Spencer, Andrew P. (2020) *Advanced techniques in laser interferometry for current and future gravitational wave detectors.* PhD thesis.

https://theses.gla.ac.uk/81266/

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

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

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

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

Enlighten: Theses
https://theses.gla.ac.uk/
research-enlighten@glasgow.ac.uk

Advanced Techniques in Laser Interferometry for Current and Future Gravitational Wave Detectors

Andrew P. Spencer

Submitted in fulfillment of the requirements for the
Degree of Doctor of Philosophy

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



August 2019

Abstract

The current generation of gravitation wave detectors have been highly successful in observing gravitational wave signals from the inspiral and merging of astrophysical compact binary objects such as black-holes and neutron stars. While these gravitational wave detectors remain in operation, plans are being made to improve the sensitivity of these detectors and to build more sensitive detectors.

This thesis covers a series of experimental work in the field of laser interferometry for gravitational wave detector instrumentation, including an investigation of polarising optics for a speedmeter interferometry and topics in laser system development for current and future gravitational wave detectors.

The concept of the speedmeter interferometer is considered as a gravitational wave detector with lower quantum noise than the Michelson interferometers currently employed. The practical limitations of one such speedmeter scheme; the polarisation Sagnac is investigated with a single arm prototype experiment.

For the polarisation Sagnac speedmeter it is important to demonstrate that polarising optics currently available are of a high enough quality to not limit the performance of a full-scale detector. This can be done by individual testing of the polarising beam-splitter (PBS) and quarter waveplate (QWP) optics, as is shown in Chapter 3, and with a prototype scale interferometer demonstration in the Glasgow 10 m prototype facility as described in Chapter 4.

In addition to polarising optics, low levels of birefringence in the cavity mirrors can place a significant limitation on the performance of the polarisation Sagnac speedmeter as this scheme relies on a conservation of the polarisation state of the light within the cavities to ensure the optical beams are directed along the desired path through the interferometer. Any light coupling to a different path though the interferometer will degrade the performance of the interferometer as a speedmeter.

In Chapter 4 and 5 this cavity birefringence effect is measured with a 10 m suspended linear cavity experiment. The polarisation of the light transmitted by the cavity was measured for different rotations of one cavity mirror to measure the

effect of mirror cavity birefringence on the polarisation state of the light circulating in the cavity.

In Chapter 6 the combination of polarisation imperfections is considered in terms of the polarisation Sagnac speedmeter to determine if this scheme can be considered a viable option for future gravitational wave detection. It is shown how the cavity birefringence effect can be combined with the birefringence level of the QWP to cancel out unwanted levels of birefringence for each arm of the polarisation Sagnac speedmeter. Under with circumstances the performance of the interferometer in terms of polarisation optics is limited only by the finite extinction ratio of the PBS optic.

In Chapter 9, work to upgrade the laser system of the LIGO Livingston detector is described. This work consisted of the installation of an additional laser amplifier to the laser system that would double the power available to the interferometer. Resulting in an estimated factor of $\sqrt{2}$ improvement in the detector sensitivity.

In Chapter 10, two investigations into the laser stabilisation system of the LIGO Livingston are presented. This includes an investigation of a persistent issue with the frequency stabilisation system that had been limiting the control bandwidth of the stabilisation loop. This issue was found to be related to the analogue servo electronics for this system. Another issue with the intensity stabilisation system that occurred during the second observing run of the Advanced LIGO detectors in 2017 was investigated as dropouts in this control loop were causing the full interferometer to lose lock.

Chapter 11 describes the installation and testing of a new laser system with emission at 1550 nm wavelength for the Glasgow 10 m prototype facility. This include the development and troubleshooting of a frequency stabilisation scheme for the laser with significantly higher intrinsic frequency noise when compared to the standard 1064 nm laser used in Advanced LIGO and other gravitational wave interferometry experiments. With the described stabilisation system the noise of the laser is reduced to $0.1\,\frac{\mathrm{Hz}}{\sqrt{\mathrm{Hz}}}$ between 100 Hz and 1 kHz with a unity gain point at 500 kHz.

In Chapter 12 a digital control system for laser frequency stabilisation is considered and an algorithm for filter design optimisation is described and tested.

Contents

Abstract **i**

Acknowledgements **xix**

Preface **xx**

Declaration **xxii**

I Introduction to Interferometry for Gravitational Wave Detection **1**

1 Gravitational Wave Interferometry **2**

1.1 Gravitational Wave Detection 2
1.2 The Gravitational Wave Detector Network 4
1.3 Noise Sources in Gravitational Wave Detectors 7
1.4 Mathematical Description of Light 11
1.4.1 Interaction of Light and Optics 11
1.4.2 Polarisation 13
1.4.3 Gaussian Beams 14
1.4.4 Modulation 18
1.4.5 Interaction of Light with a Moving Mirror 22
1.5 Interferometry Basics 23
1.5.1 Michelson Interferometer 23
1.5.2 Fabry-Pérot Cavity 25
1.5.3 Power Recycling 27
1.5.4 Signal Recycling 28
1.5.5 Dual Recycling Fabry-Pérot Michelson Interferometer . . . 28
1.5.6 Optical Readout and Signal Detection 28

2 Quantum Noise **31**
2.1 Quantum Noise of Light 31
2.2 Quantum Noise in Interferometers 33
2.3 Standard Quantum Limit 35
2.3.1 Schemes for Reducing Quantum Noise 36
2.4 Speedmeter Interferometry 38
2.4.1 Speedmeter Topologies 39
2.5 Practical Limitations to Speedmeter Designs 46
2.5.1 Speedmeter Mechanism 46
2.5.2 Infrastructure 48

II Experimental Investigation of the Polarisation Sagnac Speedmeter Topology **49**

3 Characterisation of Polarisation Optics **50**
3.1 Polarising Optics for the Sagnac Speedmeter 50
3.2 Birefringence of Non-Polarising Optics for the Sagnac Speedmeter 54
3.2.1 Cavity Mirrors 54
3.3 Characterisation of Polarising Optics 56
3.3.1 PBS 57
3.3.2 QWP 58
3.4 Summary 61

4 Polarisation Sagnac Experiment **64**
4.1 Design of the Single Arm Polarisation Sagnac Speedmeter Experiment 65
4.1.1 Suspension Design 65
4.1.2 Ellipticity Readout 71
4.2 The Single Arm Polarisation Sagnac Speedmeter Experiment . . . 71
4.2.1 Locking the Cavity 71
4.2.2 Ellipticity Measurements 78
4.2.3 Rotating the ETM 80
4.3 Results of the Single Arm Polarisation Sagnac Speedmeter Experiment 80
4.4 Summary 81

5 Results of the Linear Cavity Experiment **83**
5.1 Jones Model of Linear Arm Cavity 83
5.2 Fitting a Jones Model to the Results of the Linear Cavity Experiment 85

5.2.1 Inferring Input Beam Parameters from Auxiliary Monitors 85
5.2.2 Results with No Resonant Cavity 86
5.2.3 Fitting Model to Cavity Data 86
5.3 Inferring Mirror Parameters 89
5.4 Consideration of other Birefringence Effects 93
5.5 Summary 96

6 Modelling a Polarisation Sagnac Speedmeter 98
6.1 Jones Model of Polarisation Sagnac Speedmeter 98
6.1.1 Optical Paths in the Polarisation Sagnac Speedmeter Interferometer 98
6.1.2 Limits on the Speedmeter Performance Relating to Individual Optics 100
6.2 Polarisation Sagnac Speedmeter with Realistic Optics 101
6.2.1 Effect of Polarisation Beam Splitter Extinction Ratio on the Sagnac Speedmeter 101
6.2.2 Effect of Imperfect Quarter Waveplate on the Sagnac Speedmeter 102
6.2.3 Effect of Cavity Birefringence on the Sagnac Speedmeter . 104
6.2.4 Combined Effect of Imperfect Polarising Optics on the Sagnac Speedmeter 105
6.3 Michelson Interferometer with Cavity Birefringence 109
6.4 Summary 111

7 Conclusion 113
7.1 Future Outlook For Speedmeter Interferometry 114

III Laser Development for Current and Future Gravitational Wave Detectors 116

8 Laser Technology Introduction 117
8.1 Laser Fundamentals 117
8.2 Laser Technology for Gravitational Wave Interferometry 119
8.2.1 Source Lasers 119
8.2.2 Laser Amplifiers 121
8.2.3 Laser Modulators 122
8.2.4 Wavelength choice for gravitational wave detectors 124
8.3 Laser Noise and Stabilisation Techniques 126

8.3.1 Passive Noise Reduction . . . 127
8.3.2 Laser Intensity Stabilisation . . . 128
8.3.3 Laser Frequency Stabilisation . . . 128
8.4 The Advanced LIGO PSL System . . . 132
8.4.1 Laser Amplification . . . 134
8.4.2 Intensity Stabilisation . . . 135
8.4.3 Frequency Stabilisation . . . 135
8.4.4 Premode Cleaner . . . 138
8.4.5 PSL output . . . 139

9 70 W Laser Amplifier for LIGO 140
9.1 High Power Laser Options for Advanced LIGO . . . 140
9.1.1 High Power Oscillator . . . 141
9.1.2 Laser Amplifier for O3 . . . 142
9.2 The neoVAN Laser Amplifier . . . 142
9.3 Integration into the PSL System . . . 143
9.4 Testing and Installation . . . 144
9.4.1 Alignment to the Amplifier . . . 144
9.4.2 Mode-matching to the Amplifier . . . 146
9.4.3 Beam Quality of the Amplifier Output . . . 146
9.4.4 Alignment to the PMC . . . 147
9.4.5 Power Stability . . . 149
9.5 Detector Sensitivity at Higher Laser Power . . . 150

10 Laser Stabilisation for aLIGO 152
10.1 Advanced LIGO Table-top Frequency Stabilisation Scheme . . . 152
10.1.1 Modifications to the TTFSS . . . 154
10.2 Investigating High Frequency Stability of the TTFSS . . . 155
10.2.1 Modelling the TTFSS . . . 155
10.2.2 Previous Investigation of TTFSS Instability . . . 158
10.3 Conclusion of Investigation into TTFSS Instability . . . 159
10.4 Advanced LIGO Inner-Loop Intensity Stabilisation Servo . . . 160
10.5 Lock Losses Due to ISS Instability . . . 160
10.5.1 ISS Dropouts and Lock Loss . . . 161
10.6 Investigation into ISS Dropouts . . . 162
10.6.1 Timing of ISS Dropouts . . . 162
10.6.2 ISS Loop Gain . . . 163
10.6.3 Lock Loss Tool . . . 165

10.7 Conclusion of Investigation into ISS Instability 168

11 Frequency Stabilised 1550 nm Laser 169
11.1 The 1550 nm Laser System 169
11.1.1 Diode Laser 171
11.1.2 Fibre Laser Amplifier 172
11.1.3 Auxiliary Optics 173
11.1.4 Performance of the Laser System 176
11.2 Frequency Stabilisation of a 1550 nm Laser 179
11.2.1 Reference Cavity 179
11.2.2 Stabilisation Scheme Design 180
11.2.3 Stabilisation Loop Response 183
11.2.4 Analogue Servo Filter Design 184
11.3 Results 188
11.3.1 Limitations of the Stabilisation Servo 188
11.3.2 Frequency Noise Measurements 192
11.4 Summary 196

12 Digital Stabilisation Servo 198
12.1 High Speed Digital Filtering 199
12.1.1 Discretisation 199
12.1.2 Implementation of digital filters 200
12.1.3 Aliasing 200
12.1.4 dSPACE Digital Processor 201
12.2 Laser Frequency Stabilisation with a Digital Servo 202
12.2.1 Digital Filter Design for the Laser Stabilisation System . . 202
12.3 Laser Stabilisation Results 205
12.4 Digital Filter Optimisation 207
12.4.1 Optimisation Algorithm 208
12.4.2 Results 218

13 Conclusion 219
13.1 Summary of LIGO Laser Work 219
13.2 Summary of 1550 nm Laser Work 220

IV Appendices 221

A Hardware for Linear Cavity Experiment 222

A.1 Vacuum System . . . 222
A.2 Seismic Isolation . . . 222
A.3 Suspensions . . . 225
A.3.1 Suspension Modes . . . 225
A.3.2 Suspension Control . . . 225

B Electronics for the 1550 nm Laser System 230
B.1 Cavity PZT Driver . . . 230
B.2 Laser PZT Driver . . . 230
B.3 Photodiodes . . . 232
B.3.1 Polarisation Optics Characterisation Photodiodes . . . 232
B.3.2 Reference Cavity Locking Photodiode . . . 232
B.3.3 10 m Cavity Locking Photodiode . . . 232
B.3.4 10 m Cavity Transmittion Photodiode . . . 233
B.4 Notch Filter . . . 233
B.5 AA/AI Filter . . . 235

C Jones Calculus 239
C.1 Jones Calculus for Optics and Light . . . 239
C.1.1 Light Fields . . . 239
C.1.2 Optics . . . 239
C.2 Linear Cavity Model . . . 241
C.3 Speedmeter Model . . . 242
C.4 Derivation of Limits to Speedmeter . . . 242
C.4.1 PBS . . . 242
C.4.2 QWP . . . 243

D Digital Servo 244
D.1 Simulink Model . . . 244

List of Tables

4.1 Key properties of cavity optics for the single arm polarisation Sagnac speedmeter experiment 65
4.2 Key cavity parameters of the single arm polarisation Sagnac speedmeter experiment 66

5.1 Birefringent parameter of the cavity mirrors based on the results of the linear cavity experiment. 92

11.1 Comparison of important datasheet parameters of the OP27 and TL071 op-amps 192

C.1 Jones Vectors to describe polarisation states of optical fields. . . . 240

List of Figures

1.1 Gravitational wave detector concept. 3
1.2 Gravitational wave detections made during the first and second observing runs of the advanced gravitational wave detectors. . . . 4
1.3 Modelled strain sensitivity of Advanced LIGO and principle noise sources. 9
1.4 Basic interaction of light field with a partially reflective mirror. . . 12
1.5 Basic interaction of light field with a beam-splitter with four ports. 12
1.6 Gaussian beam representation of a laser. 15
1.7 Transverse profiles of HG_{nm} laser beams of up to sixth order Hermite-Gaussian modes. 16
1.8 Transverse profiles of LG_{pl} laser beams of up to sixth order Laguerre-Gaussian modes. 17
1.9 Sketches of amplitude (a) and phase (b) modulation of a carrier field (dashed line). 18
1.10 Diagram explaining amplitude modulation. 19
1.11 Diagram explaining frequency modulation. 21
1.12 Basic interaction of light field with a partially reflective mirror in motion. 22
1.13 A Michelson interferometer. 24
1.14 A Fabry-Pérot Cavity. 26
1.15 Dual-recycling Fabry-Pérot Michelson interferometer for gravitational wave detection. 29
1.16 Simulated sensitivity of a 4 km scale Michelson interferometer with different interferometric enhancements. 30

2.1 Representation of a quantised carrier state as a time series (a) and in quadrature form (b). 32
2.2 Representation of a quantised vacuum state as a time series (a) and in quadrature form (b). 32

2.3 Quantum noise in an interferometer. 34
2.4 Simple representation of a speedmeter. 38
2.5 The speedmeter zoo. 39
2.6 The Ring Sagnac Speedmeter. 40
2.7 The Polarisation Sagnac Speedmeter. 42
2.8 The sloshing speedmeter. 43
2.9 The sloshing Sagnac speedmeter. 44
2.10 The polarisation circulation speedmeter. 47

3.1 Basic principle of a polarising beam-splitter. 52
3.2 Basic principle of a quarter waveplate. 53
3.3 Experimental setup to measure the extinction ratio of a PBS. . . 56
3.4 Extinction data recorded with s and p polarised light in transmission and reflection from the PBS normalised to the total input power. 58
3.5 Measured polarisation extinction ratio of the PBS in transmission and reflection as a function of the angle of incidence on the PBS optic. The manufactures specification of a transmission extinction ration of 2500 is shown as a grey dashed line. 59
3.6 Experimental setup to measure the birefringence of a QWP. . . . 60
3.7 Results of the QWP characterisation at normal incidence. 61
3.8 Results of the QWP characterisation at normal incidence with the QWP fast axis between 35° and 55°. 62
3.9 Results of the QWP characterisation at difference angles of incidence (AOI) 63

4.1 The single arm polarisation Sagnac speedmeter experiment in the Glasgow 10 m prototype system. 65
4.2 Suspensions for the polarisation Sagnac speedmeter. 66
4.3 Modification of the suspension mass to hold the PBS to widen the aperture. 68
4.4 Modification of the QWP suspension to allow control of the optics roll degree of freedom. 69
4.5 Test setup for aligning the QWP suspension. 70
4.6 Optical layout of the readout scheme for ellipticity. 72
4.7 Cavity locking and laser stabilisation scheme. 73
4.8 Signals to lock the 10 m cavity with and without active laser noise reduction. 74

4.9 Transfer functions of the component filters used to lock the 10 m cavity. 75
4.10 The transfer function of the filter designed to compensate for the suspension modes of motion of the cavity locking servo (a) and the effect of applying this filter to the in-lock error point noise (b) between 10 Hz and 100 Hz. 76
4.11 In-lock error point noise power spectrum. 77
4.12 In-lock transfer functions of the 10 m cavity locking loop. 77
4.13 Transmitted DC power from the locked 10 m cavity. 79
4.14 Sum and difference in DC power transmitted from the locked 10 m cavity in orthogonal polarisations 80
4.15 Rotating the ETM. 81
4.16 Measured ellipticity induced by the resonant cavity for the six orientations of the ETM mirror. 81

5.1 Linear cavity formed by two mirrors the input test-mass (ITM) and the end test-mass (ETM). Each mirror has a birefringence η with fast axis orientated at an angle θ. 84
5.2 Experimental setup of the linear cavity experiment. 86
5.3 Data recorded with no cavity. 87
5.4 Plots showing time variation of measured ellipticity (a), cavity detuning (b), and environmental temperature (c) for all six data sets. 88
5.5 Cavity locking error point for data recorded with the ETM at 30° in terms of cavity detuning in Hz 90
5.6 Jones model of linear cavity fitted to experimental data. 90
5.7 Parameter estimation for birefringent mirrors in a linear cavity Jones model. 91
5.8 Parameter estimation for birefringent mirrors in a linear cavity Jones model. 92
5.9 Ellipticity data as a function of ETM angle with the fitted result of the Jones model parameter estimation. The pink traces show the range of fits based on the parameter estimation shown in Figure 5.8 with the thick pink trace representing the best fit parameters. 93
5.10 Ellipticity as a function of ETM and ITM rotation based on the linear cavity results. 94

6.1 Optical paths through polarisation Sagnac speedmeter from input to readout. 99

6.2 Measured polarisation extinction ratio of the PBS in transmission and reflection as a function of the angle of incidence on the PBS optic. 102
6.3 Modelled effect of PBS extinction ratio on the signal ratio in a speedmeter interferometer. 103
6.4 Modelled effect of QWP birefringence on the Michelson leakage from a single arm of a polarisation Sagnac interferometer. 104
6.5 Relative levels of speedmeter and Michelson signals at the output of a polarisation Sagnac speedmeter as a function of the rotation of the QWP's in the two arms. 105
6.6 Contour plot of Michelson output signal for a single arm of a polarisation Sagnac speedmeter as a function of cavity mirror birefringence and relative alignment. 106
6.7 Effect of cavity birefringence on the amount of light leaking from a single cavity of the polarisation Sagnac speedmeter. 107
6.8 Combined effect of cavity and QWP birefringence on the amount of light leaking from a single cavity of the polarisation Sagnac speedmeter. 108
6.9 Contour plot of Michelson output signal for a single arm of a polarisation Sagnac speedmeter with measured QWP parameters as a function of cavity mirror birefringence and relative alignment. . . 110
6.10 Modelled polarisation at the output of a Michelson interferometer with birefringent cavity mirrors. 112

8.1 Basic operating principle for a four-level laser. 118
8.2 Operating principle of an electro-optic modulator as a phase compensator. 123
8.3 Operating principle of an acousto-optic modulator as an intensity or frequency modulator. 125
8.4 Visualisation of the generation of the Pound-Drever-Hall error signal from the phase modulated field reflected from a linear reference cavity. 129
8.5 Example of Pound-Drever-Hall error signal. 131
8.6 Overview of the Advanced LIGO PSL laser amplification and stabilisation scheme. 133
8.7 LIGO front-end laser. 133
8.8 Stabilisation requirement of the ISS for Advanced LIGO. 134
8.9 Stabilisation requirement of the FSS for Advanced LIGO. 136

8.10 Frequency reference cavity. 137
8.11 Premode cleaner cavity. 138

9.1 Predicted strain sensitivity of Advanced LIGO for different input laser powers. 143
9.2 Laser amplifier for Advanced LIGO. 144
9.3 Photograph of amplifier unit in place on the LLO PSL table with the beam path overlaid. 145
9.4 Section of the PSL table layout schematic showing the alignment path of the 70 W beam to the premode cleaner. 148
9.5 Output power of the laser amplifier as tested in the TnT lab. . . . 149
9.6 Output power of the laser amplifier after installation in the LLO PSL. 150
9.7 Sensitivity of the Advanced LIGO detectors with different injected laser powers. 151

10.1 Schematic of the Advanced LIGO table-top frequency stabilisation scheme. 153
10.2 Open loop gain measurements of LIGO table-top frequency stabilisation servo. 156
10.3 TTFSS PC path. 156
10.4 Transfer function breakdown of PC path output stage. 157
10.5 ISS diffracted power during observation. 162
10.6 Cumulative sum of lock losses that could have been caused by a dropout in the ISS diffracted power for the duration of O2. 163
10.7 Open loop gain of the inner ISS loop. 164
10.8 Catalogue of ISS dropouts during 7 days in February 2017. 166
10.9 Catalogue of ISS dropouts during 10 days in June 2017. 167

11.1 Overview of Glasgow 10 m prototype vacuum system. 170
11.2 External cavity diode laser operating principle. 171
11.3 NuAmp fibre laser amplifier internal workings. 172
11.4 NuFern fibre amplifier (a) and Pure Photonics diode laser (b). . . 173
11.5 Cross-section view of the two types of polarisation maintaining fibre. 174
11.6 Laser frequency noise of the 1550 nm laser compared to a NPRO seed laser. 177
11.7 Laser intensity noise of 1550 nm laser system. 178
11.8 Long term power output of 1550 nm laser system. 178

11.9 Rigid reference cavity for the frequency stabilisation system. Side-view of the reference cavity construction (a). Reference cavity spacer in the vacuum system (b). 180
11.10Frequency stabilisation scheme for the 1550 nm laser system. . . . 181
11.11Simulated open loop gain of the feedback paths to the laser PZT (purple) and to the reference cavity PZT (blue). The shows the case with a crossover just above 0.5 Hz which gives a gain separation of 70 dB at 100 Hz between the two feedback paths. 182
11.12Response of cavity (a) as a frequency reference and the PZT (b) and EOM (c) frequency actuators. 184
11.13Frequency stabilisation servo circuit diagram for the 1550 nm laser system. The input stage of the circuit includes a transitional integrator that in combination with a passive RC low pass filter forms a 3 kHz low pass filter (this design is explained in Section 11.2.4). The PZT feedback path includes a 12 Hz low pass filter and switchable inverting stage. The EOM feedback path includes a transitional integrator and a high voltage driver stage included in a separate box. 185
11.14Transfer function of the analogue servo for laser stabilisation. . . . 185
11.15Gain of the input stage of the analogue servo for high and low gain settings. 187
11.16Open loop transfer function measurement of the laser locking loop. 189
11.17Open loop transfer function measurement of the laser locking loop between 25 kHz and 32 kHz showing mechanical resonances of the laser. 190
11.18Error point noise in-lock with different gain setting of the analogue servo with OP27 (top) and TL071 (bottom) op-amps in the PZT feedback path. 193
11.19In-loop transfer function measurements of the 1550 nm laser locked to the reference cavity. 194
11.20Modelling the laser stabilisation servo from the LISO simulation of the servo circuit and the known response functions. 195
11.21In-loop measurement of stabilised laser frequency noise with respect to the free running laser frequency noise. 197

12.1 Block diagram of a basic control system with negative feedback. . 199
12.2 Block diagram of an implementation of a digital filter as a second-order section in the 'direct-form-II' layout of filter parameters . . 200

12.3 Representation of the aliasing problem with digital signal sampling. 201
12.4 Layout of digital frequency stabilisation scheme. 203
12.5 Transfer functions of basic filters for the digital laser frequency stabilisation system. 204
12.6 Digital filter design for frequency stabilisation servo. 206
12.7 Open and closed loop transfer functions of laser locked to reference cavity with basic digital filter (blue) and a two stages of bi-quadratic filters to increase the gain below 1 kHz (green). . . . 207
12.8 Measured laser noise with laser locked to reference cavity with basic digital filter (blue) and with two bi-quadratic filters (red). 208
12.9 Bode plots of randomised filters. 209
12.10Example of an optimised filter generated for one pole and one zero complex conjugate pair open loop transfer function. 211
12.11Example of an optimised filter generated for one pole and one zero complex conjugate pair pole-zero map. 212
12.12Example of an optimised filter generated for four pole and four zero complex conjugate pairs. 213
12.13The inner workings of the particle swarm optimisation showing pole-zero frequency and Q-factors for an example optimisation process. 214
12.14Particle swarm optimisation showing pole maps for an example optimisation process 215
12.15Particle swarm optimisation showing zero maps for an example optimisation process 215
12.16Weighted figures of merit (ξ) used in the filter optimisation metric showing the upper and lower limits of the four parameters and the related figures of merit. 217
12.17Measured transfer function of optimised digital filter compared to modelled result. 218

A.1 Birds-eye view of the Glasgow 10 m prototype vacuum system. The polarisation cavity experiment used the tanks indicated with horizontal and square shading. 223
A.2 Panoramic view of the Glasgow 10 m prototype vacuum system. The open tank on the left is the ETM tank of the polarisation Sagnac experiment, and the beam tube seen on the right is the arm of the 10 m cavity. 223

A.3 Side-view of seismic isolation stack for in-vacuum breadboards for the Glasgow 10 m prototype constructed from aluminium extrusions from the Bosch company. 224
A.4 Picture of the seismic isolation of the optical breadboard. 224
A.5 Picture of the suspension during construction. 226
A.6 Suspension modes of motion 227
A.7 Picture of the suspension actuation at the upper stage. 228
A.8 Full layout of the 1550 nm laser stabilisation and 10 m suspended cavity with polarising optics experimental set-up. 229

B.1 Circuit diagram of the PZT driver circuit designed to send signals to the reference cavity PZT for length actuation. 231
B.2 Circuit diagram of the PZT driver circuit designed to send signals to the internal frequency actuator of the 1550 nm laser. 231
B.3 Calibration of characterisation PDs in transmission (blue) and reflection (red) of test optic for the high and low gain transimpedance settings. 232
B.4 Circuit diagram of the reference cavity locking photodiode. 233
B.5 Circuit diagram of the photodiode in the reflection path of the 10 m suspended cavity. AC coupled output is used to generate the cavity locking error signal. 233
B.6 Circuit diagram of the photodiodes in the transmission path of the 10 m suspended cavity. 234
B.7 Calibration of the two photodiodes used in transmission of the 10 m cavity as part of the ellipticity readout scheme. 234
B.8 Circuit diagram of the notch filter design to filter out the 12.08 MHz local oscillator signal from the 1550 nm frequency stabilisation system. 235
B.9 Simulated response of notch filte designed for a 50 Ω load impedance. The grey trace is the simulated notch filter at 12.08 MHz and the black trace is the modelled response of the circuit design presented in Figure B.8 235
B.10 Design of an anti-imaging/anti-aliasing filter for the dSPACE digital process with a sample rate of 166.7 kHz. 236
B.11 Circuit diagram of the filter with the used component values. Parallel combinations of resistors were used to produce values of resistance closer to the desired values than the available component values allowed for. 237

B.12 Measured transfer function of the anti-aliasing/anti-imaging filter designed for the digital feedback system using the circuit design presented in Figure B.11 238

D.1 Block diagram of Simulink® model for frequency stabilisation courtesy of B. Barr including five second-order sections for shaping of the feedback signal. The model also includes a signal generator and triggering for auto-locking. 245

D.2 User interface for the digital laser stabilisation servo 246

Acknowledgements

I would like to thank my supervisor Ken Strain for his support and teaching throughout this work. Without his guidance and knowledge this would not have been possible.

Special thanks to Bryan Barr for constant guidance and help during my thesis and willingness to explain things to me again and again, and for all the help in and out of the lab and for much proof-reading.

I would also like to thank Stefan Hild, who acted as second supervisor during my time as a PhD student but also was my project supervisor during my undergraduate degree and who is responsible for getting me involved and interested in gravitational wave research. Thanks also to Sebastian Steinlechner who also supervised my masters project and introduced me to lab work.

Thanks must also go to the members of the Piledriver who have helped me get through this process by providing much needed sanity preserving social distractions.

Preface

This thesis covers a series of experimental work carried out at the University of Glasgow in the 10 m Interferometry Prototype laboratory between October 2015 and March 2019, with some work carried out during a five month visit to the LIGO Livingston detector in 2017. The work was carried out by the author with contributions from others as described in this section.

In Part I of this thesis, introductory material relating to interferometry for gravitational wave detection (Chapter 1) and quantum noise (Chapter 2) is discussed.

In Part II of this thesis an experimental investigation of the viability of the polarisation Sagnac speedmeter scheme as a future gravitational wave detector is presented.

Experimental work presented in Chapters 3-5 and 11-12 was conducted with the help and supervision of B. Barr.

In Chapter 3 polarising optics are characterised with table-top experiments to measure the extinction ratio of a polarising beam-splitter and the birefringence of a quarter-waveplate optic. This work was preceded by two student projects conducted by J. Bannigan and G. McGhee to characterise polarising optics, supervised by the author along with B. Barr and K. Strain. For this thesis new data was taken characterising different optics with a different measurement set-up but utilised the experience gained from the previously undertaken projects.

Chapter 4 contains a description of the installation of hardware and the commissioning of a linear cavity experiment to measure birefringence in a 10 m long suspended optical cavity as part of the experimental validation of the polarisation Sagnac concept. The suspensions used for this experiment were constructed by the author and B. Barr to a design specification courtesy of R. Jones and N. Gordon. The modifications to the suspension deisgns for the polarisation optics were done by R. Jones. The Simulink® control model was originally built by B. Barr.

Chapter 5 contains a description of the analysis of the results from the linear

cavity experiment and in Chapter 6 these results are used to model the impact on the polarisation Sagnac speedmeter of the results from Chapters 3-5. This work, and the outlook for speedmeter interferometry, is summarised in Chapter 7.

Finally in Part III, topics in laser development for current and future gravitational wave detectors are covered. This includes work on the Advanced LIGO laser system to install a new laser amplifier and to improve the laser stabilisation during operation. This part also includes work to install and frequency stabilise a new 1550 nm laser system for the Glasgow 10 m prototype facility.

In Chapter 8 the field of laser technology for gravitational wave detection is summarised as an introduction to Chapters 9-12. This includes a description of the Advanced LIGO laser system.

The work in Chapters 9 and 10 was carried out over the course of two visits to the LIGO Livingston Observatory under the LSC Fellows programme and Caltech Visitors programme.

In Chapter 9 an upgrade of the existing LIGO laser system with the aim to double the optical power available to the interferometer is described. This equates to a significant sensitivity improvement and was a priority upgrade for the LIGO detectors between the second and third observing runs of Advanced LIGO. The work presented here reflects the contribution made to the hardware testing and installation as supervised by M. Heintze alongside T. Hardwick and M. Fyffe.

The work in Chapter 10 describes two investigations into the LIGO laser stabilisation system and the impact on the detector operations during the second observing run of Advanced LIGO. This work was carried out during visits to the LIGO Livingston site and while at the University of Glasgow with help and guidance from M. Heintze, B. O'Reilly, A. Mullavey, M. Laxen and C. Adams from LIGO Livingston and K. Strain at the University of Glasgow. Chapter 11 covers the work carried out at the Glasgow 10 m prototype to implement a new laser system with an output wavelength of 1550 nm for low-noise interferometry experiments and the frequency stabilisation scheme developed for this laser.

In Chapter 12 this work is further extended with the implementation of a high-speed digital servo for laser frequency stabilisation. The dSPACE control model was originally written by B. Barr.

This PhD was funded by the STFC Studentship award number 1653071.

Declaration

I hereby declare that all work presented in this thesis was carried out by the author unless otherwise explicitly stated. This work has not been presented in any previous application for a degree at this or any other institution.

Part I

Introduction to Interferometry for Gravitational Wave Detection

Chapter 1

Interferometry for Gravitational Wave Detection

Optical interferometers are used across many fields of science as precision measurement devices. In gravitational wave astronomy these devices are used to sense perturbations in the relative separation of a set of test masses as caused by the passage of a gravitational wave.

To generate a gravitational wave of sufficient magnitude as to afford the possibility of detection requires the occurrence of cataclysmic astrophysical events. A gravitational wave detector must be extremely sensitive as gravitational waves interact weakly with matter

In this chapter the field of gravitational wave detection is reviewed to motivate the topic of laser interferometry for gravitational wave detectors.

1.1 Gravitational Wave Detection

Gravitational waves are generated by any non-axisymmetric rotating massive objects. As the massive system rotates the gravitational distortion of space-time causes waves to be radiated. These waves travel at the speed of light outward from the source and can be observed as ripples in the local space-time of a detector [1].

