omega\tau} \left(\frac{-\kappa T}{\frac{T^2}{4} + 4\Omega^2\tau^2}\right) \\
 &= -e^{-2i\phi} \frac{\kappa\gamma}{\tau(\Omega^2 + \gamma^2)} \\
 &= -e^{-2i\phi} \frac{\iota_c\gamma}{\Omega^2(\Omega^2 + \gamma^2)} \\
 &\equiv -e^{-2i\phi} \mathcal{K}
 \end{aligned} \tag{A.18}$$

using the cavity bandwidth $\gamma = \frac{T}{4\tau}$ defined before, and further defining the cavity parameter $\iota_c = \frac{8I_c\omega_0}{mcL}$ after replacing the laser intensity I_0 with the intracavity intensity $I_c = \frac{4I_0}{T}$.

Finally we find the strain coefficient using the same process,

$$\begin{aligned}
(\dots) &= \frac{-i\sqrt{2\kappa T}e^{-i\Omega\tau}}{1 - \sqrt{R}e^{-2i\Omega\tau}} \\
&\approx \frac{-i\sqrt{2\kappa T}}{\frac{T}{2} + 2i\Omega\tau} \\
&= -\sqrt{\left(\frac{\frac{T}{2} - 2i\Omega\tau}{\frac{T}{2} + 2i\Omega\tau}\right) \frac{-2\kappa T}{\left(\frac{T}{2} + 2i\Omega\tau\right)^2}} \\
&= -\sqrt{\left(\frac{1 - i \tan \phi}{1 + i \tan \phi}\right) \frac{-2\kappa \frac{T}{4\tau}}{\frac{T^2}{4\tau} + \Omega^2\tau}} \\
&= -\sqrt{-e^{-2i\phi} \frac{-2\kappa\gamma}{\tau(\gamma^2 + \Omega^2)}} \\
&= -\sqrt{2\mathcal{K}}e^{-i\phi}.
\end{aligned} \tag{A.19}$$


In a more compact notation we have

$$\vec{\beta}(\Omega) = -e^{-2i\phi(\Omega)} \begin{bmatrix} 1 & 0 \\ \mathcal{K}(\Omega) & 1 \end{bmatrix} \vec{\alpha}(\Omega) - e^{-i\phi(\Omega)} \begin{bmatrix} 0 \\ \sqrt{2\mathcal{K}(\Omega)} \end{bmatrix} \frac{\hat{h}(\Omega)}{h_{\text{SQL}}},$$

where $\tan \phi \equiv \Omega/\gamma$, $\mathcal{K} = \frac{\iota_c \gamma}{\Omega^2(\Omega^2 + \gamma^2)}$, $\gamma = \frac{T}{4\tau}$, $\iota_c = \frac{8I_c \omega_0}{mcL}$, and $I_c = \frac{4I_0}{T}$.

Ignoring the GW strain, we can then use the relations from equations (A.6) and (A.7) to obtain

$$\begin{aligned}
\hat{\beta}_+ &= \frac{-e^{-2i\phi(\Omega)}}{2} ((2 + i\mathcal{K})\hat{\alpha}_+ + i\mathcal{K}\hat{\alpha}_-^\dagger) \\
\hat{\beta}_- &= \frac{-e^{2i\phi(\Omega)}}{2} ((2 + i\mathcal{K})\hat{\alpha}_- + i\mathcal{K}\hat{\alpha}_+^\dagger).
\end{aligned}$$



A.4 Detuned Cavity

For an input mirror detuned by phase θ the quadratures are *rotated*,

$$\begin{aligned}\vec{a} &= \sqrt{R}\mathbf{R}_{2\theta}\vec{b} + i\sqrt{T}\mathbf{R}_\theta\vec{\alpha}, \\ \vec{\beta} &= i\sqrt{T}\mathbf{R}_\theta\vec{b} + \sqrt{R}\vec{\alpha},\end{aligned}$$

where

$$\mathbf{R}_\theta = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}. \quad (\text{A.20})$$

Recall from the free-space case the input-output relation,

$$\vec{b} = e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \vec{a} - e^{-i\Omega\tau} \sqrt{2\kappa} \frac{h}{h_{\text{SQL}}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Substituting the relation for \vec{a} ,

$$\vec{b} = e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \left(\sqrt{R}\mathbf{R}_{2\theta}\vec{b} + i\sqrt{T}\mathbf{R}_\theta\vec{\alpha} \right) - e^{-i\Omega\tau} \sqrt{2\kappa} \frac{h}{h_{\text{SQL}}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Then rearranging for \vec{b} ,

$$\left(1 - e^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \sqrt{R}\mathbf{R}_{2\theta} \right) \vec{b} = ie^{-2i\Omega\tau} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \sqrt{T}\mathbf{R}_\theta\vec{\alpha} - e^{-i\Omega\tau} \sqrt{2\kappa} \frac{h}{h_{\text{SQL}}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Then defining the matrix in parentheses as \mathbf{D} ,

$$\mathbf{D} \equiv \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - e^{-2i\Omega\tau} \sqrt{R} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \mathbf{R}_{2\theta},$$

and multiplying from the left by \mathbf{D}^{-1} ,

$$\vec{b} = ie^{-2i\Omega\tau} \sqrt{T} \mathbf{D}^{-1} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \mathbf{R}_{\theta} \vec{\alpha} - e^{-i\Omega\tau} \sqrt{2\kappa} \frac{h}{h_{\text{SQL}}} \mathbf{D}^{-1} \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Substituting this into the expression for $\vec{\beta}$,

$$\vec{\beta} = -Te^{-2i\Omega\tau} \mathbf{R}_{\theta} \mathbf{D}^{-1} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \mathbf{R}_{\theta} \vec{\alpha} - ie^{-i\Omega\tau} \sqrt{2\kappa T} \frac{h}{h_{\text{SQL}}} \mathbf{R}_{\theta} \mathbf{D}^{-1} \begin{bmatrix} 0 \\ 1 \end{bmatrix} + \sqrt{R} \vec{\alpha}. \quad (\text{A.21})$$