In a gravitational wave detector (as shown in Figure 1.1) laser light is split into two beams that are each reflected from two spatially separated test masses. The orthogonal optical paths between where the beam is split and where the beams are reflected from the test mass form two 'arms' of the detector. When a gravitational wave interacts with a detector it alternately stretches and squeezes the space-time between the detectors' test masses. This effective change in the

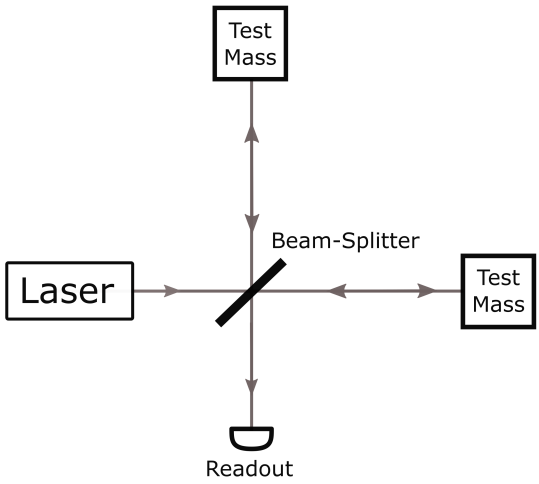


Figure 1.1: Basic model of a gravitational wave detector as a pair of test masses probed by a laser beam split at a beam-splitter. The beams returning to the beam-splitter are recombined and measured at the readout.

length of the detector arms as seen by a beam of light can be measured with an interferometric readout.

The first direct detection of a gravitational wave signal was made in September 2015 by the Advanced LIGO observatories [2–4]. Between 2015 and 2017, ten more gravitational wave signals were detected, beginning the era of gravitational wave astronomy [5].

The source of all these detections was the merger of heavy, compact astrophysical objects. Ten of these detections, including the first, were signals from binary black hole mergers. In August 2017 the merger of two neutron stars was detected as a gravitational wave event coincident with a gamma ray burst [6]. This event was followed by electromagnetic detection of the afterglow of the merger. This marked the first ever joint detection of an astrophysical event from gravitational and electromagnetic signals [7].

When two heavy compact objects orbit each other in a binary system, gravitational waves are generated, radiating energy outwards. As energy is lost from the orbiting system the radius of orbit decreases and the two objects inspiral closer and closer until the two collide. In the case of a black hole binary the two black holes merge to form a single singularity. In the case of two neutron stars the merger is extremely violent and results in an explosive collision accompanied by the ejection of a large amount of matter and light.

The catalogue [5, 9] of gravitational wave sources detected between 2015 and 2017 is shown in Figure 1.2. The pre-merger compact binaries and the post-merger remnants are shown in terms of their masses. The final state of the binary neutron

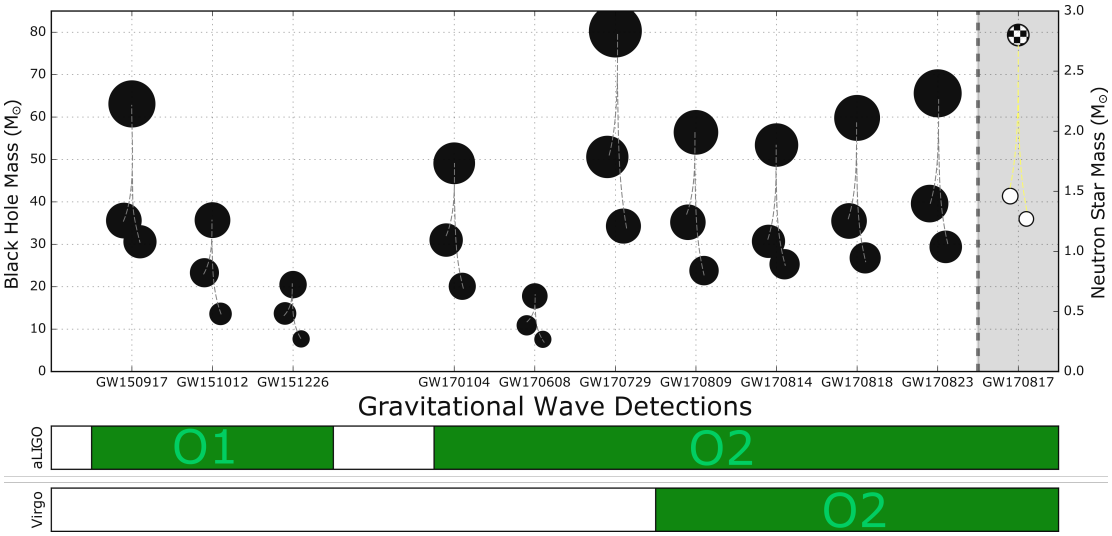


Figure 1.2: Representation of the binary merger masses of all gravitational wave detections made during the first [8] and second observing runs [9] of the advanced gravitational wave detectors. Source masses are shown on the left axis for black holes and on the right for neutron stars. The operational status of the LIGO and Virgo detectors is shown in the green bars below the plot.

star post-merger is not known (so is displayed partially shaded).

These detections constitute the beginning of gravitational wave observation as a branch of astronomy, opening a new way of observing the Universe.

1.2 The Gravitational Wave Detector Network

The first network of long-baseline ground-based gravitational wave interferometers was known as the first generation detector network. These detectors established the infrastructure of gravitational wave observatories and over the course of several science runs incrementally improved in sensitivity. The so called 'advanced' detectors improved the sensitivity of the initial detectors to the point where detection was possible. It was with the second generation of advanced detectors that the first direct observation of gravitational wave signals was made. The third generation of gravitational wave detectors will further improve the rate of detection, the range, and frequency bandwidth of the instruments. These detectors are currently being planned while the second generation of detector continue to operate and make astrophysical observations.

LIGO

The Advanced LIGO gravitational wave detectors are located at two observatory sites in the USA: in Hanford, Washington and in Livingston, Louisiana. The Advanced LIGO detectors are of a dual recycled Fabry-Pérot Michelson design with 4 km long arms. [10, 11]

The initial LIGO detectors were first generation gravitational wave detectors operated between 2002 and 2010. Over the course of several science runs the sensitivity was incrementally increased but did not reach the level required for detection. Between 2010 and 2015 a series of hardware upgrades were commissioned to upgrade the instrument to what is known as the Advanced LIGO detector which began observing runs in 2015 with an improved sensitivity that led to the first detection of gravitational waves. The Advanced LIGO upgrade [12] included a replacement of the core optical suspensions with heavier mirrors and monolithic fused silica suspensions [13] to improve the isolation of the optics and thermal noise performance (see Section 1.3). Advanced LIGO also included the additional of signal recycling [14] to the already power recycled Fabry-Pérot Michelson interferometer and an increase in the injected laser power [15].

After the end of the current schedule of observing runs is completed, there are plans in place to further improve the Advanced LIGO detectors. A series of 'soft' upgrades called Advanced LIGO+ would improve the sensitivity by a factor of two by utilising improved techniques that have been demonstrated since the commissioning of Advanced LIGO while still using the majority of the Advanced LIGO hardware and infrastructure.

In addition, there are plans to construct a third LIGO detector in India [16], this would have the same design as the Advanced LIGO detectors and would improve the source sky localisation and network up-time [17].

A more ambitious upgrade named LIGO Voyager [18] is under consideration. This would require a major infrastructure change as it includes the cryogenic cooling of the test mass suspensions. The move to cryogenic optics will require a change in the optic substrate and coating material and the laser system.

Silicon optics are begin researched for cryogenic interferometry as fused silica (that is currently used) has poor thermal noise properties at low temperatures. Using silicon optics will require a laser system with longer wavelengths as the current laser systems operating at a wavelength of 1064 nm are incompatible with silicon optics. High power stabilised laser systems at 1550 nm and 2 μm are being investigated for the LIGO Voyager project.

Virgo

The Virgo detector is a 3 km long dual recycled Fabry-Pérot Michelson detector located in Cascina, Italy [19]. The Virgo detector features a complicated suspended optic system that in addition to the final monolithic test mass suspension stage, features additional stages of inverted pendulum attenuation. This system (dubbed the 'superattenuator') suppresses the environmental seismic noise coupling to the test mass positions [20].

The Advanced Virgo detector joined the second observing run of the gravitational wave detector network and made its first detection on the 14th of August 2017 [21] when the signal from a binary black hole merger was detected in coincidence with the two LIGO detectors. A few day later, on the 17th of August, the signal from a binary neutron star merger [6] was detected for the first time. The inclusion of Virgo in the detector network allowed for accurate sky localisation of the source such that an alert could be sent to electromagnetic telescope astronomers.

GEO

The GEO detector is a 600 m long detector located near Hanover in Germany [22]. The GEO detector was run in operation all through the first and second observing runs of the advanced detectors and maintains a higher up time percentage than the LIGO and Virgo detectors. Only having a 600 m infrastructure baseline means the GEO detector is limited in the level of sensitivity it can achieve relative to the longer baseline detectors. However, the detector continues a program known as 'astro-watch' to run in observing mode while the LIGO and Virgo detectors are off-line to collect science data and maintain the network in the event of an extremely loud event.

The current iteration of the GEO detector is referred to as GEO-HF as it is optimised for high frequency sensitivity. The detector is a dual recycled Michelson interferometer without Fabry-Pérot arm cavities as this would reduce the sensitivity bandwidth. The detector topology includes a folding mirror at the arm ends to allow for a 1.2 km arm length.

The GEO detector has maintained a place as a technology demonstrator by implementing advanced techniques. The GEO detector was the first to implement signal recycling in conjunction with power recycling [23] and was the first detector to use squeezed light injection [24]. The GEO detector demonstrated the possibility of long term squeezed light injection [25] and the squeezer system was later adopted by the Virgo detector.

KAGRA

The KAGRA detector, currently being commissioned in Japan, is the first cryogenic underground gravitational wave detector. The vacuum system infrastructure is constructed in the Mozumi mines in Gifu, Japan [26]. The site seismic noise is lower than on the surface [27]. The cryogenic cooling of the single crystal sapphire test mass mirrors also lowers the thermal noise of the detector. When completed, the KAGRA detector will be a dual recycled Fabry-Pérot Michelson with 3 km arm cavities [28].

The plan is for the KAGRA detector to join the advanced detector network during the third observing run and then to improve the sensitivity towards full operating design sensitivity in-line with the other detectors.

Einstein Telescope

The proposed Einstein Telescope is a mature concept for a third generation gravitational wave detector to be built in Europe [29] with the aim to be ten times more sensitive than the design sensitivity of the advanced detectors.

The Einstein Telescope (ET) will be a 10 km long interferometer built underground to reduce the seismic noise limit [30].

The ET design features three detectors with a 60° opening angle between arms in an overlapped triangle layout to allow for sky localisation from a single site detector.

The design proposed for ET features a nested set of interferometers referred to as the xylophone design [31]. The idea is to build interferometers of different designs optimised for improved sensitivity across different frequency bands. The low frequency detector (ET-LF) will include cryogenic mirror suspensions and the high frequency detector (ET-HF) will use a high circulating laser power and room temperature optics.

Including the high and low frequency interferometers means the ET observatory will feature six interferometers on a single site.

1.3 Noise Sources in Gravitational Wave Detectors

The key fundamental noise sources that limit the design sensitivity of Advanced LIGO are shown in Figure 1.3. Quantum noise is covered in detail in Chapter 2. In this section the other noise sources shown in Figure 1.3 are described. A

review of the limiting noise sources for Advanced LIGO at the time of the first detection can be found in [3].

Seismic Noise

To isolate the test masses from non-astrophysical disturbances the interferometer is operated under ultra-high vacuum to reduced acoustic noise coupling and optical scattering. The test masses themselves are hung on multi-stage pendulums which are in turn mounted on seismic isolation platforms to decouple the motion of the mirror from the motion of the ground [32].

The motion of the test masses must be ten orders of magnitude lower than the surface ground motion. To achieve this level of isolation a combination of active seismic isolation between the vacuum tanks and the ground and between the in-vacuum optical tables and the vacuum tanks is implemented. The Advanced LIGO actively controlled seismic isolation platforms provides three orders of magnitude of suppression of the ground motion.

The Advanced LIGO quadruple suspensions [13,33] provide suppression of the ground motion coupling to the test mass by seven orders of magnitude. The suspended test masses act as quasi-free masses at frequencies above the fundamental pendulum frequency. For a single pendulum the coupling of external motion above the pendulum resonance to the motion of the lower stage rolls off as $1/f^2$ and for a quadruple pendulum the isolation has a $1/f^8$ law.

The residual coupling of seismic noise to the gravitational wave detection band is seen in the brown trace in Figure 1.3. This is passively filtered by the suspension system as well as being passively and actively filtered by the seismic isolation system.

Thermal Noise

The source of thermal noise is Brownian motion within the optical and suspension materials that couples to motion of the optic through mechanical dissipation in the material [35].

The mirrors used in gravitational wave detectors are made of stacks of alternating layers of high and low refractive index material deposited onto a large fused silica substrate. The optics are suspended in a pendulum system using thin fused silica fibres that are bonded to the sides of the substrate.

Mechanical loss in the suspension fibres [36] and in the coating layers [37], are the dominant sources of thermal noise that affect gravitational wave detectors.

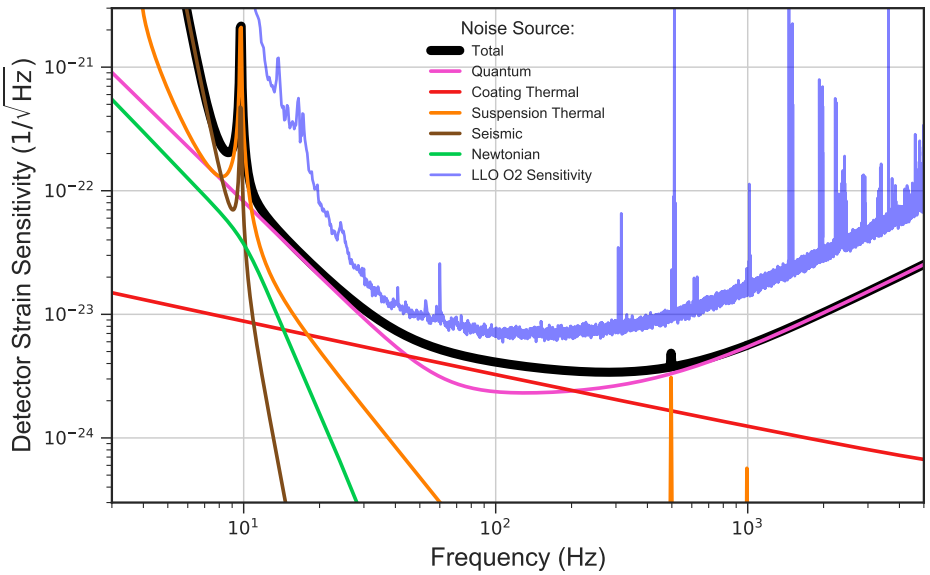


Figure 1.3: Modelled strain sensitivity of Advanced LIGO and principal noise sources. This plot was created with the interferometer modelling software GWINC [34]. The sensitivity of the LIGO Livingston detector during the second observing run is shown in the blue trace when the detector was run with lower laser power injected into the interferometer and below 50 Hz is limited by technical control noise so was not at design sensitivity. Suspension thermal noise (orange) and seismic noise (brown) limit the sensitivity at the low frequency end of the band, above 10 Hz quantum noise (pink) limits the design sensitivity and between 30 Hz and 500 Hz, a combination of coating thermal noise (red) and quantum noise limit the sensitivity.

The test mass suspension system described above is necessary to isolate the test mass motion from ground motion, but the design of the suspension has to be considered carefully as the Brownian motion of molecules couples to a length noise of the test mass cavity though a mechanism referred to as suspension thermal noise [38]. It was found that this noise source could be reduced to below the requirement for a gravitational wave detector with the use of monolithic fused silica suspensions. In such a system the suspension fibre material is the same low thermal noise material as the test masses [39, 40].

The multi-stage pendulum also acts to filter the suspension thermal noise, so above the fundamental pendulum resonance frequency the noise drops by a factor of $1/f^2$ per pendulum stage. The suspension will also have resonances in all six degrees of freedom of the mass (see Section A.3). The frequency of these resonances depends on the design of the suspension which is carefully chosen to reduce the presence of resonance peaks in the gravitational wave detection band of the instrument. Only the resonances at 10 Hz, 500 Hz and 1 kHz can be seen in the suspension thermal noise (orange trace) in the strain sensitivity plot shown in Figure 1.3.

The most limiting thermally driven noise source in gravitational wave detectors is coating thermal noise. Brownian motion in the coating of the optics couples to length noise of the cavity as the surface of the optic is perturbed by the thermally driven oscillation of the coating material's atoms.

The material choice for the multi-layer mirror coatings is important to achieve low thermal noise and high reflectivities. Higher reflectivities can be achieved with more layers and with greater difference between the two material refractive indices. The methods for depositing the coating material onto the mirror substrates also greatly affects the overall thermal noise performance of the mirror [41–43].

Newtonian Noise

Newtonian noise is the term given to the coupling of changes in the local gravitational field to the position of the detectors test masses. Transients in the local gravitational field exert a changing gravitational (or 'Newtonian') force on the suspended test masses which leads to length noise of the detectors arms that could mask gravitational wave signals at low frequencies [44].

Newtonian noise sources include seismic activity that changes the local mass density of the ground surrounding the detector, atmospheric density fluctuations and human activity close to the detectors [45].

1.4 Mathematical Description of Light

Gravitational wave interferometers use light to probe two test masses and extract information on the motion of the test masses as induced by gravitational waves. This section contains a detailed description of light that is then used in Section 1.5 to describe the basics of interferometry for gravitational wave detection.

In the most general case, a light field is described as a propagating electromagnetic wave with the mathematical form

$$\vec{E}(x,y,z,t) = E_0\,\vec{e_p}\,U_{nm}\cos(\omega t - \vec{k}\cdot\vec{r} + \varphi) \tag{1.1}$$

where E_0 is the field amplitude and ω is the angular frequency. The vector $\vec{k}$ is refered to as the wavenumber of the wave in direction $\vec{r}$ and φ is a phase term. The vector $\vec{e_p}$ described the polarisation state of the field and U_{nm} describes the cross-sectional profile of the light field (these properties are ignored for the moment and will be discussed in more detail in Sections 1.4.2 and 1.4.3). For the remainder of this description only the steady state form of an optical field propagating in direction z is considered, and the complex form of the oscillator is used so it can more simply be written as

$$E = E_0\exp(-ikz + \varphi). \tag{1.2}$$

1.4.1 Interaction of Light and Optics

With this simple description of light as an oscillating field we can calculate the propagation of light through any number of optical systems. The most basic example is the interaction of a light field with a reflective optic as shown in Figure 1.4. The optic is described by coefficients of reflection ρ and transmission τ where in the case of no optical loss, $\rho^2 = 1 - \tau^2$, and so;

$$\vec{E_1} = \rho E_{In}, \tag{1.3}$$

$$\vec{E_2} = i\tau E_{In}. \tag{1.4}$$

By convention, the reflectivity and transmissivity coefficients of individual optics are represented by upper case letters for optical powers and by lower case Greek letters for amplitudes, so $R_j = \rho_j^2$ and $T_j = \tau_j^2$.

In this formalism the relative phase between reflected and transmitted optical fields is expressed by the factor of i in the expression for the transmitted field; this convention is explained and justified in [46].

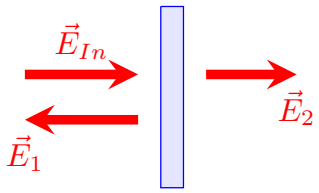


Figure 1.4: Basic interaction of light field with a partially reflective mirror.

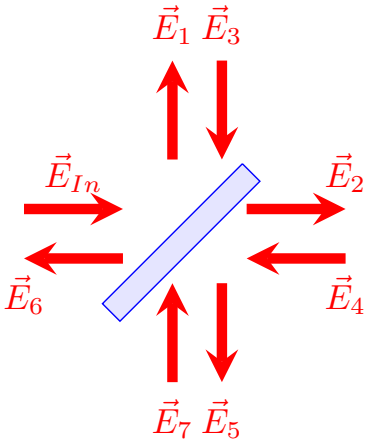


Figure 1.5: Basic interaction of light field with a beam-splitter with four ports.

For an optic with four ports, i.e. a beam-splitter (see Figure 1.5), the same convention is used such that;

$$\vec{E}_1 = \rho E_{In} + i\tau E_7, \tag{1.5}$$
$$\vec{E}_2 = i\tau E_{In} + \rho E_7, \tag{1.6}$$
$$\vec{E}_5 = i\tau E_3 + \rho E_4, \tag{1.7}$$
$$\vec{E}_6 = \rho E_3 + i\tau E_4. \tag{1.8}$$

Light propagating over a length L in the direction z through a medium of refractive index n, will acquire a phase of $\varphi = -knL$, where $n = 1$ for a vacuum and $k = \omega/c$. The negative sign in the expression of phase means that we consider a field that has propogated over a distance to have aquired a phase delay relative to the initial field.

The wavelength of a laser source is taken to be unchanging so it often convenient when considering a wave traveling over distances much longer than its wavelength, to consider a macroscopic length L' equal to the closest integer multiple of the wavelength and a microscopic detuning length δL such that $L = L' + \delta L$.

1.4.2 Polarisation

The polarisation of a light field refers to the orientation of the electric and magnetic field vectors relative to the propagation direction of the electromagnetic radiation. Conventionally it is the electric field that used to define the polarisation direction. For an electromagnetic wave propagating in free space the electric and magnetic fields are orthogonal. If the field is unpolarised, the field vector has a random orientation transverse to the propagation vector.

Consider a light field propagating in the z direction. The projection of the electric field in the x and y dimensions of a right-handed co-ordinate system define the polarisation state in relation to the surface of an optic. If the electric field oscillates only in the y dimension along dimension z then the light is vertically polarised (referred to as s-polarised). By the same convention, p-polarised light has an electric field oscillating in the x dimension and is said to be horizontally polarised.

The projection of the oscillating field in the $x - y$ plane is used to categorise the polarisation state more generally. If the field traces out a line in this plane, then the polarisation is said to linear. If the field traces an ellipse or a circle, then the light is said to be elliptically or circularly polarised [47].

Polarised light can be described mathematically using Jones Calculus [48], where any polarised light field is represented by a 2x1 Jones vector. Optical elements are represented as 2x2 Jones matrices, and the interaction of light with optics can be calculated as the multiplication of matrices and vectors.

Taking a general light field described by the oscillating electric field in Cartesian coordinates as

$$\begin{bmatrix} E_x(t) \\ E_y(t) \\ 0 \end{bmatrix} = \begin{bmatrix} E_{0x}\exp(i\phi_x) \\ E_{0y}\exp(i\phi_y) \\ 0 \end{bmatrix} \exp\Big(i(kz - \omega t)\Big), \tag{1.9}$$

then the general Jones vector is

$$\begin{bmatrix} E_{0x}\exp(i\phi_x) \\ E_{0y}\exp(i\phi_y) \\ 0 \end{bmatrix}. \tag{1.10}$$

It is convenient to use normalised Jones vectors in which case the polarisation states can be represented more simply. A general, linearly polarised field is

represented by the vector

$$\begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}, \tag{1.11}$$

and an elliptical field by the vector

$$\begin{bmatrix} \cos\theta \exp(i\phi_x) \\ \sin\theta \exp(i\phi_y) \end{bmatrix}. \tag{1.12}$$

In the special case when the amplitudes are equal and the phase difference $\phi_x - \phi_y = \pi/2$, then the light is circularly polarised.

From Equation 1.10 it can be seen that changes to the polarisation state can be caused by changes in relative phase in the x and y directions. A material or optic that causes a different phase change depending on geometric orientation is said to be birefringent. A birefringent optic can be described by the Jones matrix

$$\begin{bmatrix} \exp(i\phi_x) & 0 \\ 0 & \exp(i\phi_y) \end{bmatrix}. \tag{1.13}$$

A polariser has polarisation dependent transmission and reflection coefficients can also be described by a Jones matrix. For the trivial case of a linear polariser orientated to transmit light polarised in the x direction this matrix is

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}. \tag{1.14}$$

More detail of Jones calculus is given in Appendix C and in Chapters 3 and 6, where this framework is used to build an analytical model of an interferometer with polarising optics.

1.4.3 Gaussian Beams

In the previous section we showed the representation of a laser beam as a electromagnetic field propagating along an axis z. The transverse profile of the field was mentioned but disregarded at that stage as we only considered what is known as the plane wave approximation. This approximation is valid for most analysis of light propagation but in some cases the geometric shape of the beam needs to be considered. This is done by summing all the spatial modes of the field. For a laser beam these spatial modes follow a Gaussian profile. We can consider,

$$E(x, y, z) = \exp(-ikz) \sum_{n,m} a_{nm} u_{nm}(x, y, z), \tag{1.15}$$

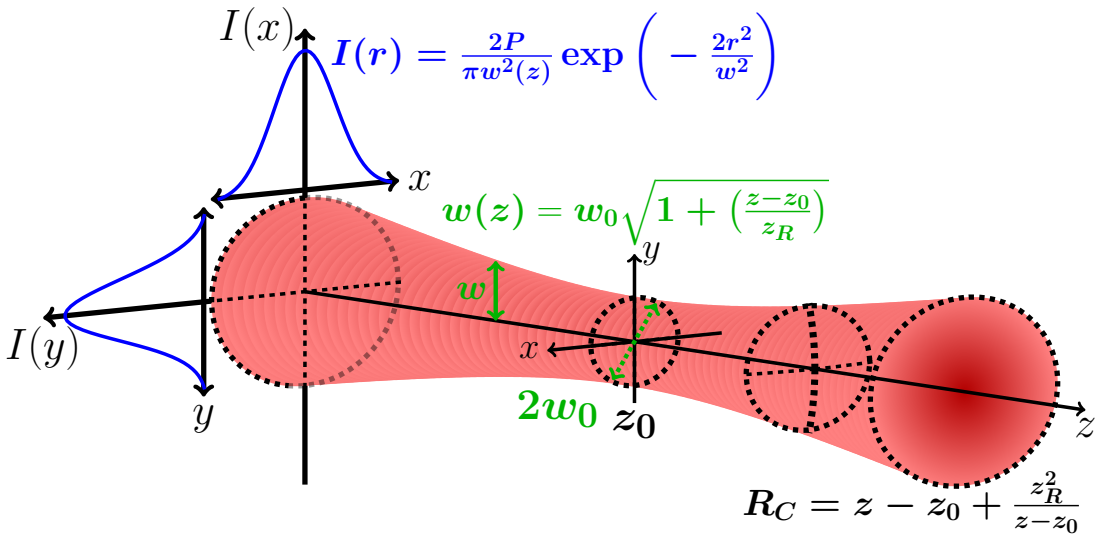


Figure 1.6: Representation of a laser as a Gaussian beam propagating in the z direction. The beam radius is defined by the equation in green. The cross-sectional intensity profile in x and y is shown to the left in blue. The black dotted lines represent the wave fronts of the beam, contours with constant phase. The radius of curvature of these phase fronts is given by R_C.

where a_{nm} is a complex amplitude term and u_{nm} is the geometric function that describes the shape of one mode of a complete set of modes of the electric field. The integers n and m are the mode order numbers.

In this case $n = m = 0$, the field is said to be of the zeroth-order or the fundamental mode, and $u_{00}(x, y, z)$ can be written as

$$u(x, y, z) = \sqrt{\frac{2}{\pi}} \frac{1}{w(z)} \exp(i\Psi(z)) \exp\left(- ik \frac{x^2 + y^2}{2R_C(z)} - \frac{x^2 + y^2}{w^2(z)} \right). \tag{1.16}$$

This beam is characterised by a number of geometric parameters with a specific terminology common to descriptions of Gaussian beams [49]. This includes a factor $\Psi(z)$ called the Gouy phase, $w(z)$ the beam radius at point z. The parameter R_C is the radius of curvature of the wave fronts. The radius of the beam can be written as

$$w(z) = w_0 \sqrt{1 + \left(\frac{z - z_0}{z_R} \right)^2}, \tag{1.17}$$

where w_0 is called the waist and represents the smallest radius of the beam, z_0 is the position of the waist along the propagation axis and z_R is the Rayleigh range. The Rayleigh range defines the boundary between near field and far field

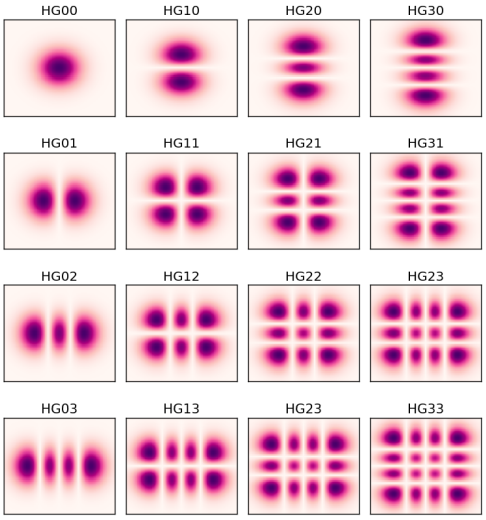


Figure 1.7: Transverse profiles of HG_{nm} laser beams of up to sixth order Hermite-Gaussian modes.

conditions. This length is specified as

$$z_R = \frac{\pi w_0^2}{\lambda}. \tag{1.18}$$

Fields can also exist as solutions to the wave equation with higher-order mode numbers $n + m > 0$. In this case the spatial parameter u_{nm} is defined using Hermite polynomials or Laguerre polynomials for Cartesian or spherical symmetry respectively. The sum $n + m$ is referred to as the order of the mode. The spatial mode shape is is given by

$$u_{nm}(x, y, z) = u_n(x, z)u_m(y, z), \tag{1.19}$$

where for a Hermite-Gaussian beam

$$u_n(x, z) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \left(\frac{\exp(i(2n+1)\Psi(z))}{2^n n! w(z)}\right)^{\frac{1}{2}} H_n\left(\frac{\sqrt{2}x}{w(z)}\right) \exp\left(-i\frac{kx^2}{2R_C(z)} - \frac{x^2}{w^2(z)}\right), \tag{1.20}$$

with $H_n(x)$ being the n^{th} order Hermite polynomial.

Laguerre-Gaussian beams are radially symmetric modes described by Laguerre polynomials $L_p^{(l)}(x)$ such that the spatial parameter can be written in terms of

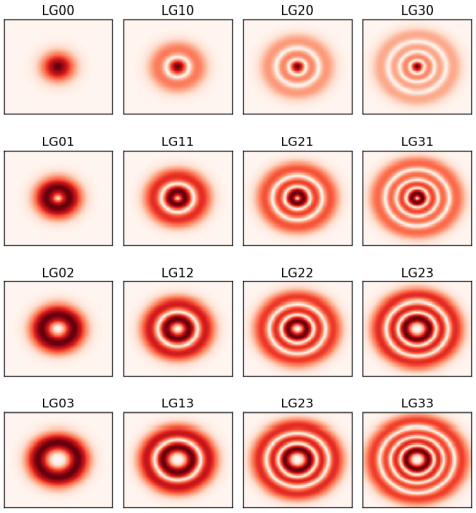


Figure 1.8: Transverse profiles of LG_{pl} laser beams of up to sixth order Laguerre-Gaussian modes.

cylindrical coordinates r, ϕ and z as

$$u_{p,l}(r,\phi,z) = \frac{1}{w(z)}\sqrt{\frac{2p!}{\pi(|l|+p)!}}\exp(i(2p+|l|+1)\Psi(z)) \times \left(\frac{\sqrt{2}r}{w(z)}\right)^{|l|} L_p^{(l)}\left(\frac{2r^2}{w(z)^2}\right)\exp\left(-ik\frac{r^2}{2q(z)}_i\, l\phi\right), \quad (1.21)$$

where q is the Gaussian beam parameter

$$\frac{1}{q(z)} = \frac{1}{R_C(z)} - i\frac{\lambda}{\pi w^2(z)}, \quad (1.22)$$

and Ψ is the Gouy phase

$$\Psi(z) = \tan\left(\frac{\mathrm{Re}(q)}{\mathrm{Im}(q)}\right)^{-1}. \quad (1.23)$$

The first 16 Hermite-Gaussian modes are shown in Figure 1.7 with Laguerre-Gaussian beam shapes shown in Figure 1.8.

A laser beam can be described as a combination of many Gaussian modes with the dominant mode typically being the fundamental mode with some higher order mode content.

The phase difference ϕ between a plane wave and a Hermite-Gaussian mode

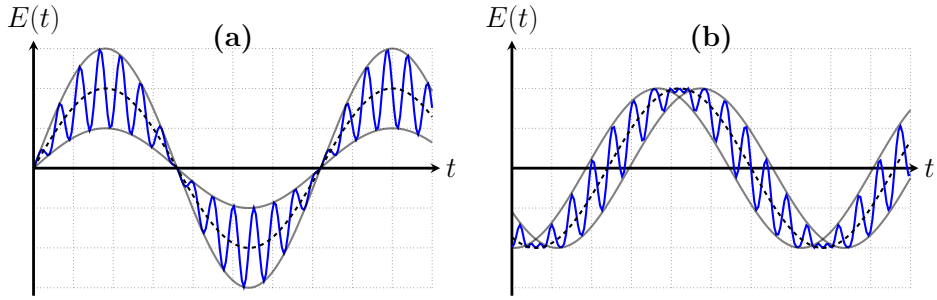


Figure 1.9: Sketches of amplitude (a) and phase (b) modulation of a carrier field (dashed line).

of order $n + m$ is given by

$$\phi = (n + m + 1)\Psi(z). \tag{1.24}$$

Therefore different order modes will have different phase evolutions along the propagation direction, this will result in different resonant conditions in optical cavities. This can be used to selectively improve the mode content of a laser beam by passing it through an optical cavity such that a single mode is resonant. Optical mode-cleaner cavities work on this principle.

1.4.4 Modulation

The next part of the mathematical toolbox needed to analyse interferometers is a description of the modulation of an optical field. Modulation is used to actuate the amplitude, phase or frequency of the laser as well as to control interferometric systems.

When discussing modulation of laser light, terminology common to radio engineering is used. The main optical field at frequency ω_0 is referred to as the *carrier*, and by modulating the carrier *sidebands* are added which are components of the field offset in frequency from the fundamental frequency.

Amplitude Modulation

An unmodulated optical field can be described by the expression

$$E = E_0 \exp(i\omega_0 t), \tag{1.25}$$

and an amplitude modulated field by

$$E = E_0 \exp(i\omega_0 t)\Big(1 + m\cos\omega_m t\Big), \tag{1.26}$$

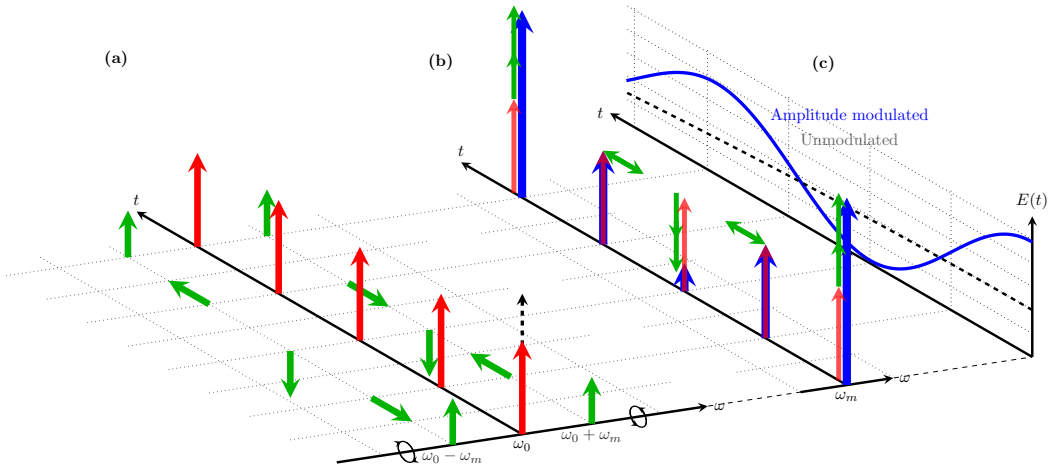


Figure 1.10: Representation of amplitude modulation of a carrier field (red) by the addition of sidebands (green) of first order only. In the reference frame of this diagram the carrier is stationary so the upper sideband which is a higher frequency to the carrier will rotated about the frequency axis with time, similarly the lower sideband being slower will rotate in the opposite direction. The result shown in (b) is the sum of the three fields (blue) which changes amplitude but not phase over time. The projection of the modulated field is shown in (c) which shows one quarter cycle of the unmodulated and modulated field with a modulation frequency tens times greater than the carrier frequency.

where ω_m is the frequency of the modulation and m is the modulation index. This can be written as

$$E = E_0 \exp(i\omega_0 t)\Big(1 + \frac{m}{2}\exp(i\omega_m t) + \frac{m}{2}\exp(-i\omega_m t)\Big). \tag{1.27}$$

Amplitude modulation imposes two sidebands on a carrier at frequencies above $(\omega_0 + \omega_m)$ and below $(\omega_0 - \omega_m)$ the carrier frequency. The phase of the sidebands relative to the carrier is as shown in Figure 1.10, such that at arbitrary starting point $t = 0$ the two sidebands and the carrier are all in phase. In the reference frame of the plot the carrier is stationary and the upper sideband rotates about the frequency axis faster than the carrier in a clockwise direction, while the lower sideband is slower so rotates anti-clockwise. In this way the sum of the three fields does not change in phase or frequency, but the amplitude changes over time.

Phase Modulation

Phase modulation of a field can generate multiple sidebands dependent on the modulation strength. Taking the same unmodulated input field defined in Equation 1.25 and modulating the phase results in a field described by the expression

$$E = E_0 \exp\Big(i(\omega_0 t + m\cos\omega_m t)\Big), \tag{1.28}$$

which can be expanded using Bessel functions of the first kind $J_k(m)$, to give

$$E = E_0 \exp(i\omega_0 t) \sum_{k=-\infty}^{\infty} i^k J_k(m) \exp(ik\omega_m t), \tag{1.29}$$

where,

$$J_k(m) = \Big(\frac{m}{2}\Big)^k \sum_{n=0}^{\infty} \frac{\Big(-\frac{m^2}{4}\Big)^n}{n!\,(k+n)!}. \tag{1.30}$$

Setting $m \ll 1$ to consider weak modulation to generate only first-order sidebands we can simplify this to

$$E = E_0 \exp(i\omega_0 t)\Big(J_0(m) - iJ_1(m)\exp(-i\omega_m t) + iJ_1(m)\exp(i\omega_m t)\Big). \tag{1.31}$$

with

$$J_{-k}(m) = (-1)^k J_k(m). \tag{1.32}$$

Approximating to first order using a Taylor expansion, a frequency modulated field can be given as

$$E = E_0 \exp(i\omega_0 t)\Big(1 + i\frac{m}{2}\Big(\exp(-i\omega_m t) + \exp(i\omega_m t)\Big)\Big). \tag{1.33}$$

In the phasor picture shown in Figure 1.11, phase modulation is represented by a carrier and two sidebands that at time $t = 0$ are exactly out of phase with each other. In the reference frame of the plot the carrier is stationary and the upper sideband rotates faster then the carrier, while the lower sideband is slower so rotates anti-clockwise. In this way the sum of the three fields is constant in amplitude but changes in phase.

Frequency Modulation

In the case of laser interferometry, we can consider phase modulation as a way of generating frequency modulation. For weaker modulation they can be considered as descriptions of the same effect as the frequency of the wave is defined as $f = \frac{d\varphi}{dt}$. To obtain a frequency modulation the phase needs to be $\omega_0 t + m\sin\omega_m t$, where the modulation index in this case is the change in frequency over the modulation frequency $m = \frac{\delta\omega}{\omega_m}$. The frequency modulated field can then be written as

$$E = E_0 \exp(i(\omega_0 t + m\cos\omega_m t)). \tag{1.34}$$

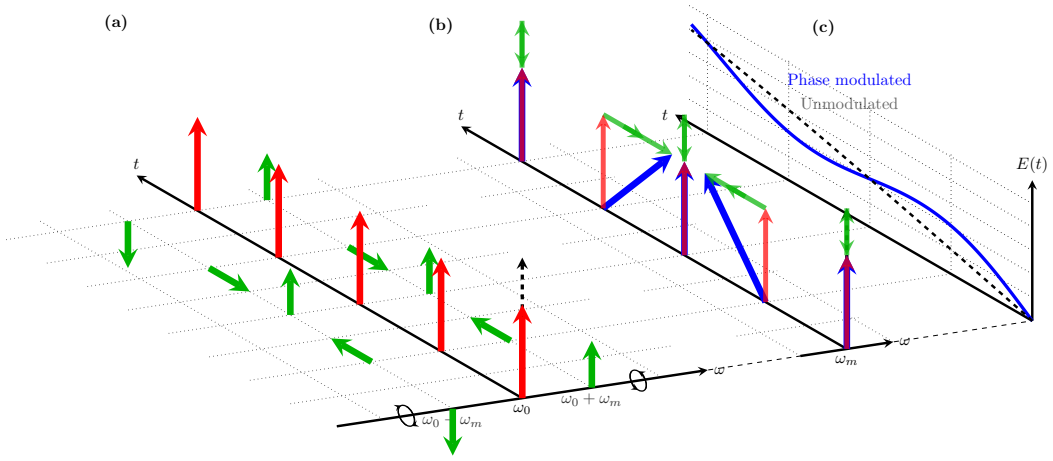


Figure 1.11: Representation of phase modulation of a carrier field (red) by the addition of sidebands (green) of first order. The height of the sidebands relative to the carrier is exaggerated for this visual demonstration, we only consider weak modulation that produces only first order sidebands. In the reference frame of this diagram the carrier is stationary so the upper sideband which is a higher frequency to the carrier will rotated about the frequency axis with time, similarly the lower sideband which is 180° out of phase with the upper sideband, will rotate in the opposite direction. The result shown in (b) is the sum of the three fields (blue) which changes phase but not amplitude with time. The projection of the modulated field is shown in (c) which shows one quarter cycle of the unmodulated and modulated field with a modulation frequency tens times greater than the carrier frequency.

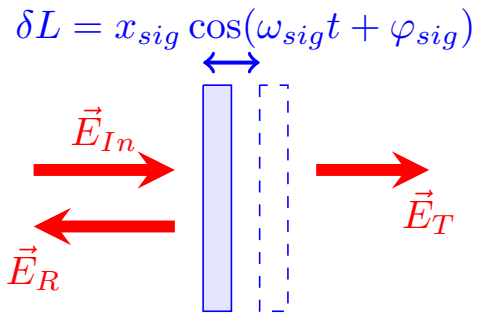


Figure 1.12: Basic interaction of light field with a partially reflective mirror in motion. Mirror motion is defined as a sinusoidal signal dithering the position along the beam propagation axis.

1.4.5 Interaction of Light with a Moving Mirror

It has been shown in Section 1.4.1 how the interaction of a field with a mirror can be described and how the path length of propagation effects the phase of the field. It was also shown how phase modulation can be described (in Section 1.4.4). By combining these elements, it can also be shown how the interaction of a field with a moving mirror introduces phase modulation to the incident field.

Considering a mirror as shown in Figure 1.12 the reflected and transmitted beams were given as $\vec{E}_R = \rho\vec{E}_{In}$ and $\vec{E}_T = i\tau\vec{E}_{In}$. Moving the mirror in the axis of propagation changes the optical path length for the reflected field, so a time varying phase difference is introduced.

The position of the mirror in motion can be given taken as $x_m = x_{sig}\cos(\omega_{sig}t + \varphi_{sig})$ where x_{sig}, ω_{sig} and φ_{sig} are the amplitude, frequency and phase offset of the signal driving the mirror position. The instantaneous phase of the reflected field will be

$$\phi_m = \frac{\omega_0 x_{sig}}{c}\cos(\omega_{sig}t + \varphi_{sig}). \tag{1.35}$$

Then the reflected field will be $\vec{E}_R = \rho\vec{E}_{In}\exp(-2i\phi_m)$ which can be expanded to

$$\vec{E}_R = \rho\vec{E}_{In}\exp(i2kx_{sig}\cos(\omega_{sig}t + \varphi_{sig})). \tag{1.36}$$

Using the Taylor expansion of e^x to the first-order and defining the modulation index as $M = 2kx_{sig}$ this can be rewritten as

$$\vec{E}_R = \rho\vec{E}_{In}\left(1 + \frac{m}{2}\exp\left(-i\left(\omega_{sig}t + \varphi_{sig} - \frac{\pi}{2}\right)\right) + \frac{m}{2}\exp\left(i\left(\omega_{sig}t + \varphi_{sig} + \frac{\pi}{2}\right)\right)\right), \tag{1.37}$$

which has the same form as used previously in Equation 1.33 to describe a phase modulated optical field. Thus, the movement of the mirror at frequency ω_{sig}

generates phase sidebands on the optical field.

In the case of a gravitational wave detector, the movement of the test mass mirrors of an interferometer due to the presence of a gravitational wave will introduce phase modulation sidebands to the light in the interferometer. These phase modulated signal sidebands can be measured as a direct readout of the gravitational wave signal [46].

1.5 Interferometry Basics

The basis of the current generation of gravitational wave detectors [50,51] is a Michelson interferometer enhanced with resonant arm cavities and coupled recycling cavities described below.

1.5.1 Michelson Interferometer

A Michelson interferometer is an optical measurement device sensitive to small differential changes in the two arm lengths of the detector. An example of a laser Michelson interferometer is shown in Figure 1.13 formed from a beam-splitter and two test mass mirrors.

The beam-splitter has four ports, commonly discussed in term of cardinal directions such that the North and East ports contain the test mass mirrors. The West port, where the probe beam is injected, is called the symmetric port as light reflected back along the input path is encoded with information about the symmetric motion (or common mode) of the two test masses. Conversely, light coupled to the Southern port of the beam-splitter is dependent on the differential (anti-symmetric) motion of the test mass mirrors. For gravitational wave detection it is the Southern port of the Michelson interferometer that is used for readout of interferometric gravitational wave signals [46].

The light traveling along each arm of length L_i will acquire a phase delay of $\Phi_i = 2kL_i$ when is returns to the beam-splitter, where $k = 2\pi/\lambda$ and λ is the wavelength of the laser light. Using the treatment laid out in Section 1.4 and assuming the beam-splitter to be an ideal 50:50 splitter such that $R = 0.5$, then the field coupled to the anti-symmetric port can be written as

$$E_{\mathrm{AS}} = \frac{iE_{\mathrm{in}}}{2}(\exp 2ikL_N + \exp 2ikL_E) \tag{1.38}$$

where subscripts N and E refer to the north and east arms of the interferometer respectively.

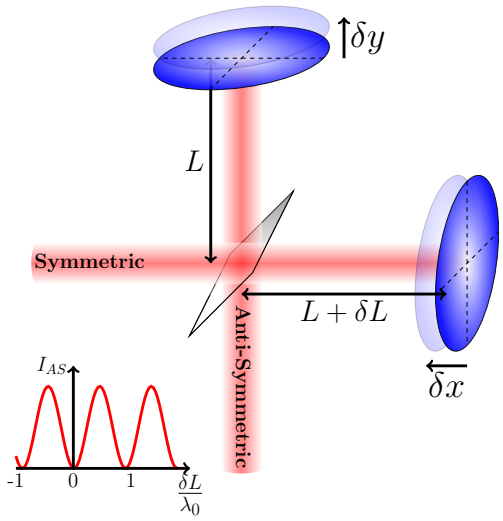


Figure 1.13: A Michelson interferometer consisting of a beam-splitter (grey) and two mirrors (blue). The western port is the symmetric or input port, the northern and eastern (N, E) ports contain the test masses at distances $L_{E,N}$ from the beam-splitter, where $L_E - L_N = \delta L$. The displacement of the test masses from equilibrium is marked as δx and δy. The intensity of light coupled to the southern, anti-symmetric port is plotted in the insert as a function of the change in length.

If the symmetric or common mode arm length is defined as $L_{sym} = \frac{L_N + L_E}{2}$ and the differential arm motion as $L_{asym} = L_N - L_E$, the the above expression can be simplified to give;

$$E_{\mathrm{AS}} = iE_{\mathrm{in}} \exp 2ikL_{sym} \cos kL_{asym}. \tag{1.39}$$

If this field is detected on a photodiode the output signal (I_{AS}) will be proportional to the product of this field and the complex conjugate of this field. Meaning the dependence on the common arm motion is removed from the detected signal which will vary with the cosine squared of the differential arm length factor, also written as

$$I_{\mathrm{AS}} \propto E_{AS} E_{AS}^{*} = P_{AS} \cos^2 kL_{asym}, \tag{1.40}$$

where P_{AS} is the power of the optical field E_{AS}.

Equation 1.40 and the insert in Figure 1.13 show that if the differential length of the interferometers is equal to a half integer number of half wavelengths, then no optical power will be coupled to the anti-symmetric port of the interferometer. This is referred to as the dark fringe condition. For the best signal to noise performance of a gravitational wave detector interferometer with power recycling (see Section 1.5.3), the arm cavities are held to an operating point close to, or on, the dark fringe.

If a gravitational wave interacts with a Michelson interferometer, the fluctua-

tions in the space-time strain will modulate the apparent lengths of the two arm cavities differentially. This will result in a modulation signal being detected at the signal readout port of the interferometer that is directly related to the gravitational wave signal. However, the coupling of gravitational wave signals to the interferometric readout is very weak. The strain of an observable gravitational waves is of the order 10^{-21} [1], therefore the strain sensitivity ($\frac{\delta L}{L}$) of a detector is required to be lower to be able to detect a gravitational wave signal with enough significance. In order to achieve this extreme level of sensitivity additional optical enhancements to the basic Michelson interferometer are included in the detector system.

1.5.2 Fabry-Pérot Cavity

Adding a partially transmissive input test mass (ITM) between the beam-splitter and the end test masses (ETM) of a Michelson interferometer forms a Fabry-Pérot cavity in each arm of the interferometer [46]. A Fabry-Pérot cavity consists of two parallel highly reflective mirrors that form a linear optical resonator. If the length of the cavity is equal to an integer multiple of half the carrier wavelength then constructive interference occurs and the cavity is said to be on resonance with the input field.

On resonance the probe light bounces between the two mirrors multiple times, building up the amount of optical power in the interferometer arm and increasing the number of times the probe light travels along the distance between the test masses. This increases the coupling strength of a gravitational wave signal to the light field and increases the sensitivity of the detector.

The amount of optical power circulating inside the cavity on resonance, and the amplitudes of the fields reflected and transmitted from the cavity, depend on the reflectivity and transmissivity of the two mirrors. Using the convention laid out in Section 1.4.1 for the reflectivity and transmissivity power and amplitude coefficients, the amount of power circulating in a Fabry-Pérot cavity can therefore be written as

$$P_{circ} = \frac{T_{\mathrm{ITM}}}{1 + R_{\mathrm{ITM}}R_{\mathrm{ETM}} - 2\sqrt{R_{\mathrm{ITM}}R_{\mathrm{ETM}}}\cos(2kL)}, \tag{1.41}$$

with $k = 2\pi/\lambda$. The length of the arms is effectively increased as the light travels between the mirrors multiple times. The number of round trips is a factor known

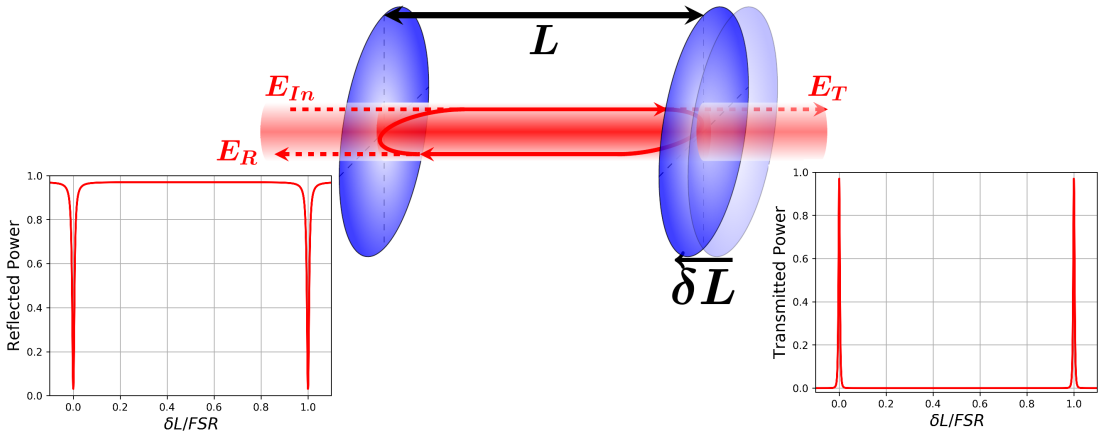


Figure 1.14: FA Fabry-Pérot Cavity formed of two mirrors of equal reflectivity and no absorption or scatter loss, for which the insert plots show the relationship between cavity detuning δL and the normalised power reflected and transmitted from the cavity. The reflected and transmitted power is plotted as a function of $\frac{\delta L}{\mathrm{FSR}}$ where FSR (free-spectral range) is the frequency spacing between resonant conditions.

as the finesse of the cavity and is given by the expression

$$\mathcal{F} = \frac{\pi}{2\arcsin\left(\frac{1-\sqrt{R_{ITM}R_{ETM}}}{2\sqrt[4]{R_{ITM}R_{ETM}}}\right)}. \tag{1.42}$$

The free-spectral range (FSR) is the spacing in frequency between cavity resonances

$$\mathrm{FSR} = \frac{c}{2L}. \tag{1.43}$$

If the length of the cavity, or the frequency of the seed light is changed the power inside the cavity will decrease and the cavity is said to be detuned from resonance. The linewidth or full width half maximum (FWHM) is defined as the frequency detuning from the resonance condition at which the power drops to half the maximum. The linewidth is the ratio of free spectral range to finesse

$$\mathrm{FWHM} = \frac{\mathrm{FSR}}{\mathcal{F}}. \tag{1.44}$$

The half-width half maximum is also often used when describing the dynamics of a resonant cavity and is often referred to as the cavity pole ($f_c = \mathrm{FWHM}/2$) as above this frequency the response of the field in the cavity to length changes drops of as $1/f$ and can be approximated as a first-order low pass filter.

The fields reflected (E_R) and transmitted (E_T) by the cavity can be calculated

as

$$E_T = \frac{\tau_{ITM}\tau_{ETM}E_I e^{-i\delta}}{1-\rho_{ITM}\rho_{ETM}e^{-i2\delta}} \tag{1.45}$$

$$E_R = \rho_{ITM}E_I - \frac{\rho_{ETM}\tau_{ITM}^2 E_I e^{-i2\delta}}{1-\rho_{ITM}\rho_{ETM}e^{-i2\delta}} \tag{1.46}$$

and the dependence on the cavity detuning δ is plotted in the insert of Figure 1.14.

Three distinct cases exist where the mirror reflectivities of the two mirrors affect the behaviour of the cavity. If the ITM reflectivity is greater than the ETM reflectivity ($R_{ITM} > R_{ETM}$) then the cavity is said to be under-coupled, and likewise the cavity is over-coupled if $R_{ITM} < R_{ETM}$. If the reflectivities of the two mirrors are equal ($R_{ITM} = R_{ETM}$) the cavity is said to be impedance matched, and as such the reflected power on resonance drops to zero and a maximum amount of power is transmitted. In the over and under-coupled cases, the amount of power reflected and transmitted depends on the values of reflectivity and transmissivity and the circulating power is greater for the over-coupled case.

For use as arm cavity enhancements to a Michelson interferometer a high level of circulating power is required so the over coupled case is utilised. In some implementations, such as mode cleaner cavities, a high throughput power is desirable so an impedance matched, or close to impedance matched, cavity may be used.

The sensitivity of a Michelson interferometer is improved by the addition of Fabry-Pérot cavities as the interaction of the probe light with the arm cavities is increased. By increasing the finesse of the cavity, the sensitivity is further improved but at the expense of bandwidth as the interferometer has lower sensitivity to signals at frequencies greater than the cavity pole. An example cavity is shown in Figure 1.14 with the reflected and transmitted intensity plotted in the inserts.

1.5.3 Power Recycling

Power recycling is a technique that increases the amount of optical power circulating inside the interferometer without limiting the bandwidth of the detector [50].

For an ideal Michelson interferometer operating at the dark fringe condition all the incident optical power will be coupled to the symmetric port of the interferometer beam-splitter and no light coupled to the readout port.

When viewed from the symmetric port the interferometer acts as a mirror reflecting light back along the input path. Placing an additional power recycling

mirror (PRM) in the symmetric port of the interferometer forms a cavity with the interferometer, reflecting the light back into the instrument to interact again with the arm cavities and increasing the circulating power [23].

1.5.4 Signal Recycling

For an interferometer operating in the dark fringe configuration, the gravitational wave signal sidebands are coupled to the anti-symmetric readout port of the interferometer. In a similar manner to power recycling, a mirror can be placed between the beam-splitter and the interferometer readout forming a coupled resonant cavity.

Different resonant conditions of this signal recycling cavity (SRC), defined by the cavity length and reflectivity of the signal recycling mirror (SRM), can improve the sensitivity of the detector across different frequency bands [52, 53].

In a resonant sideband extraction scheme (RSE) the signal recycling mirror (SRM) and the main interferometer form a cavity in which the gravitational wave signal of a selected frequency is resonant and therefore the signal will be amplified.

Gravitational wave detectors that have implemented signal recycling (GEO and LIGO) to date have done so in a broadband detuned configuration where the sensitivity between 10 Hz and 5 kHz is improved.

1.5.5 Dual Recycling Fabry-Pérot Michelson Interferometer

The combination of all these enhancements is known as a dual recycling Fabry-Pérot Michelson interferometer (DRFPMI), an example is shown in Figure 1.15 and the effect of optical recycling on the sensitivity of a 4 km scale Fabry-Pérot Michelson gravitational wave detector is shown in Figure 1.16.

1.5.6 Optical Readout and Signal Detection

To detect length changes of a Michelson interferometer due to gravitational wave signals, a photodetector is placed in the asymmetric port as a signal readout device. A photodetector measures the intensity or the power of the optical field impinging on the detector. It is not possible to directly measure the amplitude of an optical field.

The field at the readout port of a Michelson interferometer will consist of multiple component fields at different frequencies including the carrier frequency of the probe light ω_o as well as any modulation or signal sidebands. The power

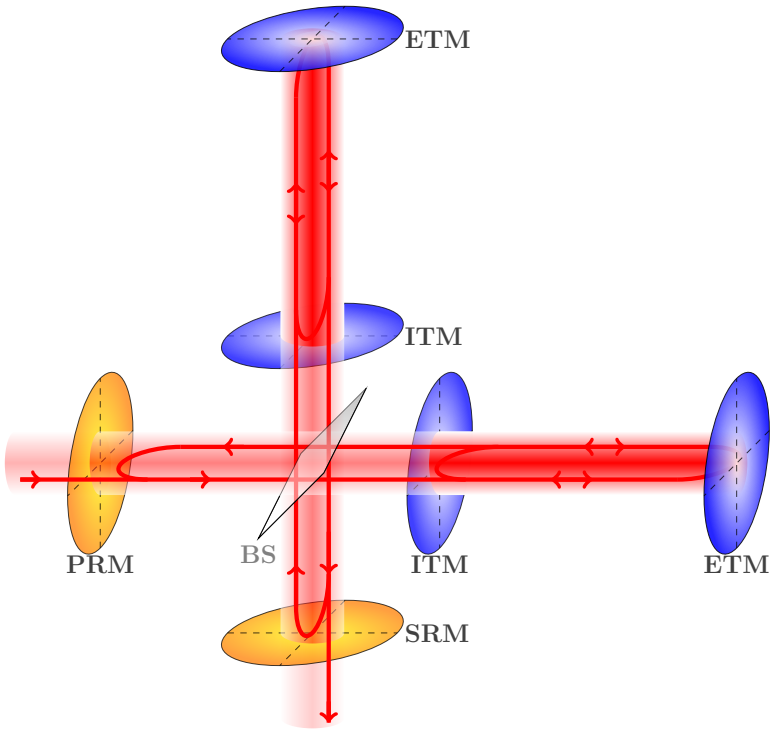


Figure 1.15: Dual-recycling Fabry-Pérot Michelson interferometer for gravitational wave detection. The Fabry-Pérot cavities are formed by the blue mirrors. The central beam-splitter is shown in grey, and the recycling mirrors in orange.

of the optical field can be calculated as the product of the field amplitude and its complex conjugate $P = EE^*$.

In the dark fringe condition, only the signal sidebands are detected at the output port but if the interferometer is slightly detunded from the dark fringe condition a small amount of the carrier light is also detected and the beat between the signal and carrier can be measured in a homodyne detection scheme known as *DC-readout*.

Consider a simple case of a field consisting of a carrier at frequency ω_0 and a pair of signal sidebands at frequencies $\omega_0 \pm \omega_{sig}$ modulated onto the carrier by the interaction of the field with the moving test masses of a Michelson interferometer. The field measured by the output photodetector will then be the sum of the beats at multiple frequencies between carrier and the two sidebands.

At DC there will be the beat between the carrier with itself, and the upper and lower sidebands with themselves. At a frequency of ω_{sig} there will be the signal from the beating of the sidebands with the carrier and at $2\omega_{sig}$ the beat between upper and lower sidebands. From this interferometric readout the gravitational wave signal can be reconstructed; this is the basic description of how interferometry is used for gravitational wave detection.

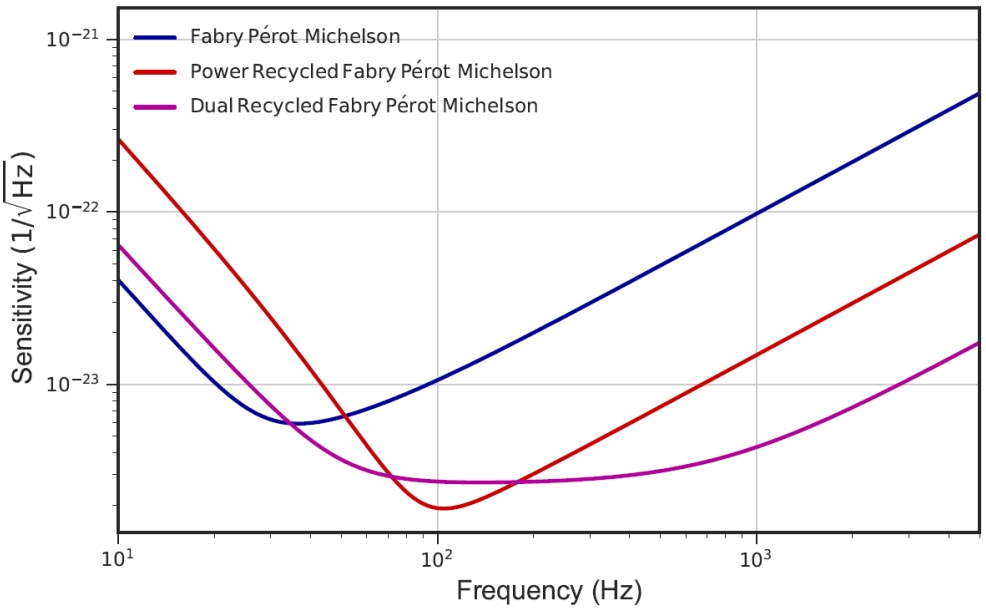


Figure 1.16: Simulated sensitivity of a 4 km scale Michelson interferometer with different interferometric enhancements. All three curves were produced with the FINESSE [54] modelling tool with Advanced LIGO design parameters. The low frequency shape of the curves is due to radiation pressure noise described in Section 2.2. The blue curve is a Michelson interferometer with Fabry-Pérot arm cavities while the red curve shows the improvement gained by adding a power recycling mirror, and the purple shows the result of the implementation of dual recycling for one choice of parameters.

Chapter 2

Quantum Noise

This chapter presents an overview of the theory of quantum noise in interferometers and focuses on one method of reducing quantum noise in a gravitational wave detector: the *speedmeter* interferometer.

In Figure 1.3 of Section 1.3 the modelled quantum noise limit of the Advanced LIGO detector is shown in the pink trace. In this plot it is shown that, when operated at design sensitivity, the LIGO detectors will be limited across much of the gravitational wave detection band by quantum noise.

In Part II of this thesis an experimental effort to demonstrate the feasibility of the speedmeter interferometer is presented.

2.1 Quantum Noise of Light

To understand quantum noise, we need a quantum picture of light and a modified description of the interaction of light with an interferometer.

The classical view was presented in Chapter 1 that treated light as an optical field propagating through the interferometer. In the quantum picture light is represented by creation ($\hat{a}^\dagger$) and annihilation ($\hat{a}$) operators in what is known as the two-photon formalism [55, 56].

In a gravitational wave interferometer, the input light field (or carrier) of frequency ω_0 probes the arm cavities. Gravitational waves modulate this carrier light and generate signal sidebands at frequencies $\omega_0 \pm \omega_{sig}$ where ω_{sig} is the frequency of the gravitational wave signal. These upper and lower sidebands are represented in the formalism by the operators

$$\hat{a}_+ = \hat{a}_{\omega_0+\omega_{sig}} \tag{2.1}$$

$$\hat{a}_- = \hat{a}_{\omega_0-\omega_{sig}}. \tag{2.2}$$

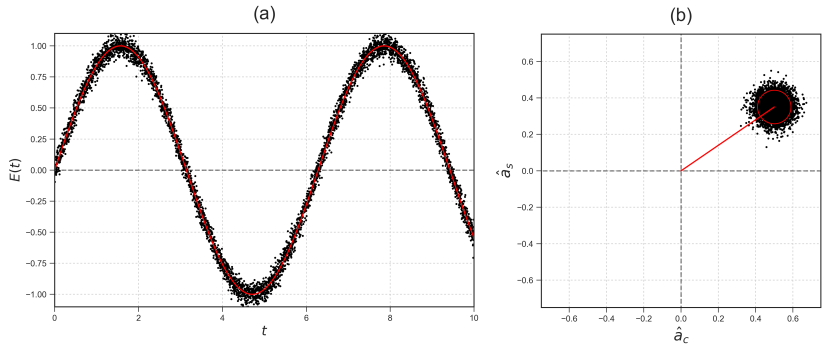


Figure 2.1: Representation of a quantised carrier state as a time series (a) and in quadrature form (b).