To leading order in $\theta \equiv \Delta\tau$ where Δ is the detuning frequency,

$$\mathbf{R}_{\theta=\Delta\tau} \approx \begin{bmatrix} 1 - \Delta^2\tau^2/2 & \Delta\tau \\ -\Delta\tau & 1 - \Delta^2\tau^2/2 \end{bmatrix}.$$

We also assume

$$T \ll 1 \implies \sqrt{R} = \sqrt{1-T} \approx 1 - T/2 - T^2/8,$$

$$\Omega\tau \ll 1 \implies e^{-2i\Omega\tau} \approx 1 - 2i\Omega\tau - 2\Omega^2\tau^2,$$

and use the definition of cavity bandwidth, $T = 4\gamma\tau \ll 1$. Together all of this gives

$$\mathbf{D} \approx \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - (1 - 2i\Omega\tau - 2\Omega^2\tau^2)(1 - 2\gamma\tau - 2\gamma^2\tau^2) \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \begin{bmatrix} 1 - 2\Delta^2\tau^2 & 2\Delta\tau \\ -2\Delta\tau & 1 - 2\Delta^2\tau^2 \end{bmatrix}$$

Keeping only terms to second-order in the small quantities κ , $\Omega\tau$, $\Delta\tau$, and $\gamma\tau$,

we have

$$\mathbf{D} \approx \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - (1 - 2\gamma\tau - 2\gamma^2\tau^2 - 2i\Omega\tau + 4i\Omega\gamma\tau^2 - 2\Omega^2\tau^2) \begin{bmatrix} 1-2\Delta^2\tau^2 & 2\Delta\tau \\ \kappa-2\Delta\tau & 1+2\kappa\Delta\tau-2\Delta^2\tau^2 \end{bmatrix}$$

$$\mathbf{D} \approx \begin{bmatrix} 2\Delta^2\tau^2+2\gamma^2\tau^2+2\Omega^2\tau^2 & -2\Delta\tau+4\gamma\Delta\tau^2+4i\Omega\Delta\tau^2 \\ +2\gamma\tau+2i\Omega\tau-4i\Omega\gamma\tau^2 & \\ -\kappa+2\Delta\tau+2\kappa\gamma\tau & 2\Delta^2\tau^2+2\gamma^2\tau^2+2\Omega^2\tau^2 \\ -4\gamma\Delta\tau^2+2i\kappa\Omega\tau-4i\Omega\Delta\tau^2 & -2\kappa\Delta\tau+2\gamma\tau+2i\Omega\tau-4i\Omega\gamma\tau^2 \end{bmatrix}. \quad (\text{A.22})$$

Noting that $\kappa = \frac{\tau l_c}{\Omega^2}$, then

$$\det \mathbf{D} \approx 4\tau^2 \left[-(\Omega - i\gamma)^2 + \Delta^2 - \frac{\Delta l_c}{2\Omega^2} \right], \quad (\text{A.23})$$

and so we can calculate \mathbf{D}^{-1} ,

$$\mathbf{D}^{-1} \approx \frac{1}{-4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta l_c}{2\Omega^2} \right]} \quad (\text{A.24})$$

$$\cdot \begin{bmatrix} 2\Delta^2\tau^2+2\gamma^2\tau^2+2\Omega^2\tau^2 & 2\Delta\tau-4\gamma\Delta\tau^2+4i\Omega\Delta\tau^2 \\ -2\kappa\Delta\tau+2\gamma\tau+2i\Omega\tau-4i\Omega\gamma\tau^2 & \\ \kappa-2\Delta\tau-2\kappa\gamma\tau & 2\Delta^2\tau^2+2\gamma^2\tau^2+2\Omega^2\tau^2 \\ +4\gamma\Delta\tau^2-2i\kappa\Omega\tau+4i\Omega\Delta\tau^2 & +2\gamma\tau+2i\Omega\tau-4i\Omega\gamma\tau^2 \end{bmatrix} \quad (\text{A.25})$$

Next we need to calculate the coefficient of the $\vec{\alpha}$ term,

$$-T e^{-2i\Omega\tau} \mathbf{R}_\theta \mathbf{D}^{-1} \begin{bmatrix} 1 & 0 \\ \kappa & 1 \end{bmatrix} \mathbf{R}_\theta + \sqrt{R} \mathbf{I} \quad (\text{A.26})$$

Defining the first term as \mathbf{A} and expanding to second order in $\Omega\tau$, $\gamma\tau$, $\Delta\tau$, and κ

(remembering that T is also small) gives

$$\begin{aligned}
\mathbf{A} &\approx -T e^{-2i\Omega\tau} \mathbf{R}_\theta \mathbf{D}^{-1} \begin{bmatrix} 1 & \Delta\tau \\ \kappa - \Delta\tau & 1 + \kappa\Delta\tau \end{bmatrix} \\
\mathbf{A} &\approx \frac{T e^{-2i\Omega\tau}}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \mathbf{R}_\theta \begin{bmatrix} 2\gamma\tau + 2i\Omega\tau & 2\Delta\tau \\ \kappa - 2\Delta\tau & 2\gamma\tau + 2i\Omega\tau \end{bmatrix} \\
\mathbf{A} &\approx \frac{T e^{-2i\Omega\tau}}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \begin{bmatrix} 2\gamma\tau + 2i\Omega\tau & 2\Delta\tau \\ \kappa - 2\Delta\tau & 2\gamma\tau + 2i\Omega\tau \end{bmatrix} \\
\mathbf{A} &\approx \frac{1}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \begin{bmatrix} 8\gamma^2\tau^2 + 8i\Omega\gamma\tau^2 & 8\gamma\Delta\tau^2 \\ -\frac{4\iota_c}{\Omega^2}\gamma\tau^2 - 8\gamma\Delta\tau^2 & 8\gamma^2\tau^2 + 8i\Omega\gamma\tau^2 \end{bmatrix}. \quad (\text{A.27})
\end{aligned}$$