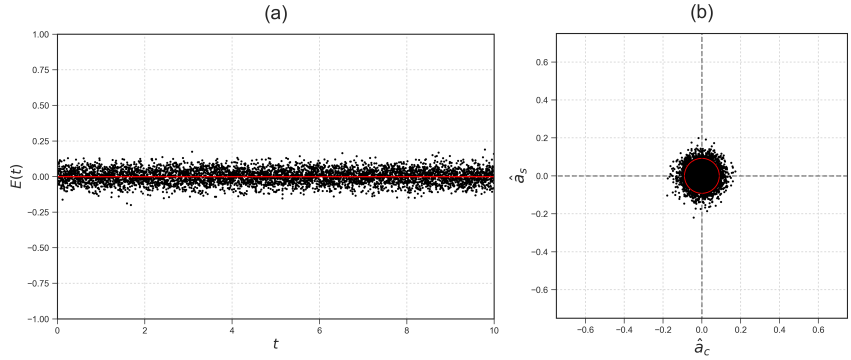


Figure 2.2: Representation of a quantised vacuum state as a time series (a) and in quadrature form (b).

The operators for the signal sidebands can be rewritten in terms of quadrature operators $\hat{a}_c$ and $\hat{a}_s$ such that in the frequency domain

$$\hat{a}_c(\omega) = \frac{\hat{a}_+ + \hat{a}_-}{\sqrt{2}} \tag{2.3}$$

$$\hat{a}_s(\omega) = \frac{\hat{a}_+ - \hat{a}_-}{i\sqrt{2}}, \tag{2.4}$$

and the time domain operators can be calculated using the Fourier transform. An arbitrary electric field vector can then be represented in the quantum picture by the operator

$$\hat{E}(t) = \sqrt{\frac{4\pi\hbar\omega_0}{\mathcal{A}c}}(\hat{a}_c(t)\cos\omega_0 t + \hat{a}_s(t)\sin\omega_0 t), \tag{2.5}$$

where $\mathcal{A}$ is the effective cross-section area of the light and spatial and polarisation vectors are neglected [57].

A consequence of this representation is that there is an inherent variance or uncertainty related to each operator due to Poisson statistics of the quantised state. A coherent state operator will have a mean value that follows a sinusoidal

oscillation at the carrier frequency ω_0 but there is a variance of the operator in both quadratures that constitutes noise on the amplitude and frequency of the carrier. This is shown graphically in Figure 2.1(a) in a time series and in Figure 2.1(b) in phase space. In the phase space picture, the state appears as a 'ball on a stick.' The stick is the vector representing the mean values of the two quadrature amplitudes and the ball represents the variance or uncertainty in the value of the two quadrature amplitudes [58]. These variations translate to noise on the carrier field.

For a vacuum state, the mean value of the operator in both quadratures is zero but the variance is non-zero. Figure 2.2 represents this graphically as a time series (a) and in phase space (b) as a 'ball' centred on zero with no 'stick'. In the time series this results in noise centred around zero with some non-zero RMS value.

2.2 Quantum Noise in Interferometers

In an interferometer, quantum noise affects the sensitivity of the instrument either as fluctuations in the amplitude or frequency of the carrier field or as vacuum fluctuations coupling to the readout.

The vacuum state fluctuations are background noise and as such can couple into an interferometer via any open port of the instrument. This including the readout port, the partially transmissive end test masses and sources of optical loss within the interferometer. Vacuum fluctuations couple to the signal readout as shot noise and also interact with the suspended test masses contributing to back-action noise [59].

The presence of quantum noise in an interferometer affects the sensitivity via two mechanisms. Phase fluctuations of the beam at the readout port result in a level of shot noise on the detected signal as there is a statistical variation in the number of photons detected at any one time. This is known as *quantum shot noise*. At the test masses the variation in the number of photons incident on the suspended mass results in a variation in the momentum transferred from the light to the optic in the opto-mechanical radiation pressure force. This results in *quantum radiation pressure noise* disturbing the test mass position.

The two forms of quantum noise can be expressed in terms of a limit to the strain sensitivity [58]:

$$h_{shot} = \frac{1}{L}\sqrt{\frac{\hbar c\lambda}{\pi^2 P}} \tag{2.6}$$

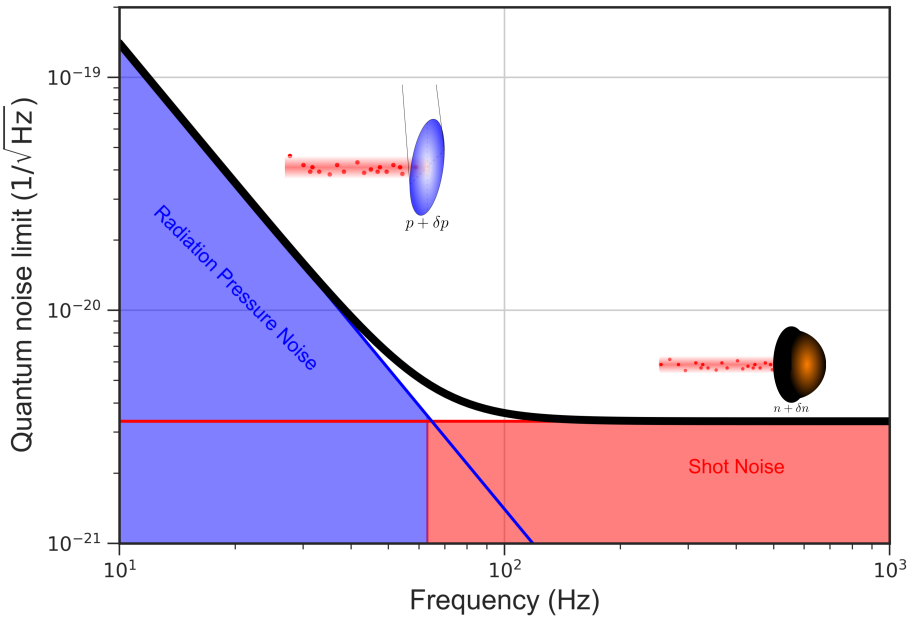


Figure 2.3: Quantum noise in an interferometer. Inserts show conceptual representation of quantum noise. The fluctuation δn in the number of photons $n + \delta n$ measured by a photodetector resulting in shot noise (red). The random momentum kicks $p+\delta p$ imparted to a suspended mirror by the beam resulting in radiation pressure noise (blue).

and

$$h_{RP} = \frac{1}{mf^2L}\sqrt{\frac{\hbar P}{2\pi^3 c\lambda}}. \tag{2.7}$$

Quantum noise can be reduced by changing the design parameters of the interferometer.

Radiation pressure noise can be decreased by increasing the mirror mass (m) at no cost to shot noise, however there is a technological limit to size (and therefore mass) of optics that can be manufactured.

By increasing the effective power (P) of the laser inside the arm cavities of the interferometer the shot noise is reduced but at the expense of greater radiation pressure noise. Also, there is a technical limitation to the amount of coherent laser power that can be generated and injected into an interferometer [60].

Increasing the length (L) of the interferometer would improve both shot noise and radiation pressure noise but once more this is limited by practical concerns as the length of the detector is constrained by the feasibility and cost of the detector infrastructure. Increasing the cavity length ultimately limits the bandwidth of the detector.

The plot in Figure 2.3 shows the quantum noise of a typical gravitational wave detector as an example. The quantum shot noise is independent of frequency, whereas the radiation pressure noise falls of as $1/f^2$. In the plot of the fundamental limits to the design sensitivity of Advanced LIGO in Figure 1.3 the quantum noise increases with frequency above 100 Hz this is due to the response of the detector with dual recycling decreasing as these frequencies are above the arm cavity corner frequency, this is unrelated to fundamental quantum noise.

2.3 Standard Quantum Limit

It was shown graphically in Figures 2.1 and 2.2 that the magnitude of the variations (area of the ball) in orthogonal quadratures are equal for both carrier and vacuum fields. These variations are described by non-correlated spectral densities of the measurements of position and radiation pressure force associated with the free masses; such quantum measurements are bounded by the Heisenberg uncertainty principle. This means that the product of the uncertainties of the two measurements can be no less than $\hbar/2$ or that the product of the amplitude spectral densities of the two measurements is limited by the uncertainty relation

$$S_x(\omega)S_p(\omega) \geqslant \frac{\hbar}{2}. \tag{2.8}$$

where $S_x(\omega)$ and $S_p(\omega)$ are the amplitude spectral densities relating to measurements of the position and momentum operators respectively.

The quantum noise limit to a laser interferometer as a quantum measurement device is the sum of the two quantum noise spectral densities associated with the shot noise and the back-action noise of the instrument. The uncertainty relation enforces a lowest possible noise level that this could be at any given frequency, this is known as the standard quantum limit (SQL) [51].

2.3.1 Schemes for Reducing Quantum Noise

The shot noise component of quantum noise is a consequence of measuring the optical field with a photodetector, a necessary part of the process of gravitational wave detection. The radiation pressure noise is a consequence of the interaction of the probe light with the test masses.

The standard quantum limit represents a limitation to the precision of measurement meters due to shot noise and radiation pressure noise but is not a fundamental limitation of all measurement devices. The SQL can also be thought of as the boundary between the classical and quantum domains. Interferometers can be constructed in such a way to have sensitivity beyond the SQL by embracing the quantum nature of light and by understanding quantum noise in an interferometer.

Additions to the interferometer such as squeezed light injection can reduce the quantum noise by manipulation of the light within the interferometer. Alternatively, quantum noise can be reduced by changing the measurement parameter of the detector.

Squeezed Light Injection

The quantum noise of light can be manipulated by a process known as squeezing [61] where the correlations between phase and amplitude noise are used to reduce noise in one quadrature at the expense of increasing the noise in the orthogonal quadrature. Considering the 'ball and stick' representation, the area of the ball cannot be decreased beyond the standard quantum limit, but the shape of the ball can be altered (or squeezed) to reduce the noise in one quadrature.

Vacuum states can couple into the interferometer from the readout port. These vacuum fluctuations interact with the test masses and are subsequently detected by the readout photodetector leading to quantum noise contaminating the signal readout. Instead of allowing vacuum fluctuations to couple into the detector an alternative squeezed field can be injected [62].

A choice must be made to squeeze either the phase or amplitude noise of the injected vacuum field. Currently, gravitational wave detectors are limited at higher frequencies by quantum shot noise but at lower frequencies the sensitivity is limited by technical noise and thermal suspension noise not by quantum radiation pressure noise. The shot noise can be reduced by squeezing the phase noise by an amount that keeps the related rise in radiation pressure noise below other noise sources.

This implementation of squeezed light to reduce shot noise has been successfully used as an enhancement of gravitational wave detector interferometers. First demonstrated in the GEO gravitational wave detector [24, 25] this quantum noise enhancement technology has been adopted by the LIGO [63] and Virgo projects [64].

This form of squeezed light injection is known as frequency independent squeezing as the squeezing angle is the same at all frequencies. The squeezing angle is the direction in the quadrature axis along which the noise ellipse is minimised or squeezed.

As phase and amplitude noise affect the quantum noise limit differently at high and low frequencies if the phase quadrature of the injected vacuum field could be squeezed at high frequencies and the amplitude quadrature at lower frequencies then the quantum noise would be reduced across the whole detection band. This is known as frequency dependent squeezing and requires the addition of a filter cavity to rotate the angle of the squeezed ellipse.

Squeezed light can be used to enhance the sensitivity of any optical interferometer either Michelson or Speedmeter.

Quantum Non-Demolition Measurements

Preforming a measurement of the position of a test-mass limits the sensitivity of further measurements at a later time because the position of the test-mass is perturbed by the opto-mechanical interaction between the optical probe field and the mechanical test mass. If instead the variable measured has a time evolution that is not perturbed, then there is no precision limit imposed. Such a measurement is known as a quantum non-demolition (QND) measurement [65].

It is possible to conceive of methods of cancelling the effect of the radiation pressure in the measurement system by constructing a QND measurement device. One such way is to change the topology of the interferometer to what is known as a speedmeter interferometer. This is discussed in detail in Section 2.4 and is the quantum noise reduction scheme examined in this thesis.

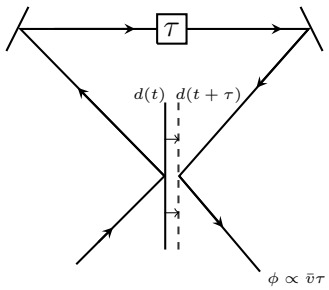


Figure 2.4: A model speedmeter, a representation of the simplest possible speedmeter. A probe interacts twice with a moving test mass. The interactions are temporally separated by a delay of τ, the result is that the probe has information about the velocity of the test mass.

2.4 Speedmeter Interferometry

In an interferometer scheme the time evolution of the momentum of the test mass is not perturbed by the interaction with the probe field over time-scales longer than the interaction. This makes the momentum of the test mass a QND variable. As a QND variable, the momentum operator of the free test-mass at a given time $p(t)$ commutes with the momentum operator at a later time $p(t+\tau)$,

$$[p(t), p(t+\tau)] = 0. \tag{2.9}$$

The momentum of a test-mass is not a directly measurable quantity of an optical instrument. Instead, velocity is used as a proxy for momentum, assuming a constant mass. In this case, the velocity of the test mass is perturbed by the probe beam during the measurement but returns to the unperturbed time evolution. So over timescales of the measurement an optical speedmeter can be considered as a quantum non-demolition measurement device [57, 66].

A model interferometric speedmeter scheme is shown in Figure 2.4. Here a laser beam is used to probe a mirror. The first reflection from the front mirror surface induces a phase shift on the light of $\delta\phi_1 = 4\pi\frac{\delta x(t)}{\lambda}$, where δx is the time dependent displacement of the mirror from equilibrium, and λ is the laser wavelength. The beam then travels along a path with time delay of τ, and is reflected again from the reverse face of the mirror, acquiring a second displacement dependent phase shift of $\delta\phi_2 = 4\pi\frac{\delta x(t+\tau)}{\lambda}$. The combined phase shift $\delta\phi$ the output probe beam,

$$\delta\phi = \frac{4\pi}{\lambda}(\delta x(t) - \delta x(t+\tau)) = \frac{4\pi}{\lambda}v\tau, \tag{2.10}$$

can then be measured as a signal directly related to the velocity (v) of the test

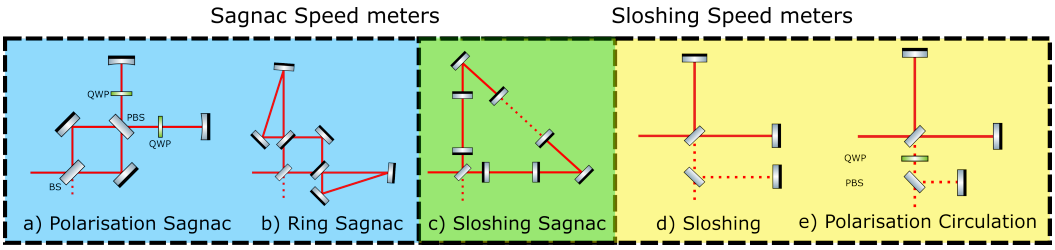


Figure 2.5: The speedmeter zoo. Simplified representation of the various speedmeter designs which are discussed in Section 2.4.1

mass mirror.

2.4.1 Speedmeter Topologies

The idea of a quantum speedmeter as a QND gravitational wave detector was first proposed by Braginsky and Khalili in [67]. The idea was then expanded to work as a laser interferometer [68, 69]. Since then, many designs of speedmeter interferometers for gravitational wave detection have been proposed.

The different speedmeter designs [57, 70] can be broadly classified as either sloshing speedmeters or Sagnac speedmeters, this distinction is used to categorise the different schemes in the simple representation of the different topologies in Figure 2.5.

Ring Sagnac

In a Sagnac interferometer two light fields travelling in counter propagating directions probe both sets of test-masses. It was shown in [71] that a Sagnac interferometer is inherently a speedmeter as the phase of the output is proportional to the difference in test mass velocity.

Sagnac interferometers are sensitive to rotations, and as such are used as gyroscopes. The rotational signal is proportional to the area encompassed by the interferometer. By folding the arms of the interferometer, the optical layout can be such that the area encompassed by the interferometer is effectively zero.

The original proposal of a Sagnac speedmeter as a QND interferometer [71] utilised ring cavities in the arms to achieve the Sagnac effect in a long base line interferometer, as in Figure 2.5(b). Light is split into two equal, counter-propagating field, the clockwise field (C) probes the north arm first at time t and then the east arm at time $t+\tau$. The anticlockwise (AC) field probes the east arm at time t and the north arm at time $t+\tau$, resulting in the phase of the light at the read-out

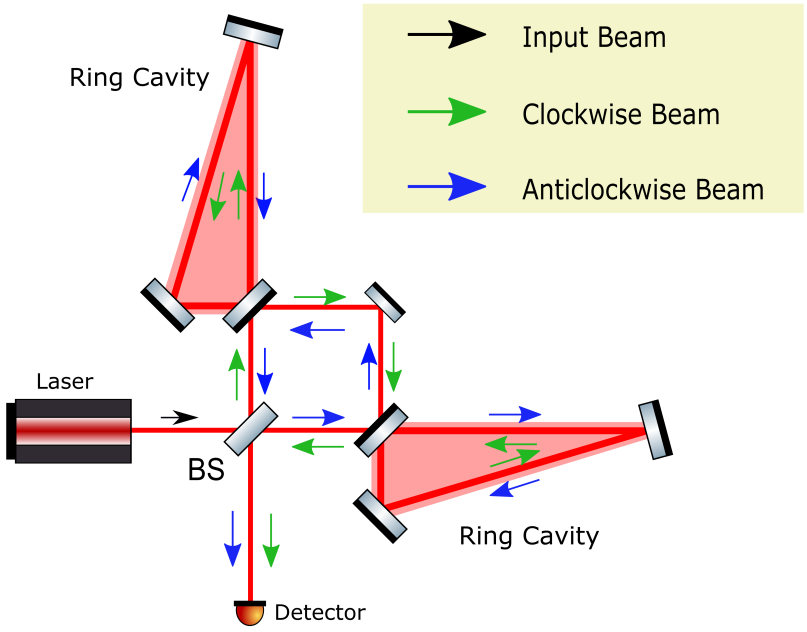


Figure 2.6: The Ring Sagnac Speedmeter. Schematic of the optical Sagnac speedmeter interferometer with ring arm cavities. In this scheme the input probe beam impinges on a beam-splitter where it is split into the clockwise (green arrows) and anti-clockwise (blue arrows) fields that sense the two arm cavities in reverse order. In this topology triangular ring cavities using three test mass mirrors are implemented.

port given as

$$\begin{aligned}\delta\phi &= \delta\phi_C - \delta\phi_{AC} \\ &\propto \Big(x_2(t) - x_2(t+\tau)\Big) - \Big(x_1(t) - x_1(t+\tau)\Big) \\ &\propto (v_2 - v_1)\end{aligned} \tag{2.11}$$

as required.

The largest concern with the Ring Sagnac scheme as a speedmeter gravitational wave detector is the problem of back-scatter in the cavity arms causing light to couple between the counter-propagating field inside the interferometer [72].

Polarisation Sagnac

An alternative to the ring Sagnac is the polarisation Sagnac which utilises polarising optics to separate and direct the counter propagation fields (shown in Figure 2.5(a) and in more detail in Figure 2.7).

The polarisation Sagnac interferometer as a gravitational wave detector was first proposed in [73] and the design was further considered in [74].

The use of polarising optics allows for the separation of the counter-propagating probe fields by using orthogonal polarisation of light. Removing the need for ring cavities, instead linear Fabry-Pérot arm cavities can still be used to enhance the sensitivity of the interferometer but are not part of the speedmeter mechanism.

The polarisation Sagnac speedmeter is the main topology investigated in Part II of this thesis.

Sloshing Speedmeter

A sloshing speedmeter is similar to a Michelson interferometer with the addition of a so called sloshing cavity at the readout port of the interferometer, as shown in Figure 2.5(d) [75]. Light leaving the interferometer 'sloshes' in an external cavity and is coupled back into the interferometer to interact a second time with the test-masses. The momentum kick from the initial interaction of the light with the test-masses is cancelled by a second interaction of the light from the sloshing cavity which has acquired a phase change of π due to the extra propagation distance. The first practical design of a speedmeter for gravitational wave detection was a sloshing speedmeter [69, 75]. More recent proposals have improved on the initial sloshing design with the use of polarising optics [76, 77] shown in Figure 2.5(e).

The phase of the light at the read out of a sloshing speedmeter has a phase

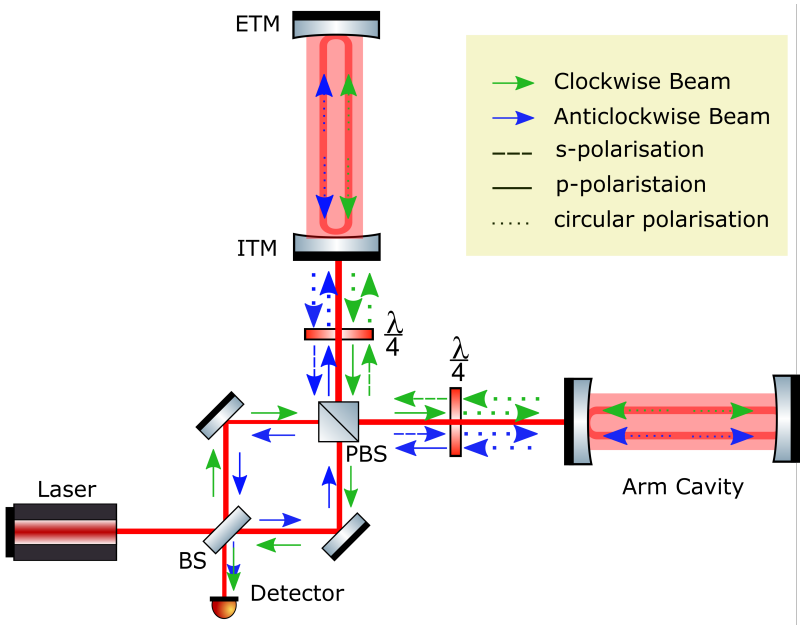


Figure 2.7: The Polarisation Sagnac Speedmeter. The input p-polarised field (solid arrows) is split into two equal counter propagating fields by the beam-splitter (BS). The clockwise beam (blue) senses the North cavity and then the East. The anticlockwise beam (green) senses the cavities in the reverse order. Each field is transmitted by the polarising beam-splitter (PBS) towards the arm cavities. The quarter waveplate between the PBS and the arm cavity ensures that the light is circularly polarised (dotted) inside the cavity, and s-polarised (dashed) when it returns to the PBS. Upon interacting a second time with the PBS each field is polarised orthogonal relative to the first interaction and so is reflected towards the second cavity. After interacting with a second QWP and cavity, the polarisation of each field when it interacts with the PBS a third time is the same as it was on the initial interaction. And so, after probing both cavities the two fields are transmitted by the PBS and interfere at the BS. The phase of the light coupled to the anti-symmetric port is proportional to the differential speed of the test masses.

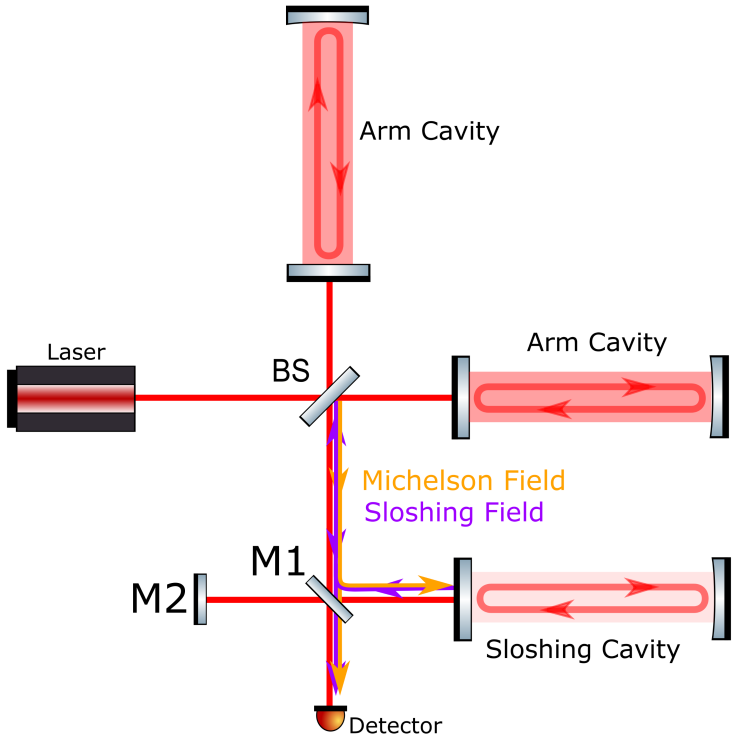


Figure 2.8: The sloshing speedmeter. The core of the instrument remains a Fabry-Pérot Michelson interferometer with an additional sloshing cavity in the anti-symmetric port. Light exiting the interferometer interacts with a partially reflective mirror (M1) that couples some light into the sloshing cavity and then back into the interferometer. There is a 180° phase difference between the Michelson field and the sloshing field. An additional mirror (M2) is included behind M1 to close this port of the interferometer to reduce the coupling of vacuum fluctuations into the interferometer.

shift proportional to the differential velocity of the two test masses,

$$\begin{aligned} \delta\phi &= \delta\phi_{Michelson} - \delta\phi_{sloshing} \\ &\propto x_1(t) - x_2(t) - \Big(x_1(t+\tau) - x_2(t+\tau)\Big) \\ &\propto (v_1 - v_2) \end{aligned} \tag{2.12}$$

as required.

Sloshing Sagnac Speedmeter

A hybrid sloshing Sagnac speedmeter has also been investigated [78]. This configuration includes features of both types of speedmeter. Light circulates around the interferometer in opposite directions, sensing both cavities. But between the two arms, however, the field sloshes in a third anti-resonant cavity. To be anti-

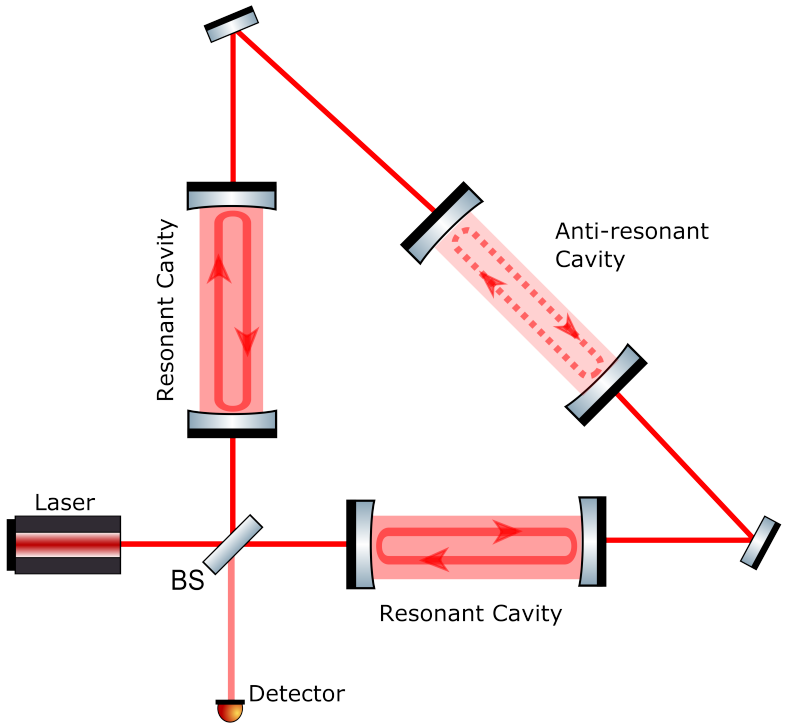


Figure 2.9: The sloshing Sagnac speedmeter. In this scheme two of the three cavities are held on resonance with the laser light, while the third is held off resonance. Light transmitted by each of the resonant arm cavities couples to the orthogonal cavity via the sloshing cavity.

resonant the length of the sloshing cavity has to be equivalent to $(N + \frac{1}{4})\lambda$.

One possible layout is shown in Figure 2.9. In this triangular configuration, the throughput of each Fabry-Pérot arm cavity is coupled into a third cavity that is anti-resonant, i.e. the length of the cavity is set so that the light interferes destructively after each round trip. The light that returns from the sloshing cavity has acquired a phase of π so when it interacts with the main arm cavities a second time it cancels the radiation pressure noise of the initial interaction.

This scheme has the added advantage of reducing the coating thermal noise in the interferometer as the effect is similar to the implementation of so-called *Khalili cavities* [79].

A Khalili cavity is an anti-resonant Fabry-Pérot cavity that replaces the ETM in an Fabry-Pérot Michelson interferometer [80]. In such a scheme the reflectivity of this composite ETM is determined by the combination of the reflectivity of the two mirrors forming the anti-resonant cavity. Coating thermal noise of an interferometer mirror scales with the number of layers in the coating as does reflectivity. To produce the higher reflective ETM's required for a gravitational wave detector interferometer requires more than 20 coating layers. However, replacing each ETM with a two mirror anti-resonant cavity relaxes the reflectivity requirement on each individual mirror meaning an overall reduction in coating thermal noise.

Polarisation Circulation Speedmeter

Two separate designs for a sloshing-type speedmeter that use polarising optics in the readout of the detector have also been proposed. These are known as the *polarisation circulation* speedmeter and the *polarisation folding* speedmeter schemes.

In both cases the core interferometer is the same as a Michelson position-meter with additional polarising optics and sloshing mirror in the readout port. Similar to the sloshing speedmeter the probe light is forced to interact twice with the test masses with a time delay between interactions.

In the polarisation folding speedmeter scheme [76] two quarter waveplates are placed either side of a signal recycling mirror in the output port of a Michelson interferometer held on the dark fringe. On the readout side of the signal recycling mirror is a polarising beam-splitter. In one port is a folding mirror and in another is the readout photodetector. The length between the signal recycling mirror and the folding mirror sets the phase difference between the two circulating field and the rotation of the quarter waveplates sets the relative polarisation of the two fields.

The polarisation circulation scheme [77] uses a similar method with a simpler optical arrangement. Here a single quarter waveplate is placed before the PBS in the dark port again with folding mirror in one port and readout in the other. Here the output field from the Michelson interferometer is rotated by the QWP from linear polarisation to circular polarisation so that half the light is reflected to the folding mirror while half is transmitted directly to the readout detector. The folding mirror reflects the light back towards the interferometer where the second pass of the QWP rotates the light back into the orthogonal linear polarisation.

In comparison to the sloshing speedmeter design, there is no requirement for an additional cavity thus reducing the technical complexity of interferometer and the required locking scheme.

2.5 Practical Limitations to Speedmeter Designs

In the ideal case, a speedmeter interferometer is sensitive only to differences in relative velocity of the arm cavity test-masses, a proxy for a quantum non-demolition (QND) variable. Measurement of a non-QND variable will degrade the benefit to quantum noise limited sensitivity of the speedmeter.

The moving test masses imparts information on the relative speed of the test masses onto the phase of the probe field as signal sidebands. In the ideal speedmeter case, information on the relative position of the test masses is cancelled by the speedmeter mechanism. In the non-ideal case, there is an incomplete cancellation of positional information leading to both speedmeter and position meter signals being measured.

The non-ideal speedmeter has to be considered when discussing the possibility of a gravitational wave detector based on the speedmeter concept. Technical limitations in addition to fundamental noise limits will always limit the sensitivity of the instrument. For the purposes of this study the focus is on technical limitations that are specific to the speedmeter interferometer, mainly effects that cause a breaking of the speedmeter mechanism leading to the imperfect cancellation of the sensed back-action noise of the detector.

2.5.1 Speedmeter Mechanism

There are multiple ways in which the speedmeter mechanism can be broken and the dominant mechanism depends on the speedmeter topology being considered.

For the ring Sagnac the main practical drawback is the high-level of scatter present in ring cavities. Photons that scatter from one field to another degrade the

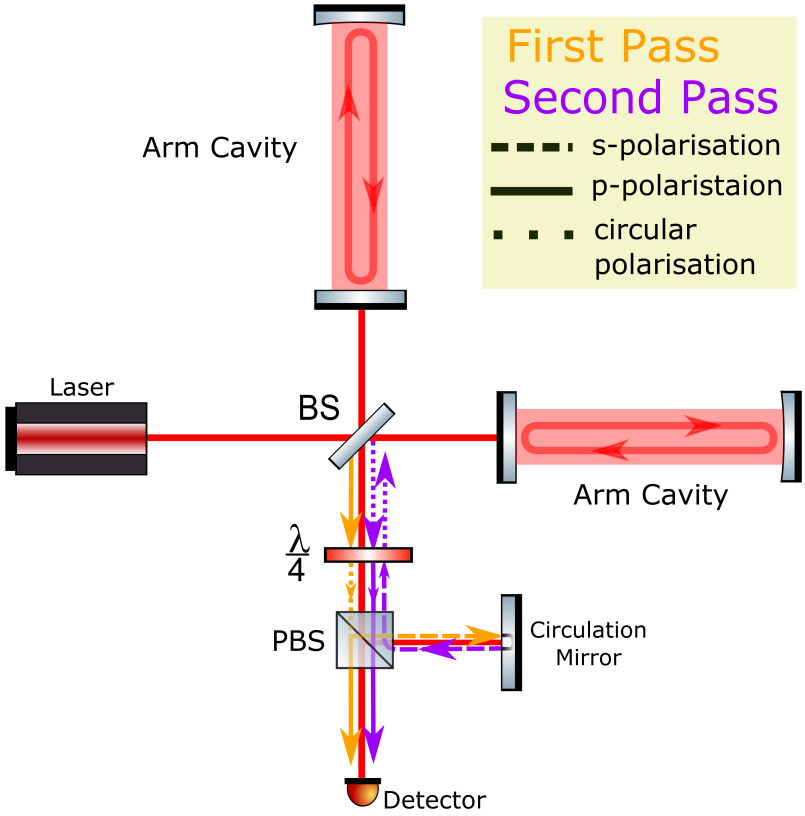


Figure 2.10: The polarisation circulation speedmeter. The core of the interferometer is a Michelson interferometer with the addition of polarising optics and a circulation mirror in the anti-symmetric output. The interferometer is illuminated with p-polarised light, at the output this is rotated into circular polarisation by the quarter waveplate (QWP), which is split at the polarising beam-splitter so some light is coupled to the detection port and some back into the interferometer via the circulation mirror. The light circulates a second time in the interferometer in a circular polarisation due to the second interaction with the QWP, and upon exiting the interferometer is rotated into p-polarisation by the QWP and is therefore transmitted by the PBS to the detection port.

speedmeter effect as they constitute a radiation pressure effect that is not cancelled coherently. A proof-of-principle experiment [81] based on the ring Sagnac will aim to show the cancellation of back-action noise and the benefits of the speedmeter. The sensitivity of the prototype is expected to be limited by scatter asymmetries [82] and higher-order laser modes [83].