We then turn to the second term in equation (A.26), using $T = 4\gamma\tau$ and expanding to second order as before:

$$\begin{aligned}
\sqrt{R}\mathbf{I} &\approx \begin{bmatrix} 1 - 2\gamma\tau - 2\gamma^2\tau^2 & 0 \\ 0 & 1 - 2\gamma\tau - 2\gamma^2\tau^2 \end{bmatrix} \\
\sqrt{R}\mathbf{I} &\approx \frac{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \begin{bmatrix} 1 - 2\gamma\tau - 2\gamma^2\tau^2 & 0 \\ 0 & 1 - 2\gamma\tau - 2\gamma^2\tau^2 \end{bmatrix} \\
\sqrt{R}\mathbf{I} &\approx \frac{4\Omega^2\tau^2 - 8i\Omega\gamma\tau^2 - 4\gamma^2\tau^2 - 4\Delta^2\tau^2 + \frac{2\Delta\iota_c\tau^2}{\Omega^2}}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (\text{A.28})
\end{aligned}$$

Putting together equations (A.27) & (A.28), and ignoring the gravitational wave

strain, gives us

$$\begin{aligned}
\vec{\beta} &= \frac{1}{4\tau^2 \left[(\Omega - i\gamma)^2 - \Delta^2 + \frac{\Delta\iota_c}{2\Omega^2} \right]} \\
&\cdot \begin{bmatrix} 4\Omega^2\tau^2 + 4\gamma^2\tau^2 - 4\Delta^2\tau^2 + \frac{2\Delta\iota_c\tau^2}{\Omega^2} & 8\gamma\Delta\tau^2 \\ \frac{4\iota_c}{\Omega^2}\gamma\tau^2 - 8\gamma\Delta\tau^2 & 4\Omega^2\tau^2 + 4\gamma^2\tau^2 - 4\Delta^2\tau^2 + \frac{2\Delta\iota_c\tau^2}{\Omega^2} \end{bmatrix} \vec{\alpha} \\
\vec{\beta} &= \frac{1}{\Omega^2 \left[(\Omega - i\gamma)^2 - \Delta^2 \right] + \frac{\Delta\iota_c}{2}} \\
&\cdot \begin{bmatrix} \Omega^2(\Omega^2 + \gamma^2 - \Delta^2) + \frac{\Delta\iota_c}{2} & 2\gamma\Delta\Omega^2 \\ -2\gamma\Delta\Omega^2 + \gamma\iota_c & \Omega^2(\Omega^2 + \gamma^2 - \Delta^2) + \frac{\Delta\iota_c}{2} \end{bmatrix} \vec{\alpha} \tag{A.29}
\end{aligned}$$

Sideband Picture

To obtain the sideband transfer functions, we substitute eq. A.29 into eqs. A.6 & A.7, giving

$$\begin{aligned}
\hat{\beta}_+ &= \frac{1}{2} \frac{(2\Omega^2(\Omega^2 + \gamma^2 - \Delta^2) + \Delta\iota_c - 4i\gamma\Delta\Omega^2 + i\gamma\iota_c) \alpha_+ + i\gamma\iota_c \alpha_+^\dagger}{\Omega^2 \left[(\Omega - i\gamma)^2 - \Delta^2 \right] + \frac{\Delta\iota_c}{2}} \\
\hat{\beta}_- &= \frac{1}{2} \frac{(2\Omega^2(\Omega^2 + \gamma^2 - \Delta^2) + \Delta\iota_c - 4i\gamma\Delta\Omega^2 + i\gamma\iota_c) \alpha_- + i\gamma\iota_c \alpha_-^\dagger}{\Omega^2 \left[(\Omega + i\gamma)^2 - \Delta^2 \right] + \frac{\Delta\iota_c}{2}}.
\end{aligned}$$

A.5 Optomechanical Filter

The mechanical susceptibility is defined as $\chi(\Omega) \equiv \frac{\hat{x}(\Omega)}{F(\Omega)}$. The equation of motion for a damped harmonic oscillator driven by a time-varying force $F(t)$ is given by

$$F(t) - m\omega_m^2 x - m\gamma_m \dot{x} = m\ddot{x}. \tag{A.30}$$

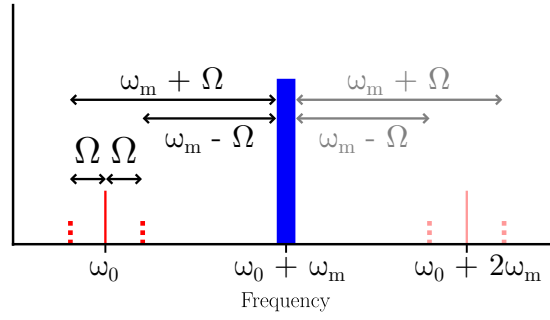


Figure A.1: The frequencies present in the optomechanical filter setup.

Transforming into the frequency domain we have,

$$\tilde{F}(\Omega) - m\omega_m^2 \tilde{x} - im\gamma_m \Omega \tilde{x} = -m\Omega^2 \tilde{x}, \quad (\text{A.31})$$

giving the susceptibility

$$\chi(\Omega) = \frac{1}{m(\omega_m^2 - \Omega^2 + i\gamma_m \Omega)}. \quad (\text{A.32})$$

In the optomechanical filter setup there are two main frequencies of interest; the laser frequency $\omega_0 + \omega_m$, and the cavity resonance frequency ω_0 . The sidebands of the laser are labelled as $\pm\omega_m \pm \Omega$ and of the cavity are labelled as $\pm\Omega$. We want to construct the relevant transfer function for the upper cavity sideband Ω . Since $\omega_0 < \omega_0 + \omega_m$ this involves making a frequency transformation of the lower output sideband to the relevant frequency (i.e. $\Omega < 0$).

First we write the general sideband input-output relation for the lower output sideband with $\Omega \rightarrow -\Omega$,

$$E_{\text{out}}(-\Omega) = T_+(-\Omega)E_{\text{in}}(-\Omega) + T_-(-\Omega)E_{\text{in}}(\Omega), \quad (\text{A.33})$$

so for the lower output sideband the transfer functions are swapped between the upper and lower input sidebands. Then we relate to the cavity sideband frequency Ω

by reading off Figure A.1,

$$E_{\text{out}}(\Omega - \omega_m) = T_+(\Omega - \omega_m)E_{\text{in}}(\Omega - \omega_m) + T_-(\Omega - \omega_m)E_{\text{in}}(-\Omega + \omega_m), \quad (\text{A.34})$$

so the upper output sideband for the cavity is a mixture of the upper input sideband for the cavity and the lower input sideband of the upper mechanical frequency at $\omega_0 + 2\omega_m$.