For the Polarisation Sagnac speedmeter the quality of the polarising optics will strongly impact the speedmeter mechanism. This is the main consideration of this thesis and the investigation is presented in full in Part II.

2.5.2 Infrastructure

In addition to the technical limitations of the different speedmeter topologies, it is important to also consider the complexity of implementing each design. Several topologies (such as the Polarisation Sagnac, Polarisation Circulation Speedmeter) could be built within the vacuum envelopes of existing gravitational wave detectors and the triangular sloshing Sagnac could be realized in the Einstein Telescope vacuum envelope. Whereas, installing an additional sloshing cavity would prove difficult and expensive.

In general, one advantage of a sloshing type design is that all of the additional technical complexity occurs at the output of the main interferometer. A sloshing type speedmeter could be operated as a Michelson interferometer with the option of operating as a speedmeter with only changes to the readout optical layout.

The main technical drawback of the sloshing speedmeter interferometers is the need for an extra cavity and the requirement for a large amount of circulating optical power [69].

Of the schemes that require the least change to the Michelson interferometer, most rely on polarising optics to form the speedmeter. It would be possible for one of these designs to be built inside the existing vacuum envelopes of the advanced gravitational wave detectors. These are the polarisation Sagnac [73], the polarisation folding [76], and the polarisation circulation speedmeter [77].

Part II

Experimental Investigation of the Polarisation Sagnac Speedmeter Topology

Chapter 3

Characterisation of Polarisation Optics for Speedmeter Interferometry

Of the topologies of speedmeter interferometer discussed in Section 2.4.1, the *Polarisation Sagnac* and the *Polarisation circulation* speedmeters rely on polarising optics as core elements of the interferometer. This chapter covers the characterisation of polarising optics for the polarisation Sagnac experiment (see Chapter 4) and discusses the effect of imperfect polarisation in an interferometer. In Section 3.1 a description of each type of optic is given and bench-top experiments to characterise the polarising optics are presented in Section 3.3. The implications of imperfect polarising optics on the sensitivity of an optical speedmeter are discussed in Section 3.4.

3.1 Polarising Optics for the Sagnac Speedmeter

A polarisation Sagnac speedmeter (see Figure 2.7) includes a polarising beam-splitter between the two arms of the interferometer and a quarter waveplate between the beam-splitter and each linear cavity to ensure that light from the two counter-propagating fields is directed in the desired path around the interferometer.

Polarising Beam-splitter

A polarising beam-splitter (PBS) is a polarising optic used to split light into two beams of orthogonal linear polarisations (see Figure 3.1).

The plane of polarisation of the optic is defined by the plane that contains the normal to the optical surface of the polariser and the incident beam. Light of linear polarisation parallel to this plane will be transmitted by the polariser and light in the orthogonal polarisation will be reflected.

Light polarised parallel to the plane of polarisation is defined as p-polarised light, and light of the orthogonal linear polarisation is said to be s-polarised [47].

A PBS is characterised by the polarisation extinction ratio (PER) in transmission and reflection. The transmission extinction ratio is defined as the ratio of transmitted power ($T_{s,p}$) in p-polarisation and s-polarisation

$$\mathrm{PER}_T = \frac{T_p}{T_s} \tag{3.1}$$

and similarly, in reflection

$$\mathrm{PER}_R = \frac{R_s}{R_p}. \tag{3.2}$$

A PBS optic can then be described by the Jones matrices

$$\mathbf{J}_{\mathrm{PBS}_T} = \begin{bmatrix} \sqrt{T_p} & 0 \\ 0 & \sqrt{T_s} \end{bmatrix}, \tag{3.3}$$

and in reflection

$$\mathbf{J}_{\mathrm{PBS}_R} = \begin{bmatrix} \sqrt{R_p} & 0 \\ 0 & \sqrt{R_s} \end{bmatrix}. \tag{3.4}$$

If we assume that the extinction ratios PER_T and PER_R are much larger than 1, the Jones matrices for the PBS can be re-written in the form

$$\mathbf{J}_{\mathrm{PBS}_T} \simeq \begin{bmatrix} \sqrt{1 - \frac{1}{\mathrm{PER}_T}} & 0 \\ 0 & \sqrt{\frac{1}{\mathrm{PER}_T}} \end{bmatrix}, \tag{3.5}$$

and in reflection

$$\mathbf{J}_{\mathrm{PBS}_R} \simeq \begin{bmatrix} \sqrt{\frac{1}{\mathrm{PER}_R}} & 0 \\ 0 & \sqrt{1 - \frac{1}{\mathrm{PER}_R}} \end{bmatrix}. \tag{3.6}$$

The angle of incidence on the polarising interface is defined by the angle between the normal to the polarising surface and the incident beam vector. The extinction ratio of a polarising optic depends on the angle of incidence. Most polarisation beam-splitters are specified for an angle of incidence of 45°, but the best extinction ratio may be at a different angle of incidence.

When light is incident on a boundary between two materials of different refractive index, the fraction of light reflected depends on the polarisation of the light and the angle of incidence. The coating layer is designed to work for a

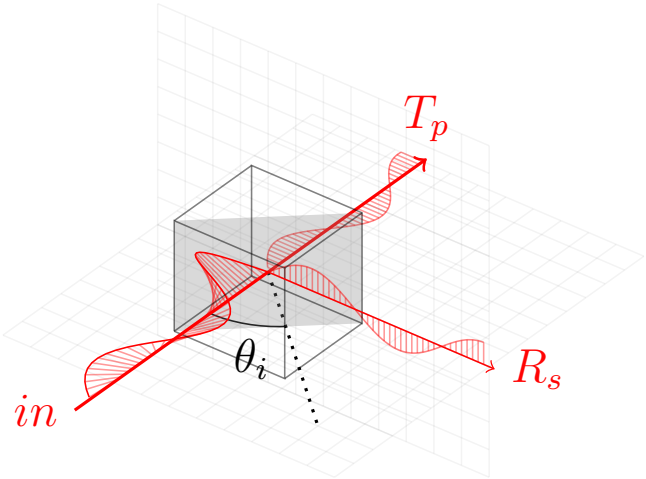


Figure 3.1: Basic principle of a polarising beam-splitter. The input field of arbitrary polarisation is split into two beams of orthogonally polarised light in reflection and transmission of the optical surface (gray shaded area).

specific wavelength of light and angle of incidence [84].

In a PBS, the angle of incidence of the beam at the interface between a dielectric coating and the optical substrate is equal to Brewster's angle [49], so the reflected beam is polarised parallel to the plane of incidence.

PBS optics can be packaged as PBS cubes or as flat thin-film polariser optics. To make a PBS cube two prisms are bonded together with one of the contacted surfaces coated with the multi-layer dielectric that forms the polarising surface. The other surfaces typically have anti-reflection coatings. A thin-film polariser is fabricated by depositing the polarisation coating onto a plate substrate.

For the polarisation Sagnac speedmeter investigation an order of thin-film polarisers were purchased from the company Layertec [85]. These optics had a specified extinction ratio in transmission of 2500 for light of 1550 nm wavelength at an angle of incidence of 45°. The fused silica substrates have a diameter of 25 mm with an anti-reflection coating on the surface opposite to the polarisation coating.

Quarter Waveplate

A quarter waveplate is an optical component built around a birefringent material. A birefringent optic exhibits a different index of refraction in different angular orientations. The orientations corresponding to maximum and minimum index of refraction known as the *fast* and *slow* axes respectively. The birefringent element of a QWP is of a certain thickness that ensures that the phase retardation imparted on light aligned to the slow axis is one quarter wave (90°) relative to light

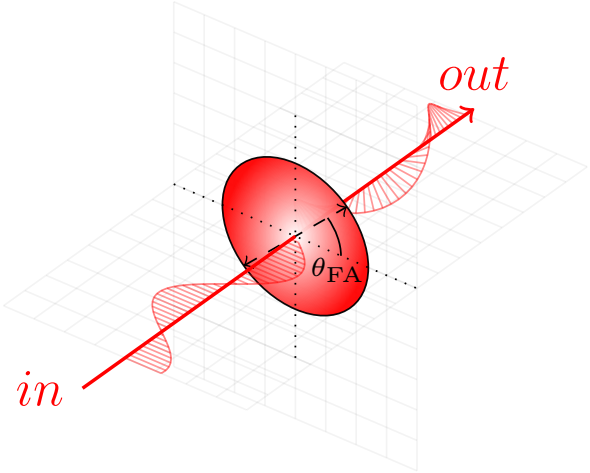


Figure 3.2: Basic principle of a quarter waveplate. The input field of linear polarisation is rotated into a circular polarisation by the QWP with fast axis orientated at 45° to the input polarisation.

aligned to the fast axis.

A QWP can be used to transform linearly polarised light to circularly polarised light if the fast axis is orientated at 45° relative to the polarisation of the light. This process is the maximum amount of rotation a QWP can produce (see Figure 3.2). Similarly, a QWP can also convert circularly polarised light into a linear polarisation.

A generalised QWP of birefringence η is represented by the Jones matrix

$$\mathbf{J'}_{\mathrm{QWP}} = \begin{bmatrix} e^{i\frac{\eta}{2}} & 0 \\ 0 & e^{-i\frac{\eta}{2}} \end{bmatrix}, \tag{3.7}$$

and a QWPwith fast axis orientated at an angle θ with respect to the horizontal axis of the reference frame is represented by the Jones matrix in Equation 3.7 pre- and post-multiplied by rotation matrices

$$\mathbf{J}_{\mathrm{QWP}} = \mathbf{Rot}(-\theta) \cdot \mathbf{J'}_{\mathrm{QWP}} \cdot \mathbf{Rot}(\theta) \tag{3.8}$$

$$= \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \sin\theta \end{bmatrix} \cdot \begin{bmatrix} e^{i\frac{\eta}{2}} & 0 \\ 0 & e^{-i\frac{\eta}{2}} \end{bmatrix} \cdot \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}. \tag{3.9}$$

An ideal QWP with fast axis orientated at 45° relative to the horizontal plane of reference can be represented as a Jones matrix of the form

$$\mathbf{J}_{\mathrm{QWP}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix}. \tag{3.10}$$

An imperfect QWP with birefringence η, orientated at 45° is represented by

the matrix

$$\mathbf{J}_{\mathrm{QWP}} = \begin{bmatrix} \cos\frac{\eta}{2} & i\sin\frac{\eta}{2} \\ i\sin\frac{\eta}{2} & \cos\frac{\eta}{2} \end{bmatrix}. \tag{3.11}$$

The thickness of the birefringent layer defines the retardation of the waveplate for a specified wavelength of light. A QWP optic can either be produced as a zero-order or multiple-order waveplate. A zero-order waveplate has a thickness equivalent to one quarter wave, whereas a multiple-order waveplate has a thickness equivalent to an integer number of wavelengths plus one quarter wave to produce the same result.

Multiple-order plates can be easy to fabricate but have a high sensitivity to temperature changes.

Zero-order plate optics of a single quarter wave thickness would be hard to fabricate. Instead, a zero-order waveplate can be formed by combining two multiple-order plates with relative alignment such that they cancel out all but one quarter wave effect.

The waveplates purchased for the polarisation Sagnac speedmeter experiment were optically contacted zero-order waveplates of a specified tolerance of $\frac{\lambda}{4} \pm \frac{\lambda}{300}$ for 1550 nm light with a 20 mm aperture from Union Optic [86].

3.2 Birefringence of Non-Polarising Optics for the Sagnac Speedmeter

As well as the polarising optics, birefringence in the non-polarising optical elements will also impact the performance of the polarisation Sagnac speedmeter as unwanted birefringence will change the polarisation state of the light in the interferometer.

3.2.1 Cavity Mirrors

In a Fabry-Pérot cavity the light reflects from the cavity mirrors at normal incidence so there should be no polarising effect due to the reflection of the beam from the optical surface. However, it is expected that mirror coatings of the type of mirror used in optical interferometers will exhibit some level of inherent birefringence.

To produce the highly reflective mirrors needed for a gravitational wave interferometer, multi-layer dielectric mirrors are used. These consist of a stack of thin

layers of alternating high and low index of refraction material deposited onto a substrate. A fraction of the incident light is reflected at each surface of the stack and the thickness of each layer is nominally one quarter of a wavelength, so that the reflected light for each surface interferes constructively. The more layers utilised, the higher the reflectivity of the overall mirror, and the greater the difference in index of refraction between each layer, the higher the fraction of light reflected at each surface.

Birefringence in the optical coating of mirrors can occur due to thermal stress between the multi-layer of the coating during the fabrication process [87].

Mirror coating birefringence has been measured by multiple experiments, mainly as investigations into noise sources and limitations to magnetic vacuum birefringence measurements. These projects such as the PVLAS [88] and BMV [89] experiments use high sensitivity Fabry-Pérot interferometers to detect changes in the birefringence of gas and vacuum states induced by strong magnetic fields as a test of quantum electrodynamics [90].

Measurements of mirror coating birefringence from multiple investigations were reviewed in [91]. The results reported were in respect to mirrors fabricated with different methods for a range of wavelengths. So from the results only a broad estimation can be made of the level of birefringence to be expected.

At the high end, coating birefringence of order 10^{-3} [92] and 10^{-4} [93] radians were measured. Whereas much lower levels were reported in [94] and [95], of order 10^{-6} radians per reflection, and as low as 10^{-7} in [87].

It is most likely that the level of birefringence in a multi-layer coating depends on the fabrication process. In one case [96] the birefringence was measured at multiple points across the mirror surface and the magnitude and orientation of the birefringence was seen the vary across the surface in a regular pattern. The level of phase retardation varied by a factor of two.

It is not currently known how specific aspects of the fabrication process influence the birefringence of the coating, but the working assumption is that the number of layers in the coating and the difference in thermal capacity between materials layers along with coating deposition process are the key factors.

So far, the birefringence of the cavity mirrors in a gravitational wave detector has no significant effect of the performance of the detector, but for polarisation depended speedmeter topologies it needs to be considered as a potentially limiting factor.

It was therefore decided to carry out a measurement of the intrinsic birefringence of the mirrors purchased for the polarisation Sagnac experiment as part of the linear cavity test of the polarisation Sagnac speedmeter. The effect of cavity

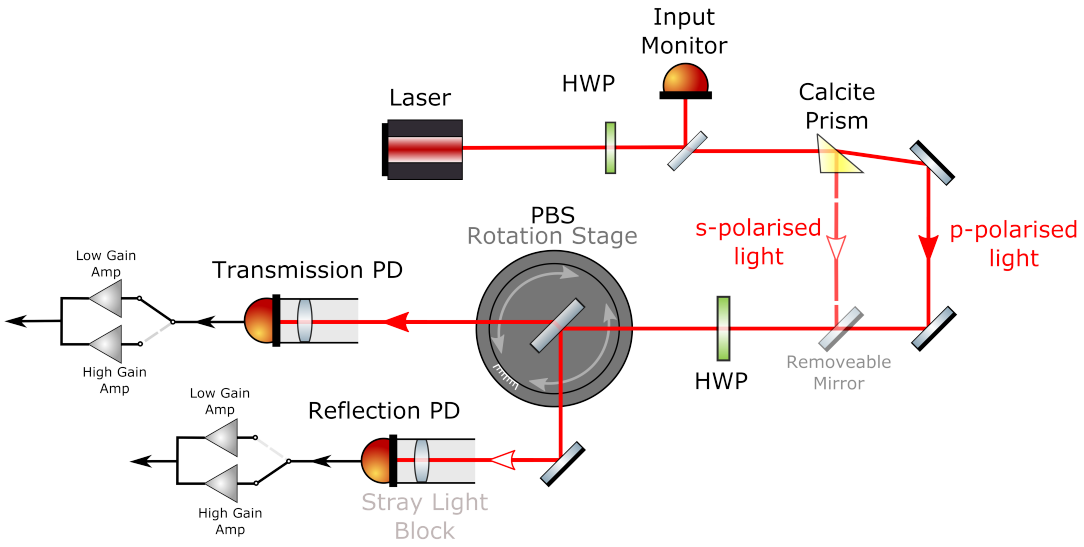


Figure 3.3: Experimental setup to measure the extinction ratio of a PBS. The PBS optic is mounted on a rotation stage so the extinction ratio as a function of the angle of incidence can be measured. The input light is prepared in a linear polarisation to match the polarisation axis of the PBS with a calcite prism and a half waveplate (HWP). The steering mirrors used for this experiment had previously been characterised and shown to have negligable polarisation effect. The input power is monitored with a PD in reflection from a 50:50 beam-splitter pick-off before the calcite prism, which is calibrated to the optical power in front of the PBS. The light incident on the PBS can either be s-polarised or p-polarised, and the transmission and reflection beams are focused onto photodiodes, an additional steering mirror is used in the reflection path. The photodiode signals are amplified, and the photodiode circuit has two gain settings suitable for amplifying milli-ampere and micro-ampere photocurrents respectively.

mirror birefringence is investigated in Chapter 4 and the results are analysed in Chapter 5. The impact of this effect on the performance of a polarisation Sagnac speedmeter is investigated in Chapter 6.

3.3 Characterisation of Polarising Optics

The PBS and QWP elements for the polarisation Sagnac linear cavity experiment were characterised with table-top experiments so that realistic optical parameters can be used in design studies of polarisation dependent interferometers. These tests aimed to measure the polarisation performance of the individual elements. Previous experiments were conducted as undergraduate student projects to characterise the polarisation optics; in this chapter new measurements are presented of the optical performance of the PBS and QWP used in the linear cavity experiment described in Chapter 4.

3.3.1 PBS

The experimental setup to measure the polarisation extinction ratio in transmission and reflection of the thin-film polarised PBS is shown in Figure 3.3.

To measure the polarisation extinction ratio of a PBS with a bench top experiment it is necessary to align the orientation of the optic to the probe light. In the test setup used here, the light is prepared in a linear polarisation using a calcite prism. Calcite prisms can be made with an extremely high polarisation extinction ratio [97]. In addition, a half waveplate (HWP) is used before the PBS to align the polarisation of the light to the PBS so as to maximise the extinction ratio. In this way the best performance of the optic could be investigated in a way that was not limited by misalignment between the calcite input polariser and the PBS being characterised.

The reflected and transmitted beams were recorded with separate photodiodes. Both photodiodes were calibrated for high and low gain settings that were designed to probe high extinction ratios by being able to measure light power of a few milliwatts or a few micro-watts.

A third photodiode (Input Monitor in Figure 3.3) was used as an input monitor in reflection of a beam-splitter in the input path of the beam to measure the optical power at the input to the experiment. Between measurements the relationship between optical power measured with this photodiode and optical power directly in-front of the PBS was measured by temporarily placing a calibrated power meter before the rotation stage on which the PBS was mounted. In this way the input monitor could be used as a calibrated readout of the power in-front of the PBS during the measurements.

The angle of incidence was read off a micrometer vernier scale leading to a very small uncertainty of $5'$ in the angle of incidence. The photodiode signals that measured the transmitted and reflected powers were recorded in CDS, a digital data acquisition system at a rate of 32 kHz. For each measurement approximately 5 seconds of recorded photodiode voltage was averaged, producing an accurate measure of the photodiode voltage with a standard percentage error of approximately 0.1%.

The largest source of uncertainty in the measurement of extinction ratio therefore is the calibration uncertainty of each photodiode. The photodiode used in this experiment were calibrated by measuring the power incident on the photodiode verses the photodiode output voltage for a range of laser powers. The power was measured with a calibrated power meter. The resulting plot of voltage against power was used to produce a calibration from a linear fit to the data (see

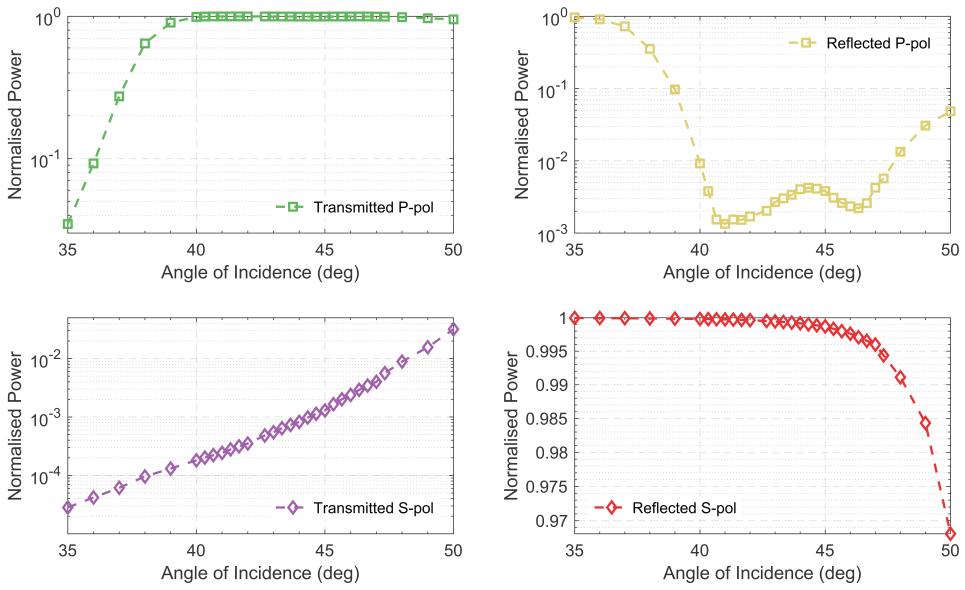


Figure 3.4: Extinction data recorded with s and p polarised light in transmission and reflection from the PBS normalised to the total input power.

Appendix B Figure B.3). The error associated with this fit was taken as the calibration uncertainty for each photodiode. The photodiode in reflection was found to have a calibration uncertainty of 0.15% in the low gain state and 0.55% in the high gain state. The photodiode in transmission was found to have a calibration uncertainty of 0.6% in the low gain state and 0.25% in the high gain state. This results in an uncertainty on the transmission extinction ratio of 0.85% and an uncertainty of 0.7% on the reflection extinction ratio. The photodiode voltage was also measured with no light on the detector to be able to subtract the effect of DC offset voltages from the recorded data.

The extinction ratio in transmission and reflection as a function of angle of incidence on the PBS was measured for both s-polarised light and p-polarised light and the results are shown in Figures 3.4 and Figure 3.5. At an angle of incidence of 45° the extinction ratio in transmission was measured to be 770 ± 6.5, and 262 ± 1.8 in reflection. The optimal angle of incidence was found to be around 41°, were the transmission extinction ratio was measured to be better than 4000, and 700 in reflection.

3.3.2 QWP

The experimental setup to measure the birefringence of a quarter waveplate is shown in Figure 3.6. The birefringence of the QWP is measured as the induced

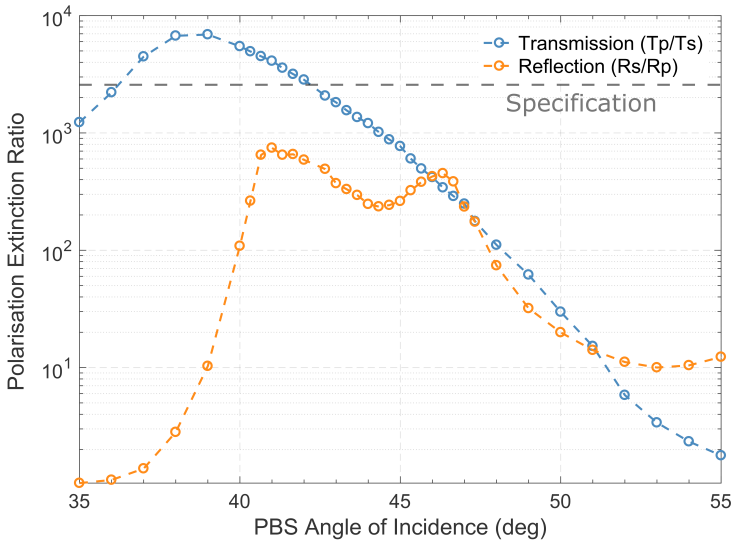


Figure 3.5: Measured polarisation extinction ratio of the PBS in transmission and reflection as a function of the angle of incidence on the PBS optic. The manufactures specification of a transmission extinction ration of 2500 is shown as a grey dashed line.

change in polarisation as a function of the rotation of the QWP. The PBS characterised in Section 3.3.1 was used to analyse the polarisation of light after the QWP.

In this setup an ideal quarter waveplate when at an angle of 45° relative to the linearly polarised light should produce the maximum amount of polarisation change and result in an equal amount of optical power reflected and transmitted by the PBS. The extinction ratio of the PBS having already been characterised can be taken into account when analysing the results of this experiment.

The results of the QWP characterisation for the optic at normal incidence are shown in Figures 3.7 and 3.8. Figure 3.8 shows that the birefringence of the QWP is more than a quarter wave. In the plot the modelled result of a QWP with birefringence of $\frac{\lambda}{4} + \frac{\lambda}{324}$ (or 1.590 radians) is fitted to the experimental data using the Jones matrices described in Section 3.1.

Similar to the PBS characterisation the measurement of angle of incidence and photodiode voltage have negligable uncertainty associated with the respective measurements. The angle of the QWP is read from a rotation stage accurate to 1°. The largest source of uncertainty is due to the fitting of the model to the data. A basic difference of least squares method is used to determine a quality factor (or χ^2) of the fit. The best fit value shown in Figures 3.7 and 3.8 has a quality of fit of $\chi^2 = 2.08 \times 10^{-9}$; this was used to establish a region of confidence to estimate the uncertainty on this result. In this way, the birefringence of the

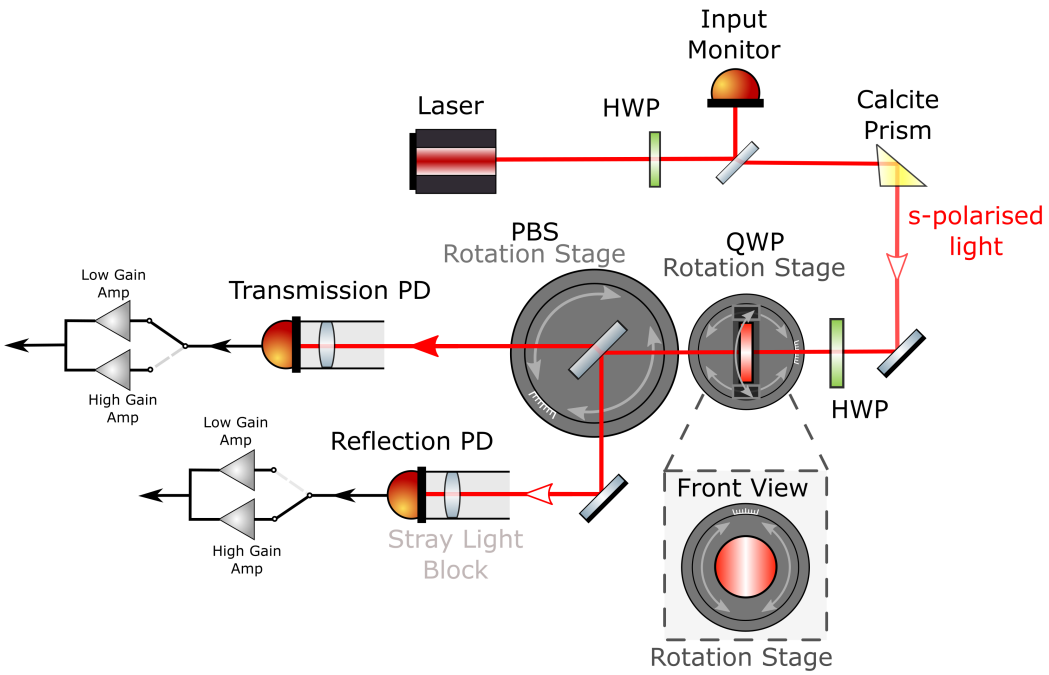


Figure 3.6: Experimental setup to measure the birefringence of a QWP, the setup is similar to the PBS characterisation setup shown in Figure 3.3 with the addition of a QWP on two rotation stages, one to change the roll angle of the optics and one to change the yaw. In this way the transmission and reflection of the PBS can be measured as a function of the QWP fast axis orientation and angle of incidence. The input beam is prepared in a linear s-polarisation meaning the light is almost completely reflected at the PBS. An ideal QWP at 45° relative to the input polarisation, would rotate the linear polarisation into circular polarisation such that an equal amount of light is detected in reflection and transmission of the PBS.

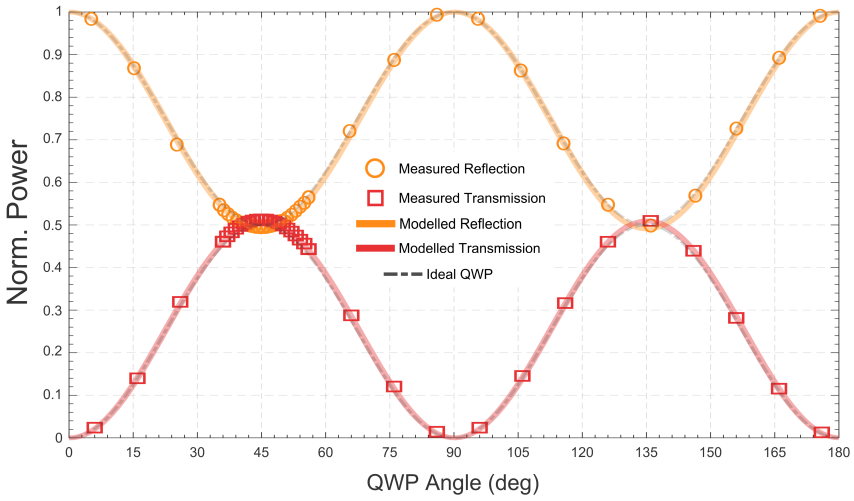


Figure 3.7: Results of the QWP characterisation at normal incidence. The power in reflection and transmission of the PBS is normalised to the total input power and plotted as a function of QWP angle for a half turn of the optic through 180°.

QWP can be said to be 1.590 ± 0.002 radians within a 95% statistical confidence based on the fit to the data.

The angle of incidence of the QWP affects the birefringence as the optical path length of the beam in the birefringent optic is increased. As the angle of incidence is increased a greater phase retardation is produced. The measured performance of the QWP at three different angles of incidence is shown in Figure 3.9.

3.4 Summary

In a polarisation Sagnac speedmeter interferometer the core optics include a PBS to direct light of orthogonal polarisation between the arm cavities and one QWP in each cavity to rotate the polarisation of the light before and after resonating in each cavity.

For the polarisation speedmeter interferometer experiments in the Glasgow 10 m prototype, thin-film polariser PBS optics and quarter waveplate optics were purchased and tested. In this chapter the characterisation process for each optic was described and the results presented. These optics were then used in the linear cavity test of the polarisation Saganc speedmeter that is described in Chapter 4. The results presented in this chapter are used in Chapter 6 in a model of a polarisation Sagnac speedmeter with realistic polarisation optics.

The PBS optics had a specified polarisation extinction ratio of 2500. The optic measured with a bench top set up was found to meet this specification

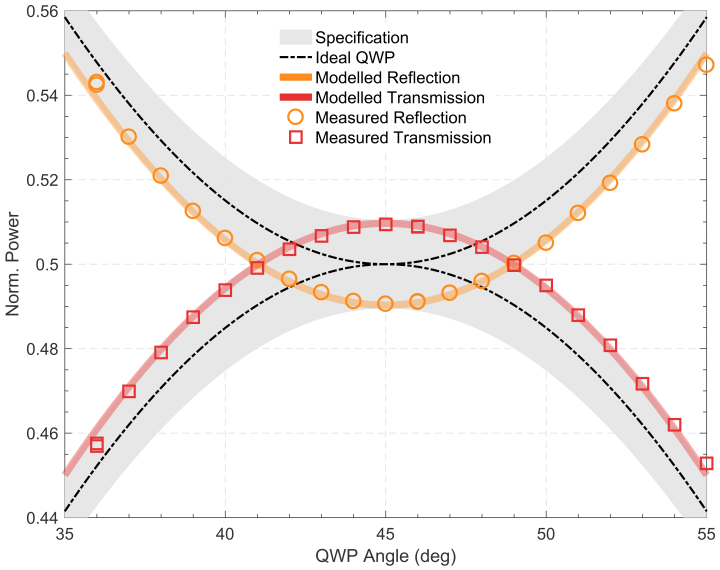


Figure 3.8: Results of the QWP characterisation at normal incidence with the QWP fast axis between 35° and 55°. A model of a QWP using the Jones matrix in Equation 3.8 fits a result to the experimental data using a birefringence of $\eta = 1.5901$. The equivalent performance of a perfect quarter waveplate with $\eta = 1.5708$ is also shown (grey dashed line) along with the specified tolarance range of $\eta = \frac{\pi}{2} \pm \frac{1}{150}$ (grey area).

in transmission with an angle of incidence of 42° with a measured extinction ratio of 2839 ± 24. The extinction ratio in reflection was lower with a maximum of 745.6 ± 5.2 at 41°. At an angle of incidence of 45° the extinction ratio in transmission was measured to be 770 ± 6.5, and 262 ± 1.8 in reflection.