Therefore, if we set $\Delta = \omega_m$, make the substitution $\nu_c \rightarrow -4\hbar g^2 \Omega^2 \chi(\Omega)$ and look at the upper cavity sideband we have, after some rearranging,

$$\begin{aligned} \hat{\beta}_+(\Omega = \Omega - \omega_m) = & \left\{ 1 + \frac{2i\gamma[-i\gamma - 2\omega_m + \Omega - \hbar g^2 \chi(\Omega - \omega_m)]}{((\Omega - \omega_m) - i\gamma)^2 - \omega_m[\omega_m + 2\hbar g^2 \chi(\Omega - \omega_m)]} \right\} \hat{\alpha}_+(\Omega - \omega_m) \\ & - \frac{2i\hbar g^2 \gamma \chi(\Omega - \omega_m)}{((\Omega - \omega_m) - i\gamma)^2 - \omega_m[\omega_m + 2\hbar g^2 \chi(\Omega - \omega_m)]} \hat{\alpha}_+^\dagger(-\Omega + \omega_m). \end{aligned} \quad (\text{A.35})$$

For the next section we will concentrate only on the upper output cavity sideband.

Resolved Sideband Regime

The approximation of interest is the resolved sideband regime, $\Omega \ll \gamma \ll \omega_m$. We use this along with the definition,

$$\gamma_{\text{opt}} = \frac{\hbar g^2}{2m\omega_m\gamma} \implies \hbar g^2 = 2m\omega_m\gamma_{\text{opt}}\gamma. \quad (\text{A.36})$$

And also,

$$\chi(\Omega - \omega_m) = -\frac{1}{m[\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)]}, \quad (\text{A.37})$$

will be useful. First we will approximate the lower cavity sideband $T_-(\Omega - \omega_m)$,

$$T_-(\Omega - \omega_m) = \frac{\frac{2i\hbar g^2 \gamma}{m[\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)]}}{((\Omega - \omega_m) - i\gamma)^2 - \omega_m \left[\omega_m - \frac{2\hbar g^2}{m[\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)]} \right]} \quad (\text{A.38})$$

$$= \frac{4i\gamma_{\text{opt}}\gamma^2\omega_m}{(\gamma^2 - \Omega^2 + 2\Omega\omega_m + 2i\gamma(\Omega - \omega_m))(\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)) - 4\gamma_{\text{opt}}\gamma\omega_m^2} \quad (\text{A.39})$$

Then multiplying by $1/(\gamma\omega_m)$ on top and bottom, pulling out a factor ω_m on the bottom, and also pulling out an additional factor of ω_m from the first parenthesis on the denominator, and then applying the resolved sideband approximation,

$$\begin{aligned} T_-(\Omega - \omega_m) &\approx \frac{\gamma}{\omega_m} \frac{4i\gamma_{\text{opt}}}{4i\Omega - 4\gamma_{\text{opt}}} \\ &= \frac{\gamma}{\omega_m} \frac{\gamma_{\text{opt}}}{\Omega + i\gamma_{\text{opt}}}. \end{aligned} \quad (\text{A.40})$$

For the upper sideband $T_+(\Omega - \omega_m)$ we can rearrange to get

$$T_+(\Omega - \omega_m) = \frac{(\gamma^2 - 2i\gamma\omega_m - \Omega^2 + 2\Omega\omega_m)(\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)) + 4\Omega^2\gamma_{\text{opt}}\gamma + 2i\omega_m\gamma_{\text{opt}}\gamma^2}{(\gamma^2 - \Omega^2 + 2\Omega\omega_m + 2i\gamma(\Omega - \omega_m))(\Omega^2 - 2\Omega\omega_m - i\gamma_m(\Omega - \omega_m)) - 4\gamma_{\text{opt}}\gamma\omega_m^2}, \quad (\text{A.41})$$

and then dividing the top and bottom by ω_m^4 ,

$$T_+(\Omega - \omega_m) = \frac{\left(\frac{\gamma^2}{\omega_m^2} - 2i\frac{\gamma}{\omega_m} - \frac{\Omega^2}{\omega_m^2} + 2\frac{\Omega}{\omega_m}\right)\left(\frac{\Omega^2}{\omega_m^2} - 2\frac{\Omega}{\omega_m} - i\frac{\gamma_m\Omega}{\omega_m^2} + i\frac{\gamma_m}{\omega_m}\right) + 4\frac{\gamma_{\text{opt}}\gamma}{\omega_m^2} + 2i\frac{\gamma_{\text{opt}}\gamma^2}{\omega_m^3}}{\left(\frac{\gamma^2}{\omega_m^2} - \frac{\Omega^2}{\omega_m^2} + 2\frac{\Omega}{\omega_m} + 2i\frac{\gamma\Omega}{\omega_m^2} - 2i\frac{\gamma}{\omega_m}\right)\left(\frac{\Omega^2}{\omega_m^2} - 2\frac{\Omega}{\omega_m} - i\frac{\gamma_m\Omega}{\omega_m^2} + i\frac{\gamma_m}{\omega_m}\right) - 4\frac{\gamma_{\text{opt}}\gamma}{\omega_m^2}}. \quad (\text{A.42})$$

The last term on the numerator is much smaller than the one preceding it. Also we use that $\Omega \ll \gamma \implies \frac{\Omega^2}{\omega_m^2} \ll \frac{\Omega\gamma}{\omega_m^2}$, and $\gamma \ll \omega_m$, so that the only terms remaining are,

$$\begin{aligned} T_+(\Omega - \omega_m) &\approx \frac{4i\frac{\gamma\Omega}{\omega_m^2} + 4\frac{\gamma_{\text{opt}}\gamma}{\omega_m^2}}{4i\frac{\gamma\Omega}{\omega_m^2} - 4\frac{\gamma_{\text{opt}}\gamma}{\omega_m^2}} \\ &= \frac{\Omega - i\gamma_{\text{opt}}}{\Omega + i\gamma_{\text{opt}}} \approx -\exp\left(2i\frac{\Omega}{\gamma_{\text{opt}}}\right). \end{aligned} \quad (\text{A.43})$$

Appendix B

Radiation pressure tests

This appendix contains tests I wrote to validate FINESSE's behaviour concerning radiation pressure calculations. For a variety of setups, I will derive analytical expressions for the electric field, present a KatScript file that describes the same setup, and compare the output of FINESSE 3 to the analytical expression.