The finite extinction ratio of the PBS will affect the performance of the polarisation speedmeter as a small amount of light will leak through the PBS between arm cavity interactions.

The QWP was measured to have a birefringence greater than a quarter wave. The measured data at normal incidence was consistent with a model of a QWP with a birefringence of $\frac{\lambda}{4} + \frac{\lambda}{324}$ or 1.590 ± 0.002 radians.

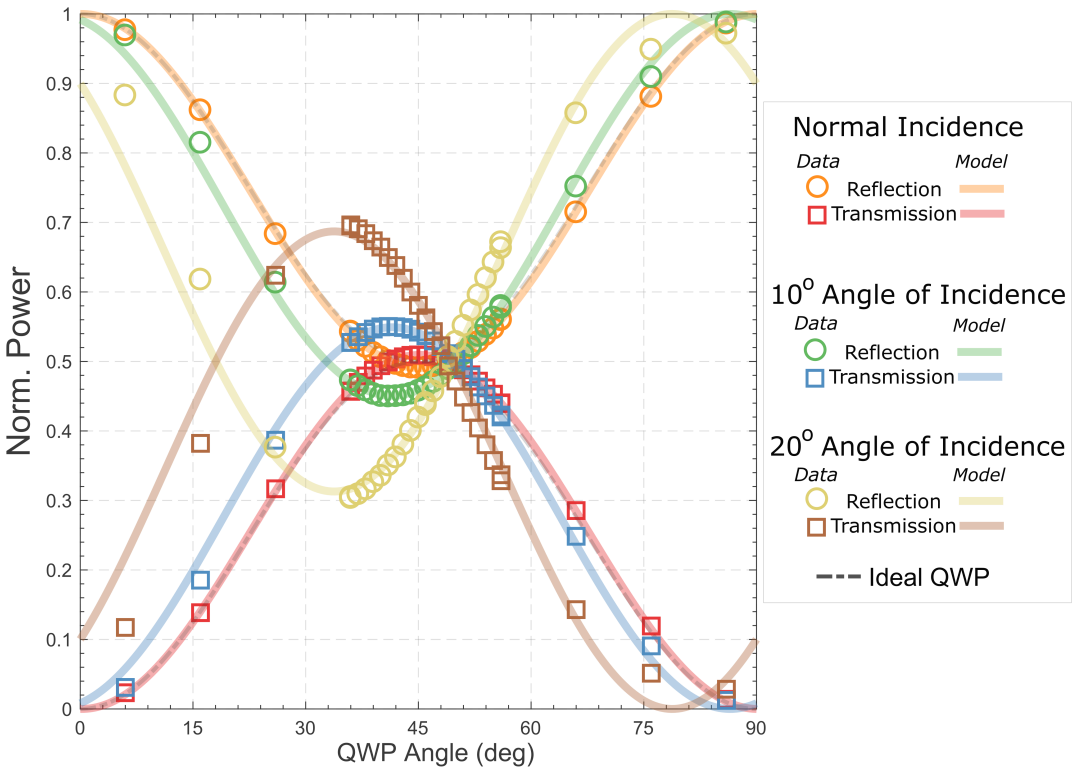


Figure 3.9: Results of the QWP characterisation at difference angles of incidence (AOI), showing the experimental results and the modelled results. It was found that the measured QWP at 10° angle of incidence had an equivalent birefringence of $\eta_{10^\circ} = 1.669$ radians and $\eta_{20^\circ} = 1.954$ radians at 20° angle of incidence.

Chapter 4

Linear Cavity Test of the Polarisation Sagnac Speedmeter Concept

The polarisation Sagnac speedmeter [73] is an interferometer topology proposed for a future gravitational wave detector. The technological requirements of such an interferometer need first to be demonstrated on a prototype scale to show the viability of such a scheme.

One vein of research was to construct a single arm of a polarisation Sagnac speedmeter with a baseline of 10 m in the Glasgow prototype system to investigate cavity birefringence effects. This allows us to determine if the polarisation of the light resonant in the linear optical cavities is influenced by low-levels of birefringence in the cavity mirrors, and compare with models that can be applied to full-scale detector topologies.

The linear cavity experiment and initial results are discussed in this chapter and a detailed analysis of the results can be found in Chapter 5.

Glasgow 10 m Prototype Facility

The Glasgow 10 m prototype laboratory is a facility dedicated to experimental demonstration of advanced interferometer techniques for future gravitational wave detectors. The clean-room facility includes two vacuum systems with 10 m and 1 m long baselines. The 10 m system is comprised of nine vacuum tanks, of a similar design to that used in the GEO gravitational wave detector (see Appendix A.1). The prototype has the capacity to house multiple experiments running in parallel. The polarisation Sagnac is one such experiment that utilises one arm and part of the central vacuum vertex (see Figure 4.1).

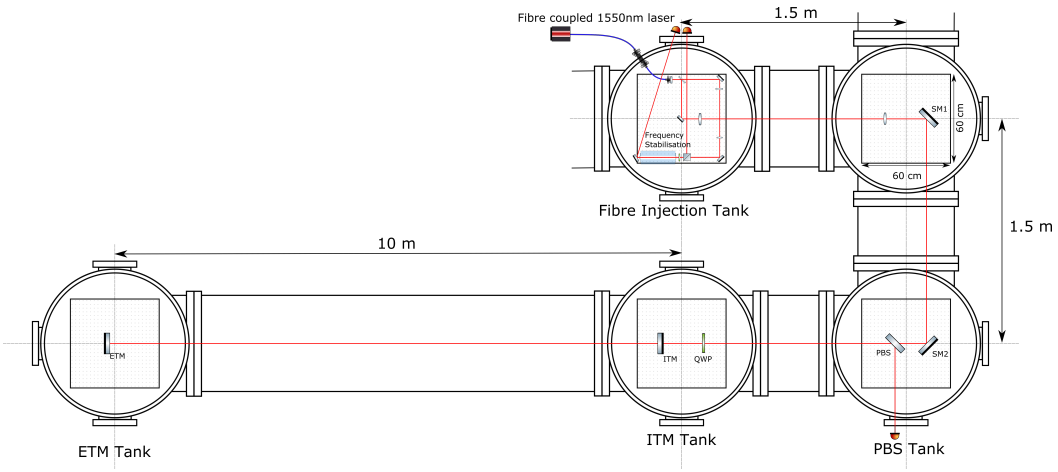


Figure 4.1: The single arm polarisation Sagnac speedmeter experiment in the Glasgow 10 m prototype system. The light is injected via a fibre feed-through shown in the upper middle tank. In this tank are steering mirrors and a reference cavity used for laser frequency stabilisation. The beam is directed to the 10 m arm by two suspended steering mirrors, and mode-matched to the cavity with a two lens telescope. In front of the cavity, are the polarising beam-splitter (PBS) and the quarter waveplate (QWP); these polarisation optics are suspended. The linear cavity is formed by the suspended input test mass (ITM) mirror and end test mass (ETM) mirrors.

Cavity Mirrors	Normalised Transmission	Curvature
ETM	500 ppm	15 m
ITM	2500 ppm	Flat with wedge of 30’

Table 4.1: Key properties of cavity optics for the single arm polarisation Sagnac speedmeter experiment

4.1 Design of the Single Arm Polarisation Sagnac Speedmeter Experiment

The optical layout of the linear cavity experiment is shown in Figure 4.1, the properties of the core optics are summarised in Table 4.1 and the cavity parameters in Table 4.2. This implementation, of what is in effect one arm of a polarisation Sagnac interferometer, would allow the testing of the suspended polarising optics and measure cavity birefringence effects and provide an assembled test-bed for future polarisation Sagnac speedmeter experiments.

4.1.1 Suspension Design

For the single arm test, six new suspensions were built and installed in the 10 m vacuum system to seismically isolate the cavity mirrors, steering mirrors and

Parameter	Value
Length	10 m
Free Spectral Range (FSR)	15 MHz
Full-width half-maximum (FWHM)	7.165 kHz
Finesse	2092
Ratio of reflected to transmitted power	44.5:55.5
Beam radius at ITM	1.868 mm
Beam radius at ETM	3.235 mm

Table 4.2: Key cavity parameters of the single arm polarisation Sagnac speedmeter experiment

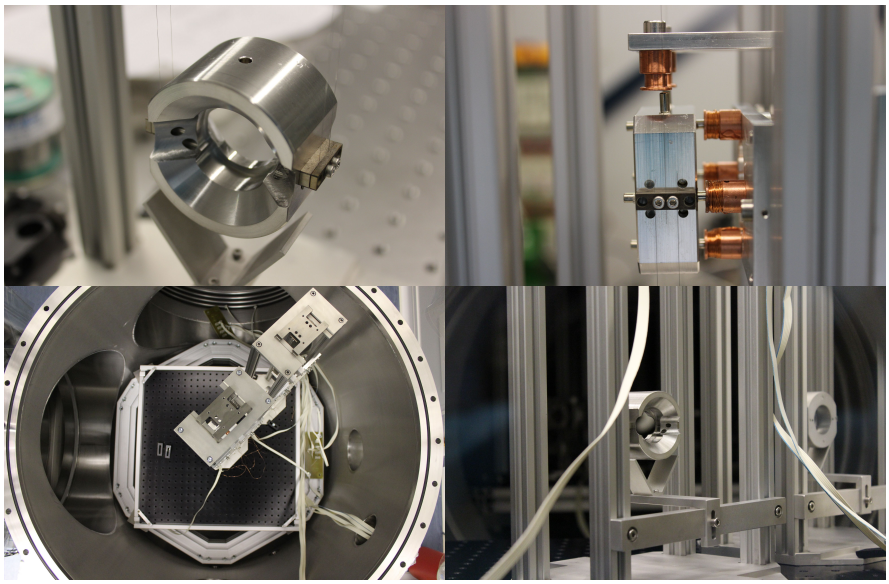


Figure 4.2: Suspensions for the polarisation Sagnac speedmeter. Top left shows a suspended optic holder test mass. Top right shows coil-magnet actuator for alignment control at the upper mass stage. Images in the bottom row show suspensions installed in the vacuum system, on 600 mm×600 mm breadboards.

the polarising optics for the linear cavity experiment. These suspensions use a double pendulum design created for previous experiments in the prototype lab (see reference [98] or Appendix A for more detail on the suspension design).

The suspensions are actively controlled with magnet-coil actuators at the upper and lower mass stage. Magnets are bonded to the suspended masses in line with coils of copper wire mounted on the suspension support frame. When a current is driven through the copper wire a magnetic field is generated to push or pull the neighbouring magnet. The coil-magnet actuators of the intermediate stage can be seen in the upper-right image in Figure 4.2. The coil-formers on the intermediate mass are made of copper and provide passive eddy current damping of the motion of the mass.

Feedback to the intermediate stage is used for alignment control, while actuation of the lower stage is used for cavity length actuation to lock the cavity to the frequency of the laser.

In addition to the construction and installation of the suspensions to the specifications of the existing designs, special modifications had to be made to the suspensions for the polarising optics.

PBS Suspension Modification

The PBS suspension was implemented with the optic at 45° to the input beam, presenting a restricted aperture and leading to a concern that the beam would be significantly clipped as it passes though the metal optic holder.

The suspended PBS is used to separate the input beam from the beam returning from the cavity. The beam must pass through the metal optic holder and be reflected from the reverse face of the optic without clipping. So three ports of the beam-splitter at 45° must present a clear aperture to the incident laser beam. In a full speedmeter layout all four ports would be used.

The PBS optic is held in a metal mass as shown in Figure 4.3. The apparent aperture of the suspended optic holder will be a combination of the ellipse of the near and far side of the metal mass as viewed at 45°. The beam transmitted through the optic will be refracted through the substrate. To ensure minimal clipping of the reflected beam the laser has to be aligned so that the spot is centred on the reverse, cavity facing, surface of the optic.

The change in optical power (P_{Loss}) of a Gaussian beam passing through a circular aperture of radius r can be calculated as

$$P_{\text{Loss}} = 1 - \frac{P_{out}}{P_{in}} = \exp\left(\frac{-2r^2}{w_0^2}\right), \tag{4.1}$$

where w_0 is the radius of the beam at the aperture. The drop in power of 'loss' from a non-circular or irregular aperture is harder to compute but can be calculated numerically by applying a clipping mask to a 2D Gaussian beam profile.

A model of the suspension aperture was generated to calculate the clipping loss and alignment tolerance for the elliptical apertures of the suspended PBS. This showed that for the unmodified suspension, a loss of 3.3 ppm can be expected for a well aligned beam.

The model allowed for arbitrary modifications to the shape of the aperture, but only modifications that could be fabricated were considered.

Figure 4.3 shows the modelled clipping masks of the modified suspensions as seen by the beam with refraction taken into account. The model showed that

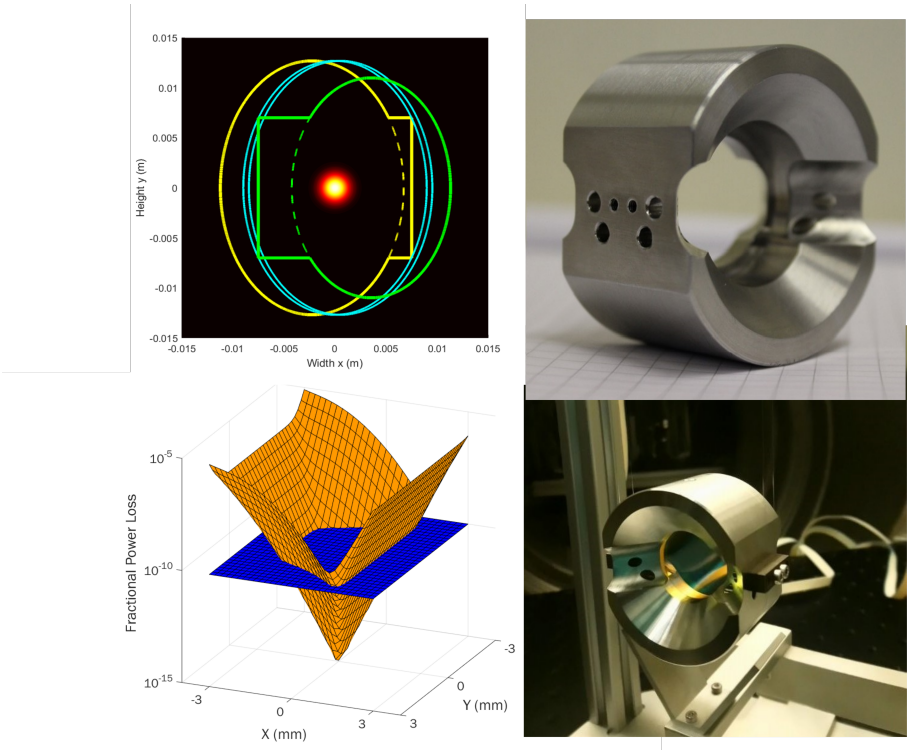


Figure 4.3: Modification of the suspension mass to hold the PBS to widen the aperture. The figure in the top left shows the modelling of the beam clipping from an irregular aperture. The green and yellow lines mark the aperture of the front and back faces of the metal mass respectively, as seen from the beam input side of the PBS with refraction taken into account. The dashed lines show the edges of the unmodified mass. The blue ellipses show the aperture of the front and back faces of the optic. The laser beam intensity profile is also shown. The figure in the lower left shows the model of alignment tolerance. The modified suspension mass is shown in the top right images and the suspended optic in the vacuum system is shown in the lower right.

using the modified suspended mass would reduce the amount of clipping loss to a level of 10^{-9} ppm for a well aligned beam. The model also showed that with the modified aperture, there was a misalignment tolerance of -1.2 mm and +2.3 mm in the x-direction, and ± 2 mm in the y-direction before the level of clipping would increase above 10^{-4} ppm, an arbitrary target set to ensure no significant loss in the application.

The metal mass was machined to remove material and widen the apparent aperture of the optic based on the model of clipping aperture. This modification in shown in the images on the right of Figure 4.3.

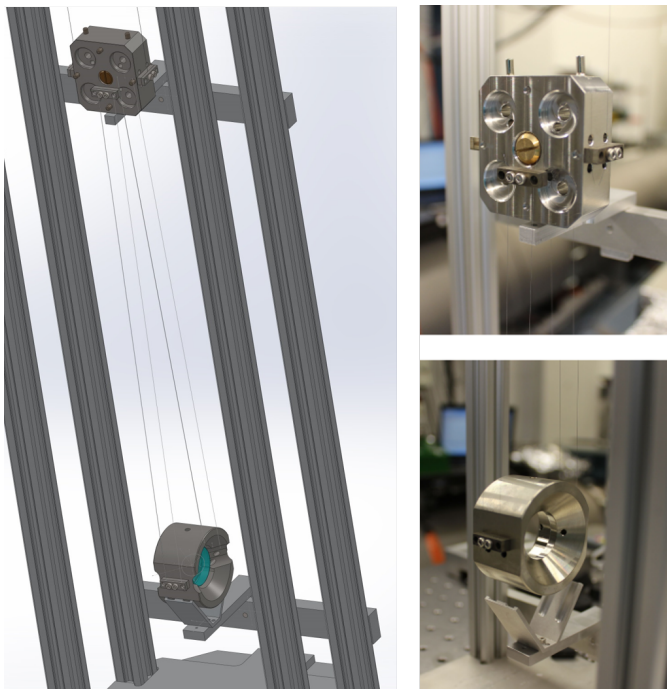


Figure 4.4: Modification of the QWP suspension to allow control of the optics roll degree of freedom. The image on the left shows the CAD design of the suspension courtesy of R. Jones. The images on the right show the modified suspension after construction with the hanging masses.

QWP Suspension Modification

The original suspension design allows for control of the pitch and yaw rotational degrees of freedom of the suspended masses with coil-magnet actuators (see Figure A.6 in Appendix A for a description of the suspension degrees of freedom).

For the suspended QWP the roll degree of freedom is of greater importance as this determines the alignment of the birefringence fast axis of the optic relative to the polarisation of the laser light. A modification to the suspension design was conceived to allow for control of the roll degree of freedom of the suspended QWP. The lower stage would be suspended at an angle of 90° relative to the intermediate mass so that control of the pitch of the intermediate mass would control the roll of the lower stage test mass.

A custom jig set-up was needed for the construction of this modified suspension, and the intermediate mass had to be modified to allow the metal wire to be clamped on the front and back face of the mass instead of to the sides, as in the original design. This modification is shown in Figure 4.4.

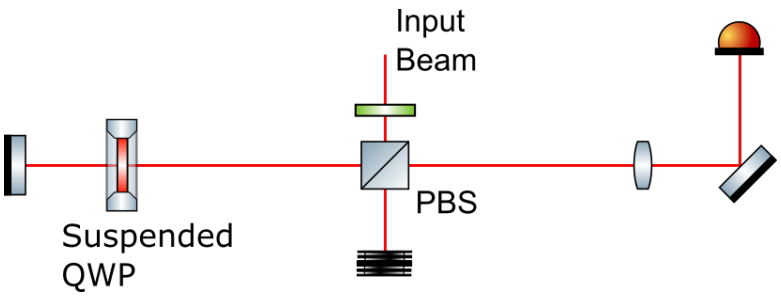


Figure 4.5: Test setup for aligning the QWP suspension.

Installing the QWP Suspension

The QWP optics needs to be suspended with the birefringence axis at 45° relative to the polarisation axis of the suspended PBS. To do this accurately was challenging as there is no fine control of the orientation of the optic in the mount and the optic can easily move during installation. The optic in the metal mass is held in place with a grub screw that pins the optics from above onto two pinning points formed in the shape of the metal mass.

The QWP suspension was removed from the vacuum system and placed on an optical bench where a test set-up was constructed. In this set-up a linearly polarised input beam double passed the suspended QWP by reflection from a folding mirror at normal incidence behind the QWP.

This mimicked the full linear cavity setup. If the QWP is mounted at the optimal angle the returning laser beam will be in the orthogonal polarisation to the input beam. The input and return beam were separated with a PBS and the return beam was measured with a photodiode. The rough angular alignment of the QWP was done by maximising the detected power, but for fine alignment a dither scheme was used.

If the QWP optic was rotated though a full 180° turn the detected power would follow a squared sinusoidal function. If the optic is set at some angle and a dither signal is applied to modulate the angle of the optics via the coil-magnet actuator, a variation in the detected power will be evident. If the angle of the QWP is set to be at the peak of the function no variation in power will be detected as the function is locally flat, provided the dither signal is of a sufficiently small magnitude.

The suspended QWP was modulated with a 1.25 Hz sine wave signal to dither the roll angle of the suspension. The photodiode readout was displayed on a spectrum analyser where a peak at the drive frequency and observed. The angle of the QWP in the suspension was then changed so as to minimise the dither peak.

With the angle of the QWP set in this way the suspension was reinstalled in the vacuum system. By setting the orientation of the suspended QWP relative to the table-top PBS the suspended QWP will now be orientated relative to the suspended PBS. This assumption is valid if the table-top PBS was mounted vertically and if the same is true for the suspended QWP and PBS that hang freely under gravity.

4.1.2 Ellipticity Readout

To investigate the effect of the of the resonant cavity on the polarisation of the light circulating in the cavity, a readout scheme was devised that looked at the polarisation of the light transmitted by the cavity. The polarisation of the transmitted light is taken as a direct measure of the polarisation of the light inside the cavity. The light in the reflection readout port is a mix of light returning from the cavity and light reflected directly from the ITM and therefore will be a mix of polarisation states.

In the ideal case, the polarisation inside the linear cavity is circularly polarised. The presence of birefringence in the cavity will rotate the light into an elliptical polarisation. The purpose of the ellipticity readout scheme is to measure the change in polarisation induced by the presence of the cavity.

The readout scheme consists of a PBS behind the ETM to split the transmitted beam into two beams of orthogonal polarisation that are then detected on separate photodiodes outside the vacuum system. The layout is shown in Figure 4.6.

The readout PBS was installed inside the vacuum system on a rigid mount and was orientated relative to the suspended PBS with the ETM removed so that no cavity was formed. The ellipticity was then measured with no cavity effects to provide the baseline to which the cavity measurements could be compared.

The two photodiode voltages are recorded along with other auxiliary monitor channels in CDS, a digital control and data acquisition system of the same type used in Advanced LIGO [99].

4.2 The Single Arm Polarisation Sagnac Speedmeter Experiment

4.2.1 Locking the Cavity

In order the run the single arm polarisation Sagnac speedmeter experiment the linear cavity needs to be locked to the frequency of the seed laser to ensure the

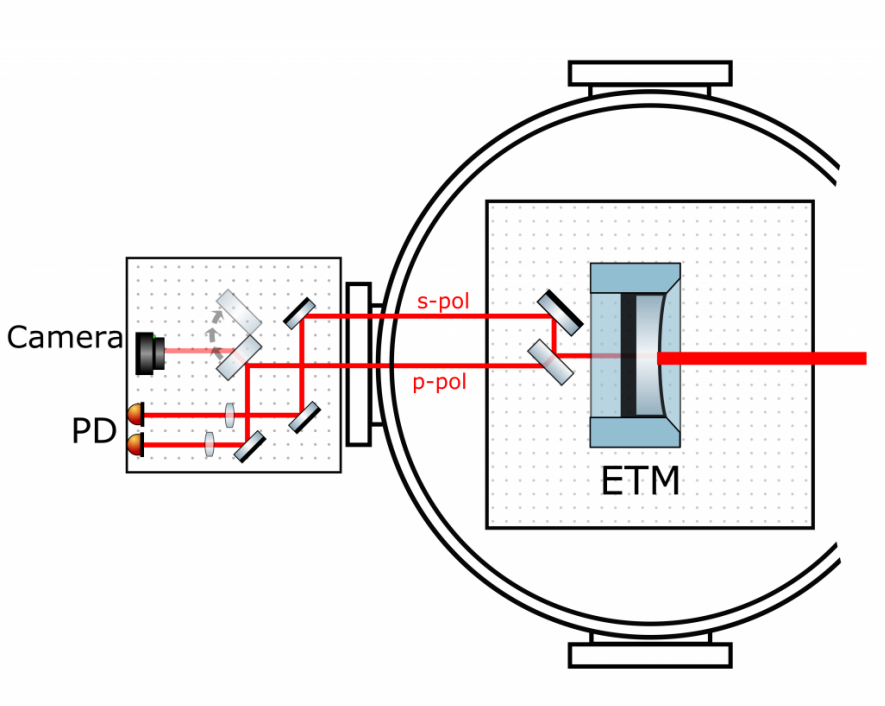


Figure 4.6: Optical layout of the readout scheme for ellipticity. On the vacuum side this includes the readout PBS aligned to the suspended PBS (not shown) and a steering mirror. This PBS separates and directs the beam transmitted though the suspended ETM to the readout breadboard on the air-side of the vacuum system. On the breadboard are additional steering mirrors to direct the two orthogonally polarised beams onto the readout photodiodes (PD). The two photodiodes are of the same electrical design. Also included is a camera sensitive to 1550 nm infrared light. The camera looks at the transmitted beam relative to the ETM mirror and is used for alignment of the 10 m cavity. The mirror in front of the camera is on a flip mount.

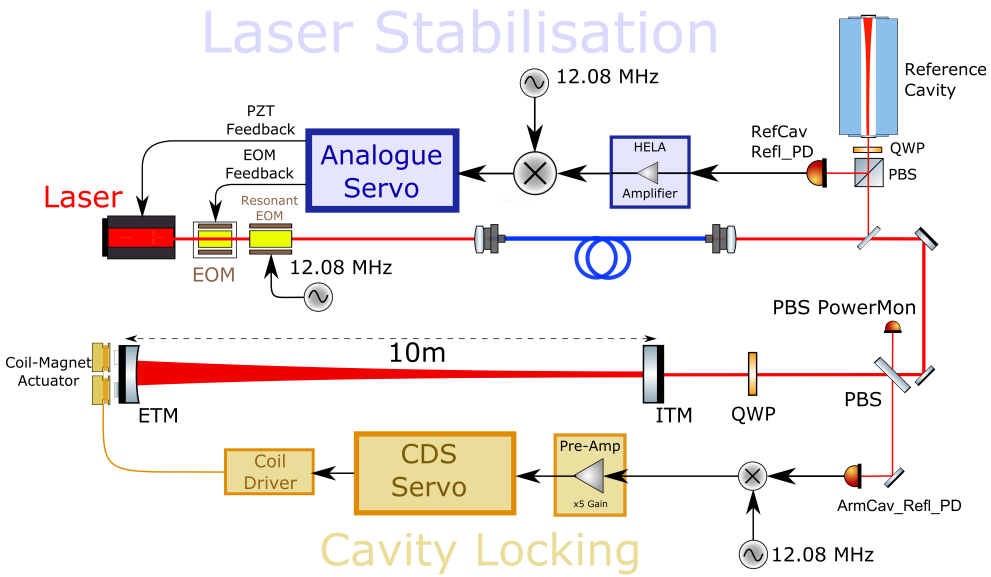


Figure 4.7: Cavity locking and laser stabilisation scheme. The frequency of the laser is stabilised to a rigid reference cavity and the 10 m cavity is locked to the stabilised laser frequency. The AC coupled signal from the PD in reflection from the 10 m cavity is demodulated with the LO signal to produce an error signal that is fed into CDS where digital filters are implemented. The output signal is applied to the coil-magnet actuators on the lower suspension stage to control the length of the suspended cavity.

light is resonant in the 10 m cavity.

The cavity is locked to the laser using a Pound-Drever-Hall stabilisation scheme [100] shown in Figure 4.7 (see Section 8.3.3 for more detail of this technique).

The laser used for the linear cavity experiment has a high level of frequency noise that requires active stabilisation. The laser and its noise performance are discussed in detail in Chapter 11, as is the frequency stabilisation scheme.

Without active reduction of the free-running laser frequency noise, it would not be possible to lock the suspended cavity due to contamination of the locking error signals with laser noise. This is illustrated in Figure 4.8 which shows the transmission of the suspended cavity and locking error signal with and without laser stabilisation.

The suspended cavity with no feedback, swings through multiple resonances with an approximate motion of one wavelength per second so each fringe is roughly 1 ms wide. In the case without active laser stabilisation the equivalent motion of the frequency of the laser is greater than the linewidth of the cavity resonance so the frequency of the laser passes though the resonance multiple times over the time scale that the cavity swings though that resonance. The laser frequency noise is imposed onto the cavity signals and dominates the observed transmission

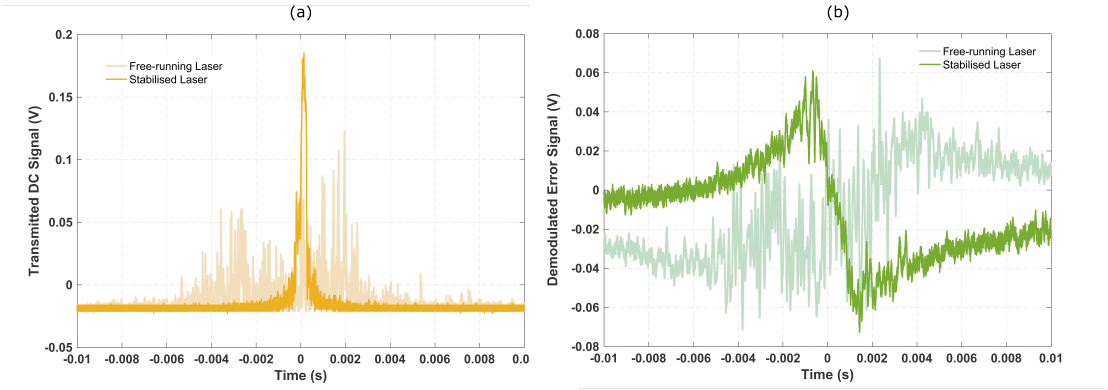


Figure 4.8: Signals to lock the 10 m cavity with and without active laser noise reduction. The plot on the left (a) shows examples of the transmitted DC signal indicating light resonant in the cavity and the right-hand plot (b) shows examples of the demodulated error signal in reflection of the cavity. This signal is filtered and used to lock the cavity on resonance with the laser.

and error signals, as is shown in Figure 4.8.

With active stabilisation the frequency noise of the laser is reduced to be below the linewidth of the cavity and so does not contaminate the locking error signal in the same way although laser noise is still evendent on the recorded error signals but the cavity can be locked to the laser.

Digital Locking Servo

The signal processing and frequency shaping for the suspended cavity control is implemented in CDS. The error signal generated from the demodulated photodiode signal is amplified and shaped to ensure that the feedback signal applied to the mass will hold the cavity locked stably.

To correctly shape the feedback signal, the actuation and sensing responses needs to be considered. Being a suspended mass in a pendulum system, actuation at frequencies from DC up the fundamental pendulum resonance frequency will have a flat response, meaning that current applied to the coils will translate linearly to an amount of positional change of the mirror mass. Above the pendulum resonance the response falls of as $1/f^2$, meaning current applied at frequencies above the resonance will produce much less motion of the mass. And at resonances the mass will move with a greater magnitude.

To form a stable control loop, the open loop response should pass though unity gain with a healthy phase margin. Below the unity gain point the length of the cavity is stabilised to the seed laser frequency. A unity gain point between 500 Hz and 1 kHz was chosen as a compromise between the suppression of low frequency motion of the cavity and the speed limitations of the digital control system.

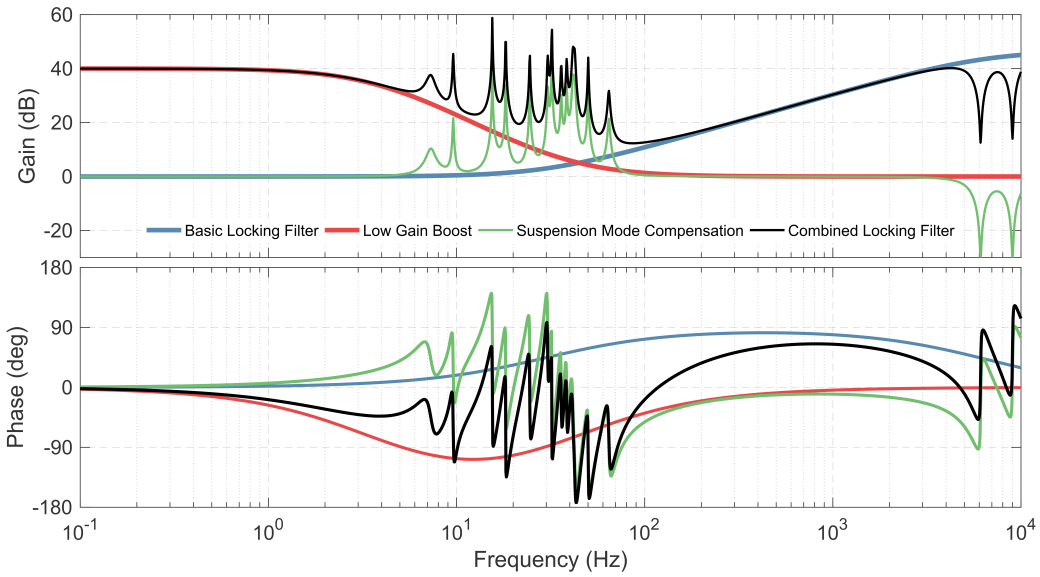


Figure 4.9: Transfer functions of the component filters used to lock the 10 m cavity. The basic transitional diffentiator (blue) is used to compensate for the pendulum response of the suspended mass. The additional transitional integrators (red) provide extra locking gain at low frequencies. The resonant gain filters (green) compensate for the 12 suspension modes. The combined locking filter with high frequency notch filters at 6 kHz and 9 kHz to damp mechanical oscillations of the mirror mass is shown in the black traces.