B.1 Mirrors

B.1.1 Single-Sided Free Mass

The layout for this setup is shown in Figure B.1.




Figure B.1: Setup for the first test case, an amplitude modulated laser with two sidebands reflecting off a free mass.

Analytics

We start with an amplitude modulated input field, such as produced by the `fsig` command in FINESSE,

$$E_i = E_0 \cos(\omega_0 t) \left(1 - \frac{A}{2} (1 - \cos(\Omega t)) \right)$$

where A is the modulation index. Writing this in the complex notation used in FINESSE gives

$$E_i = E_0 e^{i\omega_0 t} \left(1 - \frac{A}{2} (1 - \cos(\Omega t)) \right) \quad (\text{B.1})$$

$$= E_0 e^{i\omega_0 t} \left(1 - \frac{A}{2} + \frac{A}{4} (e^{i\Omega t} + e^{-i\Omega t}) \right) \quad (\text{B.2})$$

We can then separate this into three fields; the carrier field a_0 , and the upper and lower sidebands a_{\pm} , such that

$$E_i = a_0 + a_+ + a_-, \quad (\text{B.3})$$

where

$$a_0 = \left(1 - \frac{A}{2} \right) E_0 e^{i\omega_0 t}$$

$$a_+ = \frac{A}{4} E_0 e^{i(\omega_0 + \Omega)t}$$

$$a_- = \frac{A}{4} E_0 e^{i(\omega_0 - \Omega)t}.$$

Next, we consider this field when incident on a free-floating mirror. For a free mass, we have

$$F(t) = m\ddot{x}(t).$$

Taking the Fourier transform, we see that

$$\begin{aligned}\tilde{F}(\omega) &= -m\omega^2\tilde{x}(\omega) \\ \therefore \tilde{\chi}(\omega) &\equiv \frac{\tilde{x}(\omega)}{\tilde{F}(\omega)} = -\frac{1}{m\omega^2}.\end{aligned}$$

For a force of the form $F_0 \cos(\Omega t)$, we have

$$\tilde{F}(\omega) = \pi F_0(\delta(\omega + \Omega) + \delta(\omega - \Omega)),$$

from which we can find $x(t)$:

$$\begin{aligned}x(t) &= \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{\chi}(\omega)\tilde{F}(\omega)e^{i\omega t}d\omega \\ &= -\frac{\pi F_0}{2\pi} \int_{-\infty}^{+\infty} \frac{1}{m\omega^2}e^{i\omega t}(\delta(\omega + \Omega) + \delta(\omega - \Omega))d\omega \\ &= -\frac{F_0}{m\Omega^2} \frac{e^{i\Omega t} + e^{-i\Omega t}}{2} \\ &= -\frac{1}{m\Omega^2}F_0 \cos(\Omega t).\end{aligned}$$

A force of this form arises from the beating of the two amplitude modulation sidebands with the carrier. The force due to radiation pressure of light interacting with a mirror is given by

$$F_{\text{rad}} = \cos \alpha \frac{P_{1i} + P_{1o} - P_{2i} - P_{2o}}{c}, \quad (\text{B.4})$$

where $P_{\{1,2\}\{i,o\}}$ is the power of the light on the negative and positive x sides of the mirror incoming and outgoing respectively, and α is the angle of incidence. In this simple case, $\alpha = 0$, and there is no incoming beam from the positive x direction.




Figure B.2: Phase change upon reflection from a mirror in FINESSE.

The force can then be calculated to leading order in A from Equation (B.1) as

$$\begin{aligned}
 F_{\text{rad}} &= \frac{|E_i|^2 + |E_r|^2 - |E_t|^2}{c} \\
 &= \frac{(1 + R - T)}{c} |E_i|^2 \\
 &= \frac{(1 + R - T)}{c} \left(E_0^2(1 - A) + AE_0^2 \cos(\Omega t) \right),
 \end{aligned}$$

where R is the coefficient of reflectivity of the mirror. Here we are uninterested in the D.C. term as this will not contribute to the mirror's oscillation, and so have

$$\begin{aligned}
 F_0 &= \frac{(1 + R - T)AE_0^2}{c}, \\
 \therefore x(t) &= -\frac{(1 + R - T)AE_0^2}{mc\Omega^2} \cos(\Omega t). \tag{B.5}
 \end{aligned}$$

Figure B.2 shows how an incoming field changes upon reflection from the mirror. The incident field will be multiplied by the reflection coefficient of the mirror, r . The reflected light will also be phase shifted due to propagation through space by a total amount $e^{-2i\phi}$, where

$$\begin{aligned}
 \phi &= \frac{2\pi x(t)}{\lambda} \\
 &= -\frac{2\pi(1 + R - T)AE_0^2}{\lambda mc\Omega^2} \cos(\Omega t) \\
 &= -\frac{(1 + R - T)AE_0^2 \omega_0}{mc^2 \Omega^2} \cos(\Omega t).
 \end{aligned}$$

This results in the incoming beam being phase modulated by the mirror motion,

$$E_r = rE_i \exp(iB \cos(\Omega t)), \quad (\text{B.6})$$

where

$$B = \frac{2(1 + R - T)AE_0^2\omega_0}{mc^2\Omega^2},$$

which to first order gives

$$\begin{aligned} E_r &= rE_i \left(1 + i\frac{B}{2} (e^{i\Omega t} + e^{-i\Omega t}) \right) \\ &= r(b_0 + b_+ + b_-), \end{aligned}$$

where

$$\begin{aligned} b_0 &= E_i \\ b_+ &= i\frac{B}{2}E_i e^{i\Omega t} \\ b_- &= i\frac{B}{2}E_i e^{-i\Omega t}. \end{aligned}$$

Putting Equation (B.3) into this gives

$$\begin{aligned} b_0 &= a_0 + a_+ + a_- \\ b_+ &= i\frac{B}{2}E_0 e^{i(\omega_0 + \Omega)t} \\ b_- &= i\frac{B}{2}E_0 e^{i(\omega_0 - \Omega)t}, \end{aligned}$$

to leading order, where we can once again see an input field and an upper and lower

```

1  l l1 P=3
2  s s1 l1.p1 m1.p1 L=5
3  m m1 R=0.5 T=0.4
4  free_mass m1_sus m1.mech 0.5e-3
5
6  ad upper s1.p1.o &fsig.f
7  ad lower s1.p1.o -&fsig.f
8
9  sgen sig l1.amp.i 0.1
10
11 fsig(0.1)
12 xaxis(sig.f, log, 0.1, 1e7, 400)

```

Listing B.1: KatScript representation of the setup shown in Figure B.1.

sideband. The full expression for the reflected field can then be written as

$$E_r = r(a_0 + a_+ + a_- + b_+ + b_-) \quad (\text{B.7})$$

$$\begin{aligned}
&= rE_0 e^{i\omega_0 t} \left(1 + \left(\frac{A + 2iB}{4} \right) (e^{i\Omega t} + e^{-i\Omega t}) \right) \\
&= rE_0 e^{i\omega_0 t} \left(1 + \left(\frac{A}{4} + i \frac{(1 + R - T)AE_0^2 \omega_0}{mc^2 \Omega^2} \right) (e^{i\Omega t} + e^{-i\Omega t}) \right) \quad (\text{B.8})
\end{aligned}$$

Finesse

The analytics above were compared to the sidebands produced the KatScript file shown in Listing B.1. The results are shown in Figure B.3.

B.1.2 Double-Sided Free Mass, One Modulation

The layout for this setup is shown in Figure B.4.

Analytics

This is similar to Appendix B.1.1, however there are no amplitude modulation sidebands present in the input field. As a result, we only expect to see the output




Figure B.3: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.1




Figure B.4: Setup for the second test case, an amplitude modulated laser with two sidebands reflecting off a free mass causing sidebands to be generated in an unmodulated laser incident on the other side.


```

1  l l1 P=3
2  s s1 l1.p1 m1.p1 L=5
3  m m1 R=1 T=0
4  s s2 m1.p2 l2.p1 L=5
5  l l2 P=2
6
7  free_mass m1_sus m1.mech 0.5e-3
8
9  ad upper s2.p2.o &fsig.f
10 ad lower s2.p2.o -&fsig.f
11
12 sgen sig l1.amp.i 0.1
13
14 fsig(0.1)
15 xaxis(sig.f, log, 0.1, 1e7, 400)

```

Listing B.2: KatScript representation of the setup shown in Figure B.4.

containing the phase modulation sidebands. The second laser will not contribute to the oscillation of the mirror, as it is a purely D.C. force. The sidebands will also be 180° out of phase, as the mirror's motion is inverted when looking from the other side. In addition, we assume that the mirror is opaque i.e. $T = 0$. For a given input field

$$E_i = E_1 \cos(\omega_0 t),$$

we therefore have

$$E_r = rE_1 \left(1 - i \frac{(1+R)AE_0^2\omega_0}{mc^2\Omega^2} (e^{i\Omega t} + e^{-i\Omega t}) \right).$$

Finesse

The analytics above were compared to the sidebands produced the KatScript file shown in Listing B.2. The results are shown in Figure B.5.




Figure B.5: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.2

B.1.3 Double-Sided Free Mass, Two Modulations

The layout for this setup is shown in Figure B.6.

Analytics

As the second beam is now also amplitude modulated, it will contribute to the mirror's motion, so we must recalculate the radiation pressure force from eq. (B.4).

For the left and right input fields,

$$E_{\text{Left}} = E_0 \cos(\omega_0 t) \left(1 - \frac{A_0}{2} (1 - \cos(\Omega t)) \right)$$

$$E_{\text{Right}} = E_1 \cos(\omega_0 t) \left(1 - \frac{A_1}{2} (1 - \cos(\Omega t - \psi)) \right),$$




Figure B.6: Setup for the third test case, two amplitude modulated lasers with two sidebands each incident on each side of a mirror, with some arbitrary phase difference.

we then have

$$F_{\text{rad}} = \frac{1+R}{c} \left(E_0^2(1-A_0) + A_0 E_0^2 \cos(\Omega t) - E_1^2(1-A_1) - A_1 E_1^2 \cos(\Omega t - \psi) \right),$$

or ignoring the D.C. terms,

$$F_{\text{rad}} = \frac{1+R}{c} \left(A_0 E_0^2 \cos(\Omega t) - A_1 E_1^2 \cos(\Omega t - \psi) \right).$$

To write this as a single cosine, we use some trigonometry:

$$\begin{aligned} F_{\text{rad}} &= \frac{1+R}{c} \left[A_0 E_0^2 \cos(\Omega t) - A_1 E_1^2 (\cos(\Omega t) \cos(\psi) + \sin(\Omega t) \sin(\psi)) \right] \\ &= \frac{1+R}{c} \left[(A_0 E_0^2 - A_1 E_1^2 \cos(\psi)) \cos(\Omega t) - A_1 E_1^2 \sin(\psi) \sin(\Omega t) \right]. \end{aligned}$$

Writing the term in brackets in the form $D \cos(x + \theta)$, we get

$$\begin{aligned} D^2 &= (A_0 E_0^2 - A_1 E_1^2 \cos(\psi))^2 + (A_1 E_1^2 \sin(\psi))^2 \\ D &= \sqrt{A_0^2 E_0^4 - 2A_0 A_1 E_0^2 E_1^2 \cos(\psi) + A_1^2 E_1^4} \end{aligned} \quad (\text{B.9})$$

and

$$\theta = \tan^{-1} \left(\frac{A_1 E_1^2 \sin(\psi)}{A_0 E_0^2 - A_1 E_1^2 \cos(\psi)} \right). \quad (\text{B.10})$$