A transitional differentiator is implemented to shape the feedback to the suspension actuator to compensate the roll off due to the suspension and provide a form of phase-lead compensation, with a $1/f$ slope though unity gain. This is the basic filter used to lock the cavity. To increase the DC hold and stability of the lock, two transitional integrators were added to provide extra gain below 50 Hz to better hold the cavity length locked to the frequency of the laser. The combination of filters designed to lock the 10 m cavity is shown in Figure 4.9.

The suspension modes need to be considered when designing the locking loop. The suspended mirror has six coupled modes of motion (see Section A.3.1) three angular and three translational degrees of freedom in which the pendulum mass can move. A double suspension will therefore have 12 modes of motion in this suspension design these resonances are grouped between 7 Hz and 70 Hz.

To compensate for the enhanced motion of the suspension at these frequencies, resonant gain filters are implemented in the control model to reduce the residual motion concentrated in this frequency range and decrease the noise at the error point.

The suspension mode compensation filter, as seen in Figure 4.9 and in Figure 4.10 (a), was designed by measuring the power spectrum of the in-lock error point

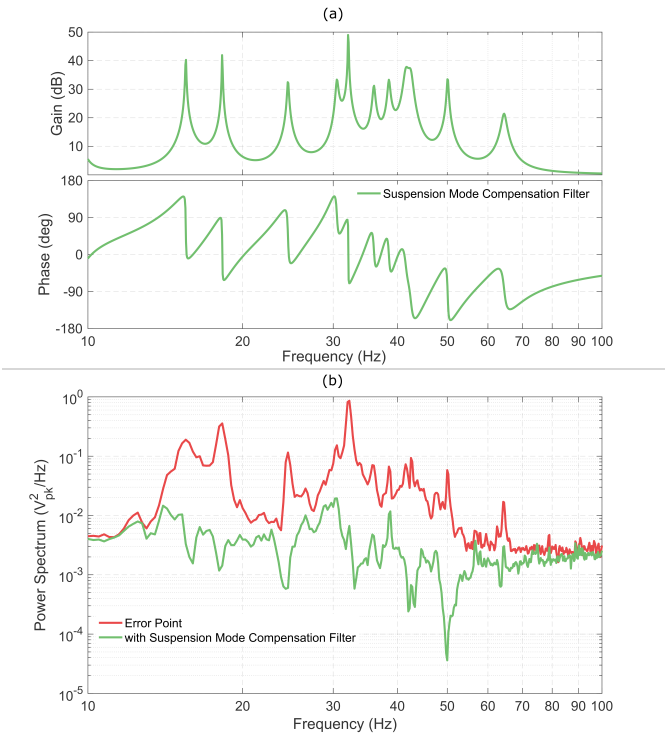


Figure 4.10: The transfer function of the filter designed to compensate for the suspension modes of motion of the cavity locking servo (a) and the effect of applying this filter to the in-lock error point noise (b) between 10 Hz and 100 Hz.

noise with the basic locking filter. In this spectrum, shown in Figure 4.10(b) multiple resonances were identified. From this measurement a filter block comprised of multiple resonant filters was designed and implemented.

In addition, mechanical resonances of the optic in the suspended mass appear at frequencies above 5 kHz. These resonance features can be damped with notch filters.

The CDS control model features a locking trigger to ensure that feedback signals are only applied to the suspension actuator when the cavity is close to a resonance. A block diagram of the control model can be found in Appendix D. The measured power transmitted by the cavity is used as a trigger. When the cavity swings though a resonance the power in the cavity increases. When this is detected the model will turn on the feedback which will actuate on the mirror position to hold the cavity on resonance so that a steady level of light is transmitted though the cavity. If the cavity moves off resonance the transmitted power will drop and the trigger will turn off the feedback as the cavity falls out of lock. This prevents the mirror getting a large kick as the cavity unlocks and helps to automatically re-lock the cavity.

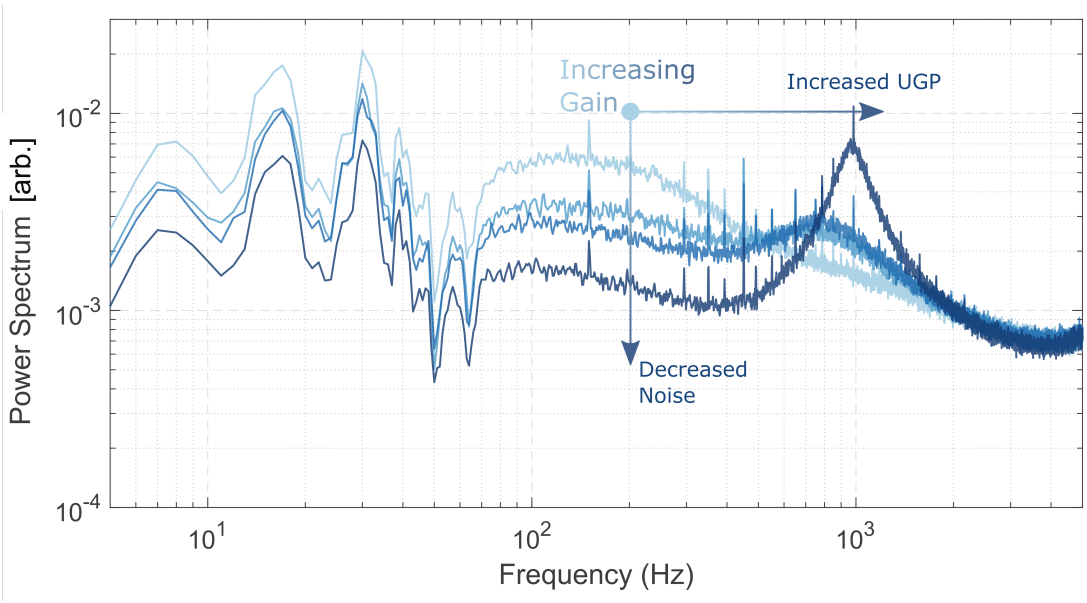


Figure 4.11: In-lock error point noise power spectrum. The shading of the traces represents the increase in loop gain over an order of magnitude with the darkest blue trace being the highest gain state. In the high gain state the unity gain instability around 1 kHz is starting to ring up due to reduced phase margin.

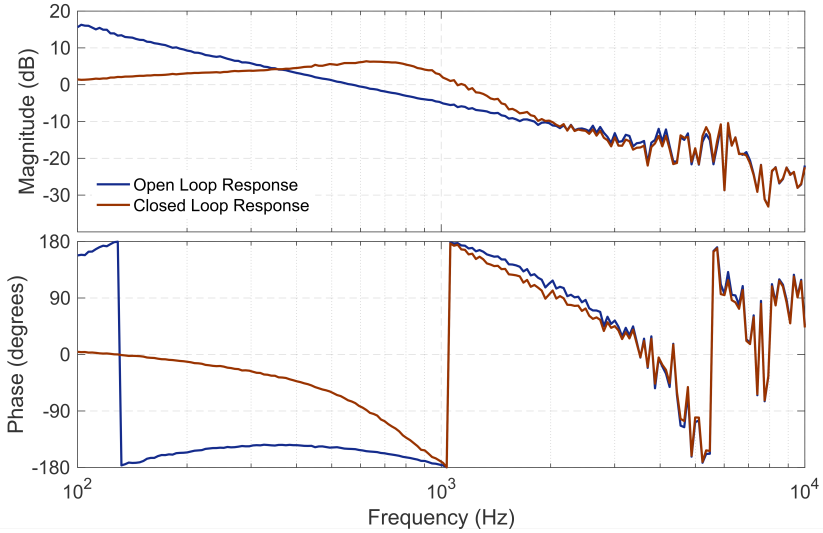


Figure 4.12: In-lock transfer functions of the 10 m cavity locking loop. The closed loop response was measured with the loop locked and unlocked and the open loop response was derived.

4.2.2 Ellipticity Measurements

With the suspended 10 m cavity locked to the frequency stabilised laser, the cavity throughput was measured with the ellipticity readout scheme described in Section 4.1.2.

The two photodiode signals were recorded in CDS for time stretches of up to two hours. By recording the cavity transmission for long periods, the time series data can then be low-pass filtered to remove laser noise and suspension mode effects from the data. The ellipticity of the transmitted beam is not a frequency dependent variable so filtering the data does not affect the polarisation analysis.

Even with the active frequency stabilisation the laser noise is still significant above 1 kHz and the laser noise is imposed on the transmitted and reflected signals of the 10 m cavity. In addition, the suspension modes of the cavity affect the recorded data at and above 0.7 Hz, the fundamental pendulum frequency of the cavity suspensions.

The result of filtering the time series data can be seen in Figure 4.13 where the raw time series data is filtered with an eighth order low-pass filter with a corner at 0.1 Hz. This removes the suspension modes and laser noise from the data.

It was noticed early on that the transmitted DC power levels varied with the environmental temperature of the laboratory, which had a typical thermal cycle of 40 minutes. It was decided to recorded data over longer time-scales to be able to explore and explain temperature driven effects in the recorded data.

From the raw time series data of the cavity transmission, the polarisation of the light in the cavity is considered by calculating the ellipticity (Ψ) defined as the normalised difference in the power of the two beams, given by the expression

$$\Psi = \frac{P_p - P_s}{P_p + P_s}. \tag{4.2}$$

The sum, difference and ellipticity calculated from a stretch of recorded data is shown in Figure 4.14 as an example of the data output.

In order to isolate the effect of the cavity on the polarisation state, a similar measurement was done without the ETM in the beam path. In this way, the beam transmitted though the ITM was measured with the same ellipticity readout scheme as the cavity data, and the measured ellipticity of the beam with no cavity was used in combination with the data recorded with the cavity so that the final result shows only the ellipticity induced by the resonant cavity and excludes imperfections in the QWP, the relative orientation of the suspended PBS and the readout PBS, substrate birefringence in the transmissive optics and birefringence

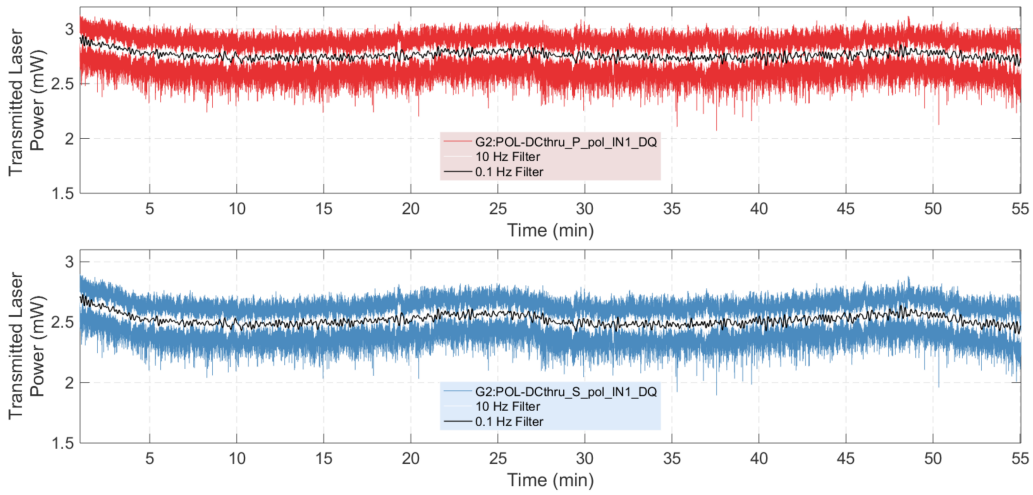


Figure 4.13: Transmitted DC power from the locked 10 m cavity, showing raw data (colour traces) and the data filtered with a 10 Hz (white) and 0.1 Hz (black) eighth order low-pass filter. The output of both photodiode channels of the ellipticity readout scheme are shown for a 55 minute time series, with the p-polarised beam power on the top subplot (red raw data) and the s-polarised beam on the lower plot (blue raw data).

of the steering mirrors.

In addition to the ellipticity readout, auxiliary channels were recorded to fully characterise the experiment. The reflected DC signal from the cavity and the error point and feedback signals were recorded to monitor the state and stability of the cavity locking, as were the reference cavity reflected DC and error signal for the laser stabilisation.

The laser input was monitored with two photodiodes to monitor both the laser power and polarisation state of the light seeding the cavity. The laser light is injected into the vacuum system via a fibre feed-through. On the air side of the system a small fraction of the light injected into the fibre is picked off at a fibre tap and measured with a fibre coupled photodiode. This output is recorded and shows the optical power injected into the system to illuminate the 10 m cavity.

A second monitor photodiode was added to record the power of the laser beam reflected from the suspended PBS before the cavity. This records the power in the orthogonal polarisation to the beam injected into the cavity before the QWP. This channel was also used to monitor the change in the polarisation of the seed light.

From this recorded auxiliary data it was noticed that the optical power and the polarisation of the light delivered to the in-vacuum experiment varied with the temperature of the lab; this effect is further investigated in Section 5.4.

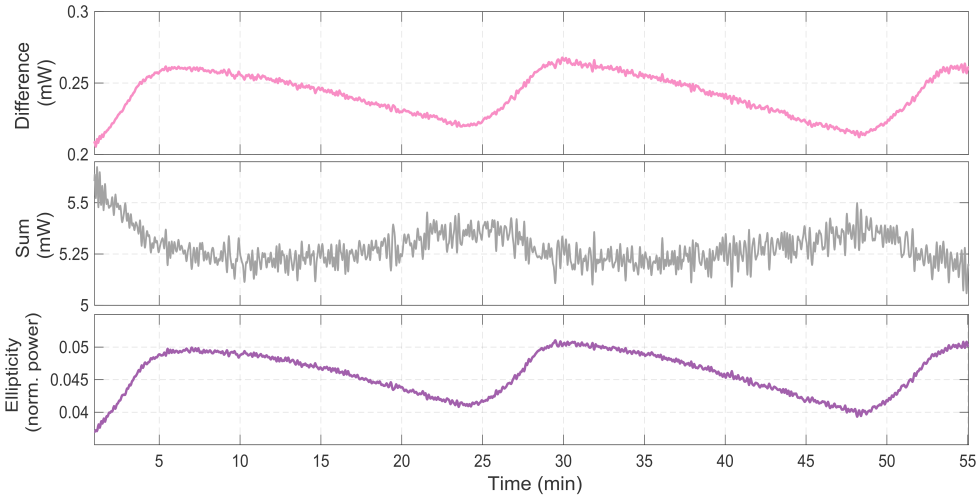


Figure 4.14: Sum (grey) and difference (pink) in DC power transmitted from the locked 10 m cavity in orthogonal polarisations measured by the two photodiode channels, and the inferred ellipticity of the light within the cavity (purple).

4.2.3 Rotating the ETM

In order to measure possible cavity birefringence effects, ellipticity measurements were taken for six different orientations of the ETM relative to the ITM. If each mirror has some intrinsic level of birefringence, then the combined birefringence of the cavity will depend on the relative alignment of the birefringent fast axes of the two mirrors.

The ETM mirror was rotated within the metal optic holder. As a reference the mirror was carefully marked with two dots of epoxy on the edge of the rear optical surface and the optic was photographed in each orientation. This is shown in Figure 4.15.

The process of rotating the ETM mirror required the vacuum system to be opened in order to access the suspended optic. After each change, the system was pumped down to vacuum and the cavity locked and data recorded. This lengthy and repetitive process restricted the number of measurements that could be completed in a reasonable time-frame.

4.3 Results of the Single Arm Polarisation Sagnac Speedmeter Experiment

As described in Section 4.2.2, the polarisation of the light injected into the cavity was measured by removing the ETM from the beam path and recording the throughput of the ITM mirror with the ellipticity readout scheme with no cavity.

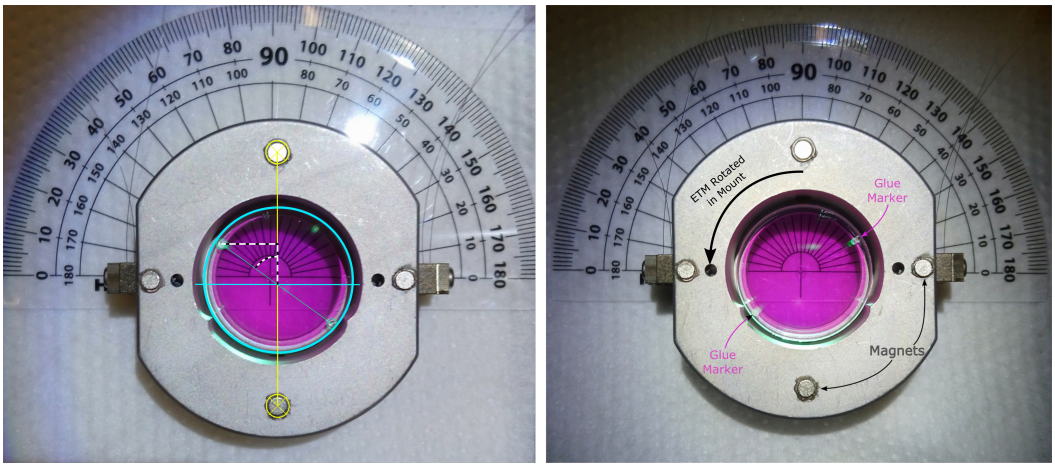


Figure 4.15: Rotating the ETM. The back surface of the optic is marked in two places with small dots of epoxy. The magnets on the metal mass are used as a fixed reference to work out the relative angle of the optic post-adjustment.

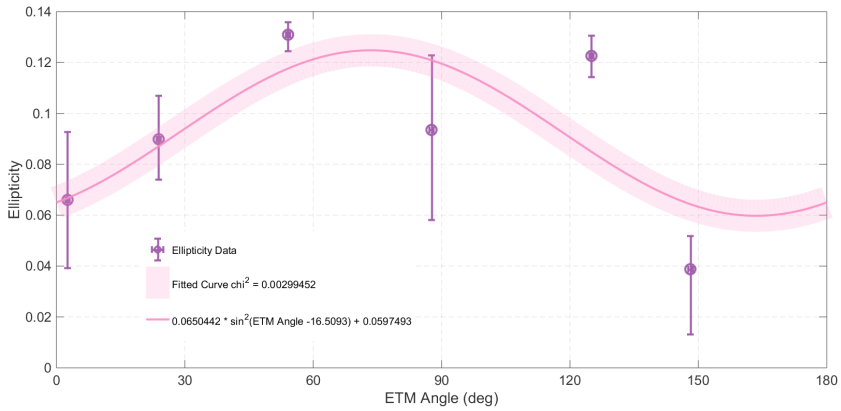


Figure 4.16: Measured ellipticity induced by the resonant cavity for the six orientations of the ETM mirror.

This was then used to correct the data recorded with the resonant cavity in place to isolate the birefringent effect of the resonant cavity.

The ellipticity as a function of ETM orientation is shown in Figure 4.16, this shows a sinusoidal function fitted to the recorded data. What can be said about the raw data is that the resonant cavity clearly has an effect on the polarisation state of the light within the interferometer. A full analysis of the data is presented in Chapter 5.

4.4 Summary

This chapter described the linear cavity experiment set up to investigate polarisation limitations in a polarisation Sagnac speedmeter scheme by utilising the

polarisation optics characterised in Chapter 3. The experiment aimed to measure birefringence effects in a single linear cavity.

The results of this experiment showed a relationship between the level of birefringence evident in the linear cavity and the relative orientation of the cavity mirrors. This is consistent with a polarisation effect due to low levels of birefringence in each mirror forming the cavity. The results of this experiment are not clearly understood on their own and further analysis of these results is presented in Chapter 5.

Chapter 5

Results of the Linear Cavity Experiment

In Chapter 4 an experiment to measure birefringence effects in a linear optical cavity was described, in this chapter the results of this experiment are analysed and compared with a model of the birefringent cavity built using Jones calculus.

5.1 Jones Model of Linear Arm Cavity

Jones calculus [48] is a mathematical framework that describes the interaction of polarised light with birefringent and polarising optics (see Section 1.4.2 and Appendix C).

A linear optical cavity on resonance with a seed light field can be considered as a composite mirror where the reflectivity and transmissivity, ρ_n and τ_n ($n = 1,2$) of the two mirrors forming the cavity combine such that the amplitude of light transmitted (E_T) and reflected (E_R) from the cavity can be expressed as

$$E_T = \frac{\tau_1 \tau_2 E_I e^{-i\delta}}{1 - \rho_1 \rho_2 e^{-i2\delta}} \tag{5.1}$$

and

$$E_R = \rho_1 E_I - \frac{\rho_2 \tau_1^2 E_I e^{-i2\delta}}{1 - \rho_1 \rho_2 e^{-i2\delta}}, \tag{5.2}$$

where δ is the cavity detuning and E_I is the input field amplitude.

To model the liner cavity experiment it is necessary to consider each mirror as a waveplate of arbitrary birefringence η_n with fast axis orientated at an angle θ_n relative to the horizontal axis.

The reflectivity and transmissivities can be replaced with Jones matrices, $\mathbf{R}_n$

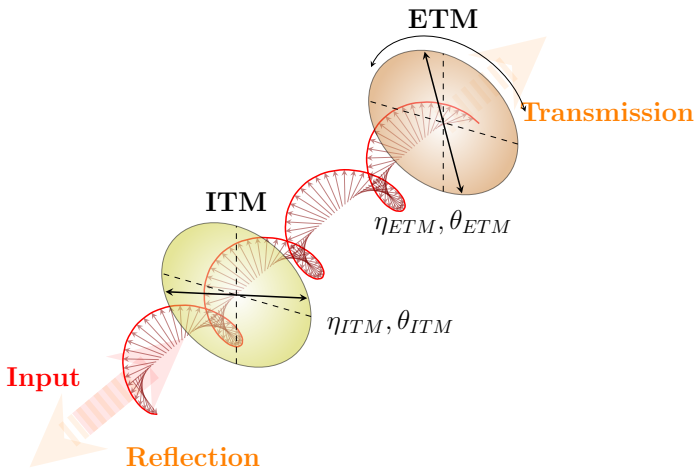


Figure 5.1: Linear cavity formed by two mirrors the input test-mass (ITM) and the end test-mass (ETM). Each mirror has a birefringence η with fast axis orientated at an angle θ.

and $\mathbf{T}_n$,

$$\mathbf{R}_n = \rho_n \, \mathbf{J}_n \tag{5.3}$$

and

$$\mathbf{T}_n = \tau_n \mathbf{J}_n, \tag{5.4}$$

where $\mathbf{J}_n$ is the Jones matrix of each birefringent element, which can be expressed as

$$\begin{aligned} \mathbf{J}_n =& \mathbf{Rot}(-\theta_n) \cdot \mathbf{J'}_n \cdot \mathbf{Rot}(\theta_n) && (5.5) \\ =& \begin{bmatrix} \cos\theta_n & -\sin\theta_n \\ \sin\theta_n & \sin\theta_n \end{bmatrix} \cdot \begin{bmatrix} e^{i\frac{\eta}{2}} & 0 \\ 0 & e^{-i\frac{\eta}{2}} \end{bmatrix} \cdot \begin{bmatrix} \cos\theta_n & \sin\theta_n \\ -\sin\theta_n & \sin\theta_n \end{bmatrix}, && (5.6) \end{aligned}$$

where $\mathbf{Rot}(\theta)$ is the rotation matrix. This can be given as a single matrix,

$$\mathbf{J}_n = \begin{bmatrix} e^{i\frac{\eta_n}{2}} \cos^2\theta_n + e^{-i\frac{\eta_n}{2}} \sin^2\theta_n & (e^{i\frac{\eta_n}{2}} - e^{-i\frac{\eta_n}{2}}) \cos\theta_n \sin\theta_n \\ (e^{i\frac{\eta_n}{2}} - e^{-i\frac{\eta_n}{2}}) \cos\theta_n \sin\theta_n & e^{i\frac{\eta_n}{2}} \cos^2\theta_n + e^{-i\frac{\eta_n}{2}} \sin^2\theta_n \end{bmatrix}, \tag{5.7}$$

which can be used in the modified expression of the cavity reflection and transmission field amplitudes

$$\vec{E_T} = \frac{\mathbf{T}_1 \mathbf{T}_2 \vec{E_I} e^{-i\delta}}{\mathbf{I} - \mathbf{R}_1 \mathbf{R}_2 \, e^{-i2\delta}} \tag{5.8}$$

and

$$\vec{E_R} = \mathbf{R}_1 \vec{E_I} - \frac{\mathbf{R}_2 \mathbf{T}_1^2 \, \vec{E_I} \, e^{-i2\delta}}{\mathbf{I} - \mathbf{R}_1 \mathbf{R}_2 \, e^{-i2\delta}}. \tag{5.9}$$

Using this formalism, a linear resonant cavity can be modelled as a birefringent element, and its effect in a larger polarisation-dependent setup can be explored in further modelling.

5.2 Fitting a Jones Model to the Results of the Linear Cavity Experiment

A complete model of the linear cavity experiment was constructed using Jones matrices to describe each polarising element, including the suspended PBS, the suspended QWP, the two cavity mirrors and the readout PBS after the ETM. Each path though the experiment from input to readout can then be modelled as the multiplication of each optic's matrix with the vector describing the input field to calculate the output field amplitude. More detail of this model can be found in Appendix C.

5.2.1 Inferring Input Beam Parameters from Auxiliary Monitors

In order the model the linear cavity experiment, the input beam power (P_{in}) and polarisation angle (θ_{pol}) of the input light need to be known. The input beam of linear polarisation is described by the Jones vector

$$\mathbf{j_{in}} = \sqrt{P_{in}} \begin{bmatrix} \cos\theta_{pol} \\ \sin\theta_{pol} \end{bmatrix}. \tag{5.10}$$

Both the power of the input beam and the polarisation angle change over the course of the measurement. The fibre injection system (described in Section 11.1.3) is susceptible to changes in the environmental temperature coupling to changes in the polarisation of the light delivered into the vacuum system. The power output of the laser system also changes with environment temperature (see Section 11.1.4).

The input beam parameters can be reconstructed from the recorded photodiode signals. The input power can be measured as the sum of power reflected and transmitted by the suspended PBS. And the polarisation angle can be measured from the ratio of the reflection and transmission from the PBS.

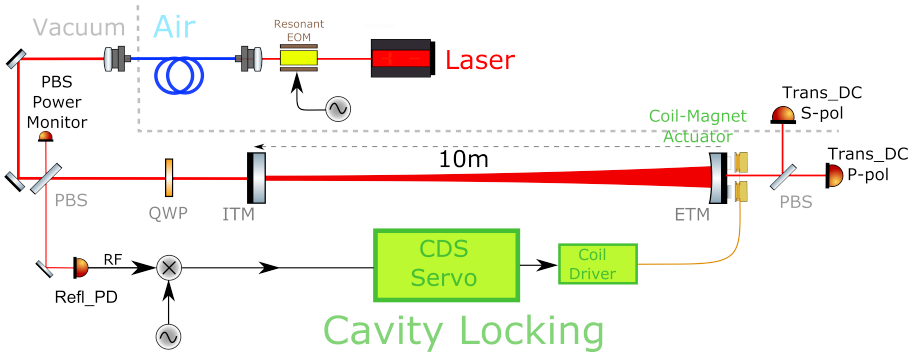


Figure 5.2: Experimental setup of the linear cavity experiment showing the four photodiode readout points. The laser beam is injected into the vacuum system via a fibre feedthrough; the beam is aligned to the 10 m cavity with two suspended mirrors, and the beam reflected from the PBS is detected with a photodiode. The beam returning from the 10 m cavity is reflected from the reverse face of the PBS and detected on a PD with DC and AC coupled outputs. The AC signal is demodulated with a local oscillator signal and filtered by the CDS digital servo to produce a locking signal. The feedback signal is applied to coil-magnet actuators that control the length of the cavity. The throughput of the cavity is measured with two photodiodes that detect light in orthogonal polarisations.

5.2.2 Results with No Resonant Cavity

To test the Jones model of the linear cavity experiment and to measure the ellipticity of the beam after the QWP, data was recorded without a resonant cavity. The ETM was removed from the cavity beam path so the throughput of the ITM was recorded with the ellipticity readout scheme.

Results of the measurements with no cavity are shown in Figure 5.3 with the Jones model of this experiment fitted to the data. The Jones model was able to match the recorded data by modelling the input beam with a time varying power and polarisation angle.

The measured ellipticity with no cavity is shown in Figure 5.3 (d). This was used to determine the parameters of the suspended QWP to be used in the model with the linear resonant cavity.

The measure of the ellipticity of the beam with no cavity showed negligible variations over the course of two hours.

5.2.3 Fitting Model to Cavity Data

The Jones model of the linear cavity was fitted to the measured data for each orientation of the ETM relative to the ITM. In each case the time variation of the input beam parameters was considered and included in the model.

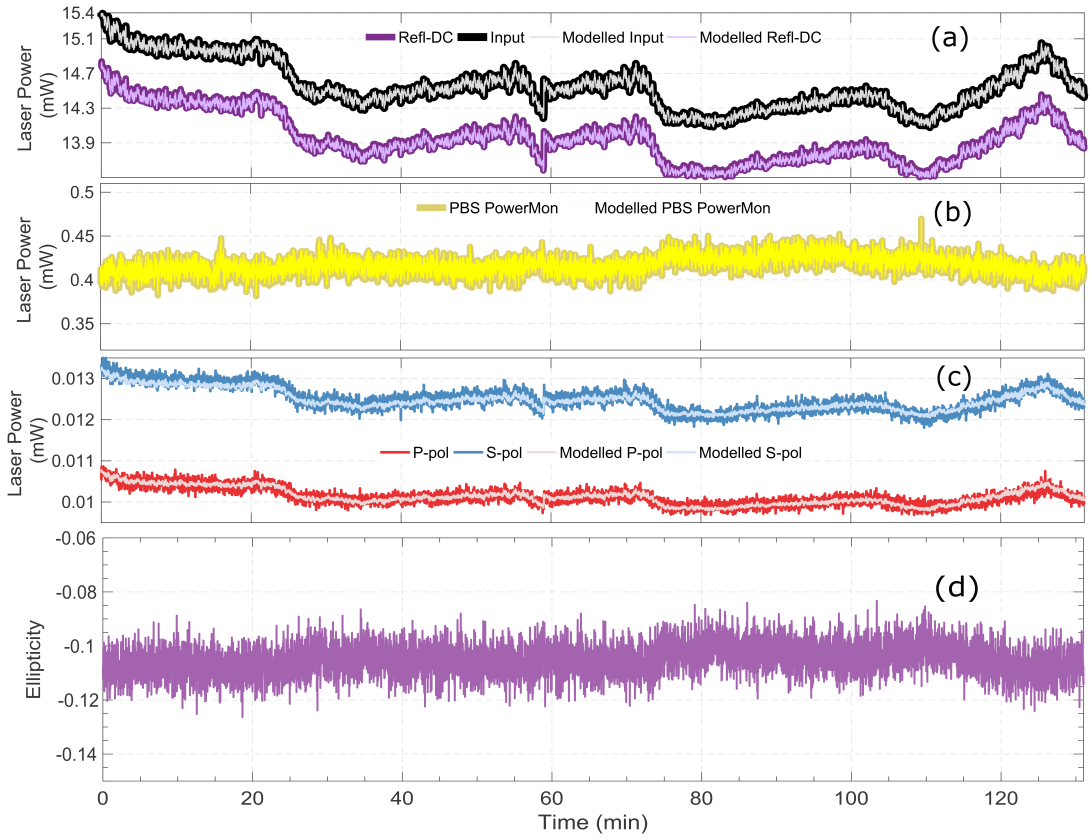


Figure 5.3: Data recorded with no cavity. Plots show the fitting of the Jones model to the experimental data time series for the total input power and the power in reflection from the arm cavity (a), the s-polarised light reflected from the PBS at the input to the cavity (b), and the transmitted power in s- and p-polarisation. Also shown is the measured ellipticity with no cavity (d).