We can now proceed in the same way as in Appendix B.1.1, to obtain an expression for the reflected field on the left hand side of the mirror,

$$E_r = r E_0 e^{i\omega_0 t} \left[1 + \left(\frac{A_0}{4} (e^{i\Omega t} + e^{-i\Omega t}) + i \frac{(1+R)D\omega_0}{mc^2\Omega^2} (e^{i(\Omega t+\theta)} + e^{-i(\Omega t+\theta)}) \right) \right]. \quad (\text{B.11})$$

To get the field on the right hand side of the mirror, we can apply the following (remembering that θ will also be affected):

$$E_0 \iff E_1$$

$$A_0 \iff A_1$$

$$\psi \implies -\psi.$$

Finesse

The analytics above were compared to the sidebands produced by the KatScript file shown in Listing B.3. For comparison with the analytical results, we must also include an extra factor of $e^{\pm i\psi}$ in the upper / lower sidebands respectively when looking at the right hand side field. This is because FINESSE measures phase relative to the first laser in the file, which in this case is the left hand side laser. The results are shown in Figures B.7 and B.8.

```

1  l  l1  P=3
2  s  s1  l1.p1  m1.p1  L=5
3  m  m1  R=1  T=0
4  s  s2  m1.p2  l2.p1  L=5
5  l  l2  P=4
6
7  free_mass  m1_sus  m1.mech  0.5e-3
8
9  ad  upper_left  s1.p1.o  &fsig.f
10 ad  lower_left  s1.p1.o  -&fsig.f
11
12 ad  upper_right  s2.p2.o  &fsig.f
13 ad  lower_right  s2.p2.o  -&fsig.f
14
15 sgen  sig1  l1.amp.i  0.1
16 sgen  sig2  l2.amp.i  0.1  phase=29
17
18 fsig(0.1)
19 xaxis(fsigs.f, log, 0.1, 1e7, 400)

```

Listing B.3: KatScript representation of the setup shown in Figure B.6.




Figure B.7: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.3 for the reflected field on the left of the mirror.




Figure B.8: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.1.3 for the reflected field on the right of the mirror.

B.2 Beamsplitter

The beamsplitter behaves similarly to the mirror, with two key differences. Firstly, the angle of incidence α in Equation (B.4) is no longer necessarily 0. This also affects the phase change on reflection due to the beamsplitter tuning, ϕ . Secondly, there are now 4 ports at which light can enter and exit, rather than two. Here we repeat similar experiments to the previous section, taking these factors into account.

B.2.1 Single-Sided Free Mass

The layout for this setup is shown in Figure B.9.




Figure B.9: Setup for the first beamsplitter test case, an amplitude modulated laser with two sidebands reflecting off a free mass at some non-zero angle.

Analytics

The steps here are much the same as in Appendix B.1.1, with the only changes being the inclusion of α in the radiation pressure force calculation, and a change in the phase shift due to mirror tuning ϕ — the total phase shift of reflected light will now be $e^{-2i\phi \cos \alpha}$, as the tuning ϕ is defined relative to the surface normal. We therefore have

$$x(t) = -\frac{(1 + R - T)AE_0^2}{mc\Omega^2} \cos(\Omega t) \cos \alpha,$$

$$\therefore \phi = -\frac{(1 + R - T)AE_0^2\omega_0}{mc^2\Omega^2} \cos(\Omega t) \cos \alpha.$$

This results in the incoming beam being phase modulated by the mirror motion,

$$E_r = rE_i \exp(iB \cos(\Omega t)), \quad (\text{B.12})$$

where

$$B = \frac{2(1 + R - T)AE_0^2\omega_0}{mc^2\Omega^2} \cos^2 \alpha,$$

```

1  l l1 P=3
2  s s1 l1.p1 bs1.p1 L=5
3  bs bs1 R=0.5 T=0.4 alpha=37
4  s s2 bs1.p2 out.p1 L=5
5  nothing out
6
7  free_mass bs1_sus bs1.mech 0.5e-3
8
9  ad upper s2.p2.o &fsig.f
10 ad lower s2.p2.o -&fsig.f
11
12 sgen sig l1.amp.i 0.1
13
14 fsig(0.1)
15 xaxis(fsig.f, log, 0.1, 1e7, 400)

```

Listing B.4: KatScript representation of the setup shown in Figure B.9.

Following the same steps as in Appendix B.1.1, the full expression for the reflected field can then be written as

$$\begin{aligned}
 E_r &= rE_0 e^{i\omega_0 t} \left(1 + \left(\frac{A + 2iB}{4} \right) (e^{i\Omega t} + e^{-i\Omega t}) \right) \\
 &= rE_0 e^{i\omega_0 t} \left(1 + \left(\frac{A}{4} + i \frac{(1 + R - T)AE_0^2 \omega_0}{mc^2 \Omega^2} \cos^2 \alpha \right) (e^{i\Omega t} + e^{-i\Omega t}) \right). \quad (\text{B.13})
 \end{aligned}$$

Finesse

The analytics above were compared to the sidebands produced by the KatScript file shown in Listing B.4. The results are shown in Figure B.10.

B.2.2 Double-Sided Free Mass, One Modulation

The layout for this setup is shown in Figure B.11.




Figure B.10: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.1




Figure B.11: Setup for the second beamsplitter test case, an amplitude modulated laser with two sidebands reflecting off a free mass causing sidebands to be generated in an unmodulated laser incident on the other side.

```

1  l  l1 P=3
2  s  s1 l1.p1 bs1.p1 L=5
3
4  l  l2 P=2
5  s  s2 l2.p1 bs1.p4 L=5
6
7  bs bs1 R=1 T=0 alpha=37
8  s  sout bs1.p3 out.p1 L=5
9  nothing out
10
11 free_mass bs1_sus bs1.mech 0.5e-3
12
13 ad upper sout.p2.o &fsig.f
14 ad lower sout.p2.o -&fsig.f
15
16 sgen sig l1.amp.i 0.1
17
18 fsig(0.1)
19 xaxis(fsig.f, log, 0.1, 1e7, 400)

```

Listing B.5: KatScript representation of the setup shown in Figure B.11.

Analytics

This is similar to Appendix B.1.2, with the inclusion of the $\cos^2 \alpha$ term from the previous section. For a given input field

$$E_i = E_1 \cos(\omega_0 t),$$

we therefore have

$$E_r = r E_1 \left(1 - i \frac{(1+R) A E_0^2 \omega_0}{m c^2 \Omega^2} \cos^2 \alpha \left(e^{i \Omega t} + e^{-i \Omega t} \right) \right).$$

Finesse

The analytics above were compared to the sidebands produced by the KatScript file shown in Listing B.5. The results are shown in Figure B.12.

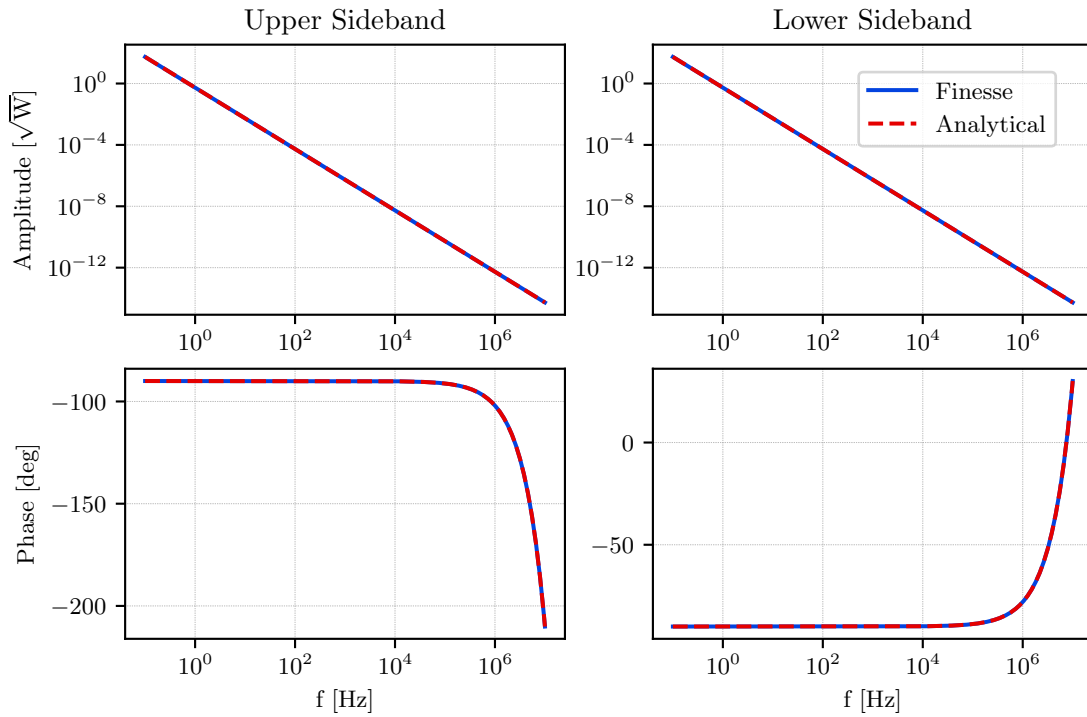


Figure B.12: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.2

B.2.3 Double-Sided Free Mass, Two Modulations

The layout for this setup is shown in Figure B.13.

Analytics

Once again, this is similar to the procedure in Appendix B.1.3, with the addition of a $\cos^2 \alpha$ term on the amplitude of the generated sidebands. The output at the top detector will therefore be

$$E_r = rE_0 e^{i\omega_0 t} \left[1 + \left(\frac{A_0}{4} (e^{i\Omega t} + e^{-i\Omega t}) + i \frac{(1+R)D\omega_0}{mc^2\Omega^2} \cos^2 \alpha (e^{i(\Omega t+\theta)} + e^{-i(\Omega t+\theta)}) \right) \right], \quad (\text{B.14})$$

with D defined as in Appendix B.1.3. The results for the right detector can be obtained with the same transformations as before.




Figure B.13: Setup for the third beamsplitter test case, two amplitude modulated lasers with two sidebands each incident on each side of a beamsplitter, with some arbitrary phase difference.

Finesse

The analytics above were compared to the sidebands produced by the KatScript file shown in Listing B.6. As in Appendix B.1.3, for comparison with the analytical results, we must also include an extra factor of $e^{\pm i\psi}$ in the upper / lower sidebands respectively when looking at the right hand side field. This is because FINESSE measures phase relative to the first laser in the file, which in this case is the left hand side laser. The results are shown in Figures B.14 and B.15.

```
1 l l1 P=3
2 s s1 l1.p1 bs1.p1 L=5
3
4 l l2 P=4
5 s s2 l2.p1 bs1.p4 L=5
6
7 bs bs1 R=1 T=0 alpha=37
8
9 s sout_left bs1.p2 out_left.p1 L=5
10 nothing out_left
11
12 s sout_right bs1.p3 out_right.p1 L=5
13 nothing out_right
14
15 free_mass bs1_sus bs1.mech 0.5e-3
16
17 ad upper_left sout_left.p2.o &fsig.f
18 ad lower_left sout_left.p2.o -&fsig.f
19
20 ad upper_right sout_right.p2.o &fsig.f
21 ad lower_right sout_right.p2.o -&fsig.f
22
23 sgen sig1 l1.amp.i 0.1
24 sgen sig2 l2.amp.i 0.1 phase=29
25
26 fsig(0.1)
27 xaxis(fsig.f, log, 0.1, 1e7, 400)
```

Listing B.6: KatScript representation of the setup shown in Figure B.13.




Figure B.14: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.3 for the reflected field at the top detector.




Figure B.15: Comparison of the upper and lower output sidebands from FINESSE and Appendix B.2.3 for the reflected field at the right detector.