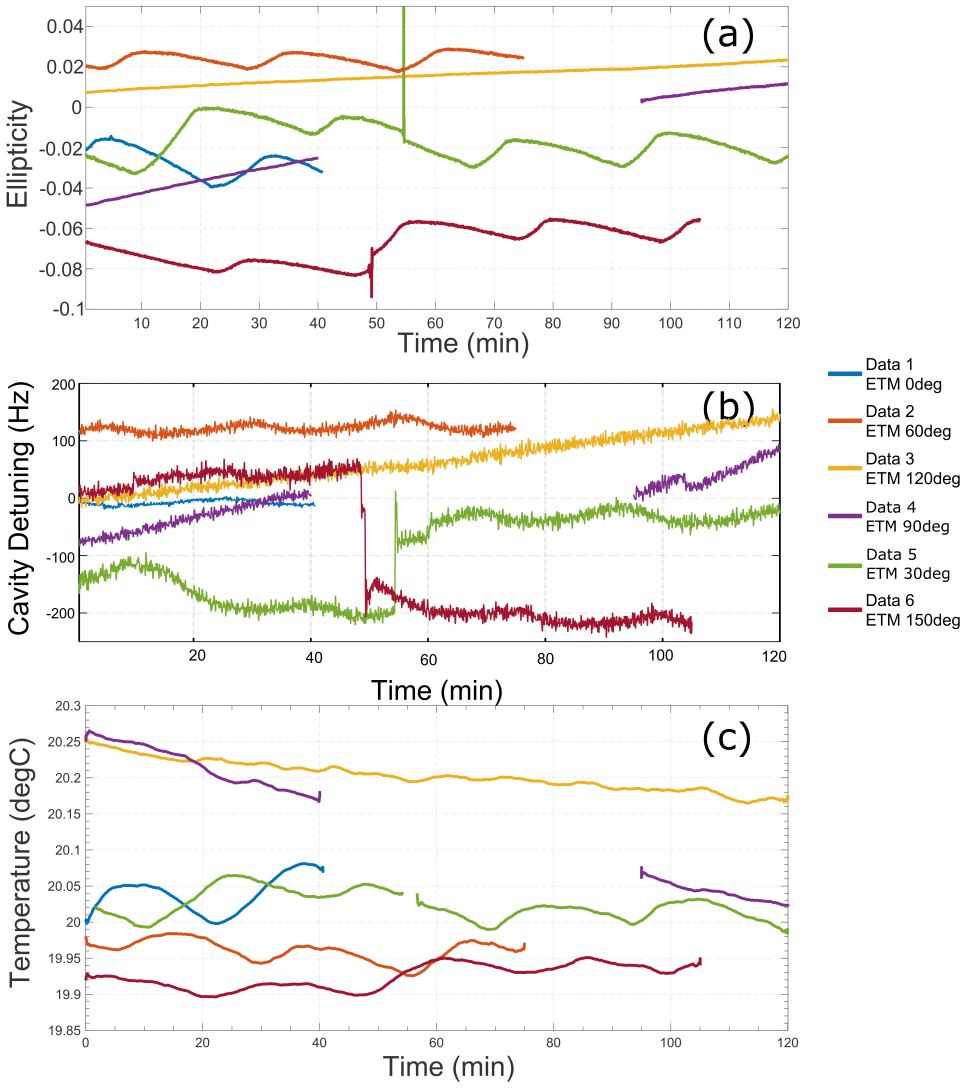


Figure 5.4: Plots showing time variation of measured ellipticity (a), cavity detuning (b), and environmental temperature (c) for all six data sets.

The model was only partially successful in matching the experimental data. It was found that along with the variations in input power and polarisation there was additional time dependence in the measured ellipticity.

Cavity Detuning

As well as the temperature dependent drifts in the input beam parameters, it can be observed from the recorded data that the cavity locking error point drifted during the course of each measurement.

The error point of the cavity locking loop is recorded as a voltage in CDS. This can be converted to give a measure of frequency using the known calibration

of the error point. By filtering the data above 0.1 Hz the offset in the locking error signal as a measure of the cavity detuning can be found.

The variation in cavity detuning for the six data sets is shown in Figure 5.4 (b).

The cavity detuning affects the amount of power reflected and transmitted from the cavity and in the case of a cavity formed by birefringent mirrors, the detuning also affects the ellipticity induced by the cavity.

If the cavity is birefringent, the resonant frequency of the cavity for orthogonal polarisations will be different. For small levels of birefringence, the two cavity resonances will overlap so light of both polarisations will be resonant simultaneously. However, the rate of change of transmitted and reflected power in each polarisation will vary differently with cavity detuning.

The cavity detuning effects were included in the Jones model of the linear cavity by making the cavity detuning δ (see Equation 5.9) a time dependent variable derived from the filtered error point signal. Including this added complexity in the model improved the matching of the model to the experimental data but did not fully explain the extra variation in the measurements.

The inferred cavity detuning over the course of the measurement made with the ETM at 30° is shown in Figure 5.5. During the course of this measurement the cavity fell out of lock. Before falling out of lock the error point had drifted such that the cavity detuning in frequency was around 200 Hz. When the cavity was locked again, the cavity detuning as closer to 50 Hz and drifted between 10-50 Hz during the following hour.

The environmental temperature of the laboratory over the period of this measurement is also shown in Figure 5.5. A comparison of the two plots shows that the drift in the cavity locking point is related to temperature of the laboratory.

Figure 5.6 shows the result of fitting the Jones model of the linear cavity to the experimental results for the corresponding time interval to the detuning and temperature data shown in Figure 5.5. In this plot it can be seen that a better fit to the experimental data is possible in the second lock segment when the cavity detuning was smaller.

5.3 Inferring Mirror Parameters

From the work to model the linear cavity it was found that at certain times the Jones model could match the experimental data well and this was typically when the cavity detuning was small. Using this information a qualitative down-selection of the ellipticity measurement data was performed based on corresponding analysis

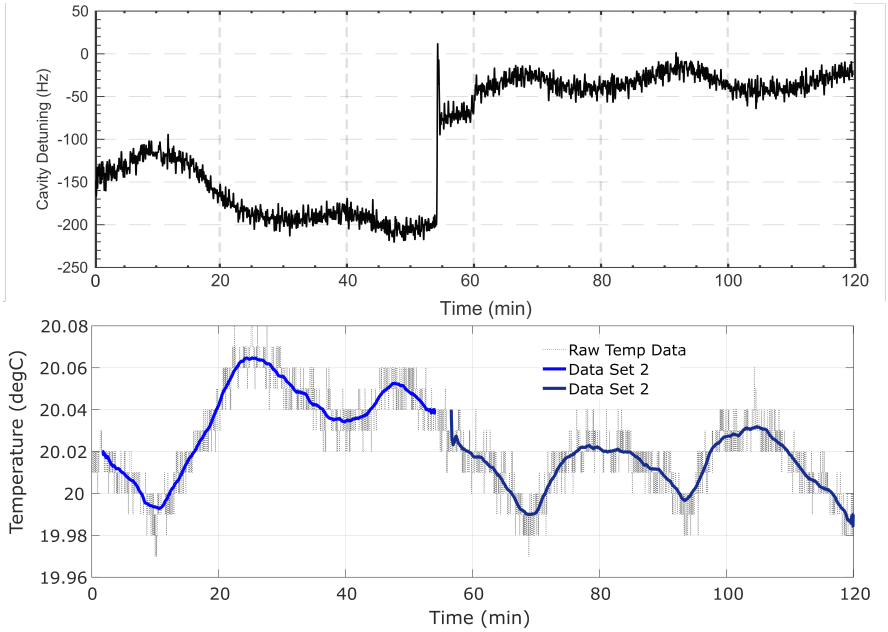


Figure 5.5: Cavity locking error point for data recorded with the ETM at 30° in terms of cavity detuning in hertz (top plot). The cavity fell out of lock after 55 minutes, after the cavity was relocked the cavity detuning was closer to zero. Also shown is the environmental temperature during the measurement (bottom plot).

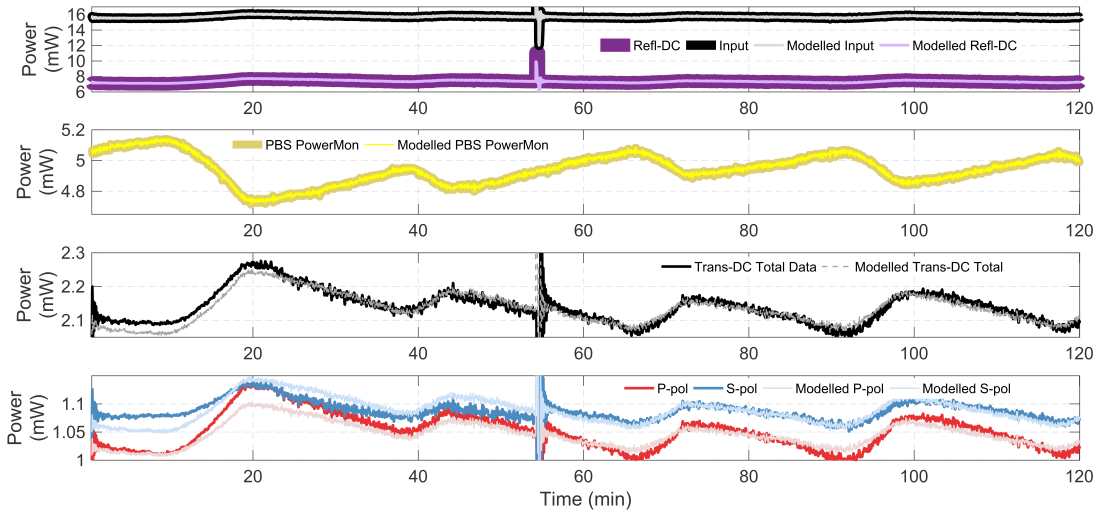


Figure 5.6: Jones model of linear cavity fitted to experimental data recorded with the ETM orientated at an angle of 30° relative to the starting point. In this case the cavity fell out of lock after 55 minutes. After a brief pause the cavity was relocked.

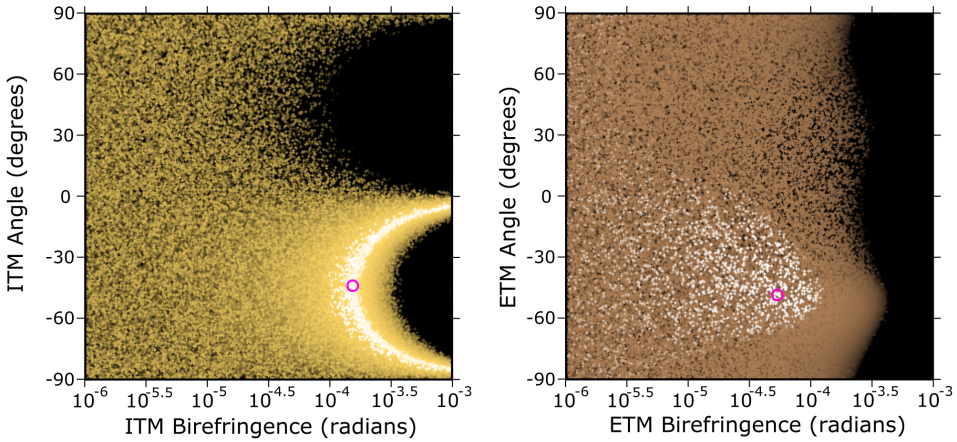


Figure 5.7: Parameter estimation for birefringent mirrors in a linear cavity Jones model. Colour map represents quality of fit to experiment data with white representing the highest quality factor of the least squares fitting and black representing the lowest quality factor normalised to the best fit highlighted with the pink marker.

of the cavity detuning discussed in the previous section to use data corresponding to times when the cavity detuning was small and the data could be reproduced with a model of the system.

With this reduced data set, a linear cavity model was used to fit to the measured ellipticity as a function of ETM angle by varying the birefringence parameters of the two cavity mirrors.

A Monte-Carlo style optimisation was done to find the best fit of the Jones model to the experimental data by choosing random combinations of ETM and ITM birefringence η and fast axis orientation θ.

For each of these simulations a weighted difference of least squares measurement, or χ^2 measurement, was taken as a quality factor of the fit.

Figure 5.7 shows the result of this parameter estimation simulation. The two plots show the parameter space of the simulation with each point being a simulation performed and the colour of each point representing the χ^2 quality factor of the fit.

By performing such a simulation, a reduced parameter space of best fits can be found that corresponds to an estimation of the birefringence parameters of the two cavity mirrors based on the experimental data.

Figure 5.8 shows the same result in a two tone colour map highlighting the region of the parameter spaces with the best fitting of the model to the data.

For the ITM there is a clear degeneracy between the level of birefringence and fast axis orientation. Between angles of $-3°$ and $-85°$ the birefringence of the

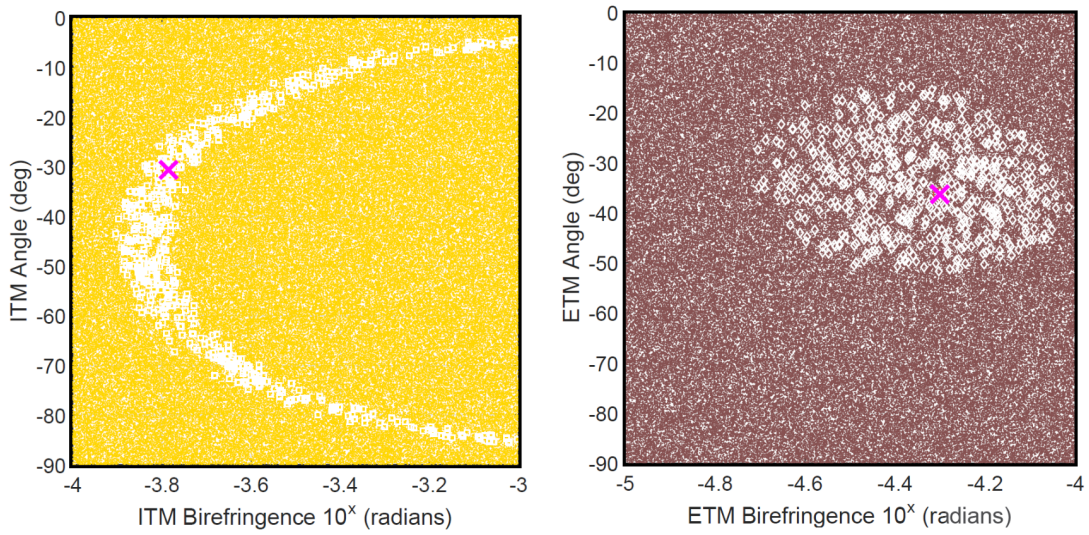


Figure 5.8: Parameter estimation for birefringent mirrors in a linear cavity Jones model. Small data points show all simulation points, the pink crosses show the simulations with a quality of fit better than threshold value to show parameter space of best fit.

	Birefringence (rad)	Fast Axis Orientation (degree)
ITM	1.99×10^{-4}	$-45°$
ETM	3.98×10^{-5}	$-30°$

Table 5.1: Birefringent parameter of the cavity mirrors based on the results of the linear cavity experiment.

mirror is predicted to be between 10^{-3} radians and $10^{-3.9}$ radians.

Because the ETM mirror was rotated, this changed the fast axis of the birefringence, breaking the degeneracy. In the case of the ETM the simulation predicts that the fast axis angle is between $-14.9\,°$ and $-51.3\,°$ with the birefringence between 10^{-4} radians and $10^{-4.7}$.

The best fit results are summarised in Table 5.1. These parameters were used for further modelling, including the plot of ellipticity as a function of ETM rotation shown in Figure 5.9 and as a function of both ETM and ITM angle in Figure 5.10.

Based on the results obtained, Figure 5.10 shows the modelled ellipticity over the full parameter space of cavity mirror orientations. This result shows that in principle there are numerous combinations of relative mirror orientations where the birefringence effect of the two mirrors can cancel out, so as to maintain the polarisation state of the light resonant in the cavity.

This is important for future work on the polarisation Sagnac speedmeter interferometer. If such an interferometer is set up without consideration or knowledge

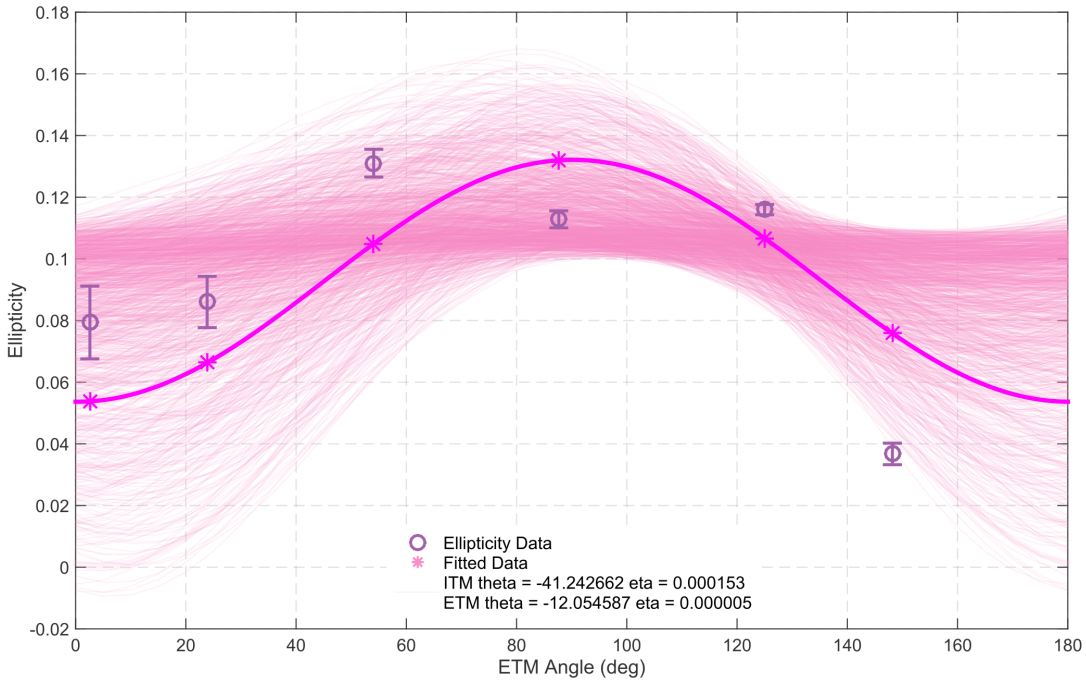


Figure 5.9: Ellipticity data as a function of ETM angle with the fitted result of the Jones model parameter estimation. The pink traces show the range of fits based on the parameter estimation shown in Figure 5.8 with the thick pink trace representing the best fit parameters.

of the birefringence of the cavity mirrors, a significant ellipticity can be introduced to the resonant light. However it is possible to orientate the mirrors to reduce the cavity birefringence effect.

5.4 Consideration of other Birefringence Effects

The ellipticity measured with no resonant cavity did not have the same variations as the results of the linear cavity experiment. It is clear that the presence of the resonant cavity effects the polarisation of the light in the experiment.

Furthermore, the measured level of ellipticity changed as the ETM was rotated. This is consistent with a mirror coating birefringence effect amplified by the resonant cavity. And while other parameters were also different between measurements, such as the cavity detuning and the environmental temperature, there was no clear correlation with the measured ellipticity.

Environmental Temperature

The temperature and humidity of the laboratory are measured with a table-top monitor and recorded in CDS.

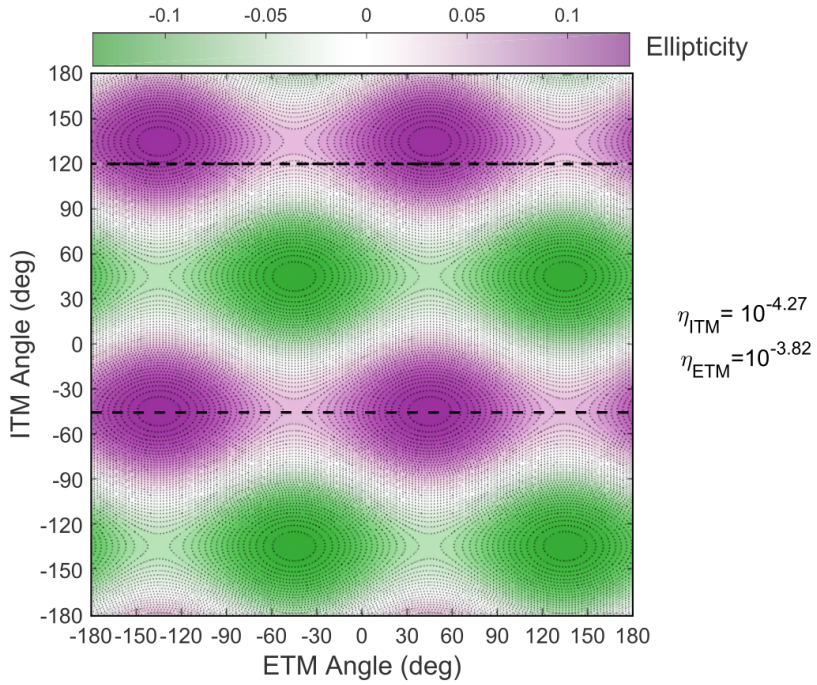


Figure 5.10: Ellipticity as a function of ETM and ITM rotation based on the linear cavity results. Black dotted line is at an ITM rotation of $-45°$ as the mirror was for the linear cavity model.

An additional in-vacuum temperature monitor was installed in the vacuum tank containing the ETM suspension. This showed that the in-vacuum temperature followed the air temperature of the laboratory over time-scales of several hours but not over time-scales of tens of minutes. This means the regular changes in the lab temperature with a typical period of 40 minutes do not translate to changes in the in-vacuum temperature.

This fact helps to rule out several factors that could only be related to in-vacuum temperature as the variations in the data from the linear cavity experiment follow the cycle of the laboratory air temperature.

PBS Orientation

The relative alignment of the readout and suspended PBS's could have a significant effect on the ellipticity measured with this scheme. The two PBS's are mounted on separate breadboards in different tanks. Over time, with changes in temperature the whole system can shift causing changes in the alignment.

However, when the ellipticity was measured with no resonant cavity the two PBS's were in place, so if this was affecting the ellipticity measurement it would have presented as a measurable effect during this test.

Stress on the ETM Mirror

As described in Section 4.2.3, each time the ETM was rotated to a new orientation the grub screw that holds the mirror in the suspended mass was loosened so that the mirror could be moved, then re-tightened with the mirror in a new orientation. Due to the difficulty of repeating this process, for each measurement set a different level of stress would have been exerted on the ETM mirror.

If the dominant birefringent effect in the linear cavity experiment was due to mechanical stress on the mirrors due to the mounting of the optic in the suspended mass, then this difference between measurements would result in a difference in measured ellipticity.

The change in stress on the ETM mirror presented in a noticeable way when trying to lock the cavity. As described in Section 4.2.1, a resonance above 4.5 kHz was present in the loop response of the suspended cavity locking system related to the mechanical mode of the cavity suspension. It was noticed that the exact frequency of this resonance changed with every change to the ETM mirror, and this was believed to be due to the variation in the force applied to the mirror.

The force exerted on the mirror is not easily measured, but if the frequency of the resonance is assumed to be proportional to square root of the force, as it would be if the system is considered to be a mass on a spring, then the force exerted on the mirror will alter the apparent stiffness of the spring. This was checked but no correlation between the measured ellipticity and the ETM mechanical mode frequency is evident.

Alignment of the Cavity

As mentioned previously in this section, the alignment of the 10 m system can change with temperature.

The alignment of the beam to the cavity will affect the amount of power coupled into the cavity if the input beam is misaligned relative to the cavity axis. A change in power in the cavity would affect the measured throughput power but not the relative power in the orthogonal polarisation; this is not consistent with the observed changes in ellipticity.

Changes in the pointing of the beam through the QWP would alter the polarisation of light as the optical path length though the QWP is changed if the angle of incidence on the optic changes. Characterisation of the QWP optic in Section 3.3.2 showed a significant change in the birefringence of the optic with a 10° change in angle of incidence.

However, alignment changes would have to be caused by changes in the tem-

perture or pressure within the vacuum system. In-vacuum temperature changes have already been reported to be unrelated to the measured ellipticity changes. During the measurements the vacuum pressure rose slightly as the system was not being pumped but the pressure remained below 10^{-4} mbar so no alignment changes from pressure change would be expected.

Spot Position on Cavity Mirrors

The change in position on the cavity mirrors could be an alternative explanation for the variation in measured ellipticity between measurements. If the birefringence of a mirror is different across the mirror surface as would be expected from [96] and from the assumption that coating birefringence is generated by thermal stress between dielectric layers of the coating and between the coating and the substrate, then changes in the spot position of the cavity beam on the mirrors would change the induced ellipticity.

The position of the beam on the ETM mirror is monitored with a camera on the outside of the vacuum system that images the rear surface of the suspended optic (see Figure 4.6). When realigning the cavity between measurements the spot position on the ETM was measured to be consistent between measurements within a tolerance of ± 1 mm.

It would be challenging to implement an equivalent monitor of the spot position of the beam on the ITM mirror due to the lack of a suitable viewport. No such measurement was attempted as the cavity could be aligned without this knowledge. There is therefore a greater uncertainty in the variation in spot position of the beam on the ITM mirror between measurements.

Further, if the ellipticity is at a minimum at the centre of the coating, where the thermal stresses would have cancelled, then the higher level of birefringence of the ITM mirror measured in the previous section could be a result of the cavity beam being off-centre during measurements. The spot position on the ITM mirror could be as much as 10 mm from the centre of the optic.

5.5 Summary

From the linear cavity experiment and the measurements of cavity birefringence effect, described initially in Chpater 4, a lot can be learned that will influence the technical considerations for future large-scale polarisation Sagnac speedmeter experiments.

If the birefringence of the cavity mirrors can be measured before installation,

then the two mirrors can be aligned to minimise the induced ellipticity. This may not always be possible since coatings for gravitational wave detector mirrors are typically deposited onto a substrate with flat sides that define the angle at which the mirror will be suspended before the coating birefringence can be measured.

Alternatively, if the birefringence can only be measured once the cavity is formed an additional birefringence compensation optic could be implemented, either as an additional waveplate or by changing the birefringence of the quarter waveplate between the arm cavity and the central PBS as part of the polarisation Sagnac scheme.

The experimental system to measure cavity birefringence effects developed here could be further utilised to measure multiple mirrors fabricated with different coating deposition processes in order to investigate the source of coating birefringence and actively move towards reducing coating birefringence.

More study of the variation of coating birefringence across the mirror surface is needed. If the effect is as predicted, then a consideration of coating birefringence needs to be factored into any beam pointing accuracy and cavity alignment tolerance calculations for kilometer base-line interferometers.

The presence of cavity birefringence effects in the linear cavity experiment clearly show that this effect needs to be considered when specifying the technical limitations of future speedmeter based gravitational wave detectors. The potential implication of this effect is investigated further in Chapter 6.

Chapter 6

Modelling a Polarisation Sagnac Speedmeter

6.1 Jones Model of Polarisation Sagnac Speedmeter

The description of imperfect polarising optics using Jones calculus was introduced in Section 3.1 for the polarisation beam-splitter (PBS) and quarter waveplate optics used in the polarisation Sagnac speedmeter interferometer. The matrix for an imperfect PBS is given in Equations 3.3 and 3.4, and the matrix for a QWP in Equation 3.8. In Section 5.1, the Jones matrix for a birefringent mirror was given in the form of the matrix given in Equation 5.7. A complete description of the Jones calculus framework used in this chapter can be found in Appendix C.

6.1.1 Optical Paths in the Polarisation Sagnac Speedmeter Interferometer

The possible on-axis paths of light through the interferometer from input to output are shown in Figure 6.1. The model of the speedmeter consists of a multiplication of the Jones matrices of the optics in each path from input to output and the expressions for the optical field amplitudes in reflection and transmission of a linear cavity contained in Equations 5.8 and 5.9, used in the previous description of the Jones model of a linear cavity in Section 5.1.

The polarisation beam splitter is central to the interferometer, directing the beams between cavities and is critical to ensure each counter propagating field probes both arms of the interferometer thereby sensing the relative velocity of the arm cavity mirrors.

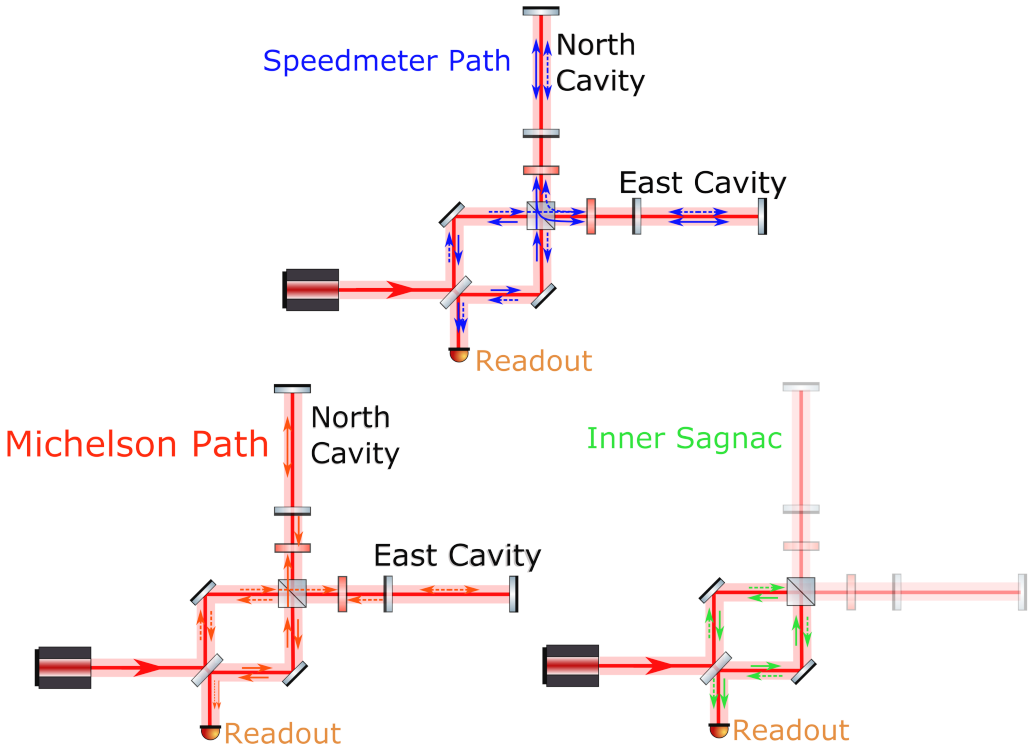


Figure 6.1: Optical paths through the polarisation Sagnac speedmeter from input to readout. The top figure shows the desired speedmeter path where the clockwise (solid) and anti-clockwise (dashed) fields probe both the north and east arm cavities in reverse order. In the bottom left figure, Michelson signals couple to the readout by probing only one cavity per field, and in the bottom right figure, fields that do not enter either cavity but couple directly back to the readout are shown.

The finite extinction ratio of the PBS will mean that for each interaction with the optic some fraction of light is coupled to the incorrect output path. Light that has probed one arm cavity will be coupled back to the main beam-splitter of the interferometer, resulting in the detection of 'Michelson signals' in the speedmeter readout, this is referred to as 'Michelson leakage' or simply 'leakage' in the remainder of this chapter.

In addition, a fraction of the light returning from each arm cavity will be transmitted by the PBS if the polarisation of the returning light is not orthogonal to the input light. This can arise from a combination of the imperfect birefringence of the QWP and unwanted birefringence in the cavity mirrors, adding to the Michelson leakage related to the PBS extinction ratio.

As well as light from one cavity leaking though the PBS due to the finite extinction ratio, a fraction of the input light will be reflected by the PBS, and never enter either arm cavity. This light will propagate back towards the central beam-splitter from both carrier fields forming an inner Sagnac and resulting in 'Sagnac signals' at the readout that are not related to the long baseline arms of the interferometer. This inner Sagnac is an additional degree of freedom of the interferometer that will need to be controlled.

A third path is possible where light travels through three cavities before returning to the signal readout. This can only occur if a fraction of the speedmeter path is reflected back into the first cavity sensed, instead of being transmitted towards the input. This three cavity interaction is not considered to be a significant effect as the amount of power propagating in such a path will be orders of magnitude lower than in other paths.

The performance of the speedmeter interferometer is considered as the ratio of optical power coupled to the readout that could contain Michelson signals versus the power that could contain speedmeter signals. Interference at the readout beam-splitter is not considered in this analysis.

6.1.2 Limits on the Speedmeter Performance Relating to Individual Optics

From modelling the Sagnac speedmeter with Jones calculus, some limits on the performance of the interferometer due to imperfection in the polarising optics can be derived. The derivation of the results presented below can be found in Appendix C.

The finite polarisation extinction ratio in transmission of the PBS (PER_T) will contribute an amount of Michelson leakage power to the interferometer output

equal to

$$P_{\mathrm{PBS}} = \frac{\mathrm{PER}_T - 1}{\mathrm{PER}_T^2}. \tag{6.1}$$

Similarly, the limitation to the speedmeter performance due to an imperfect QWP of birefringence (η) can be predicted from the expression,

$$P_{\mathrm{QWP}} = \cos^2 \eta, \tag{6.2}$$

where the QWP's in each arm have an equal level of birefringence.

6.2 Polarisation Sagnac Speedmeter with Realistic Optics

In this section the results from Chapters 3 and 5 characterising the polarising optics are used in a model to predict the limitations to a speedmeter interferometer that used these optics, or optics of similar parameters.

To quantify the limitations of polarisation optics in a polarisation Sagnac speedmeter the Jones calculus model is used to calculate the optical power at the readout beam-splitter, the ratio of power that could contain Michelson signals is compared to optical power that has followed the desired path through the interferometer and could contain speedmeter signals.

6.2.1 Effect of Polarisation Beam Splitter Extinction Ratio on the Sagnac Speedmeter

The polarisation extinction ratio of the PBS was measured as a function of angle of incidence with the result presented in Figure 3.5 as the transmission and reflection extinction ratios.

The transmission extinction ratio T_p/T_s and reflection extinction ratio R_s/R_p are the parameters by which this type of optic is typically characterised. For the purpose of the polarisation Sagnac speedmeter, it is also useful to consider the extinction ratios for each polarisation (these being T_p/R_p and R_s/T_s) as more intuitive parameters by which to consider the performance of the optic in the speedmeter. These extinction ratios as functions of the angle of incidence are shown in Figure 6.2.

The measured extinction ratios were applied in the polarisation speedmeter model with the result shown in Figure 6.3. The inner Sagnac contributes more

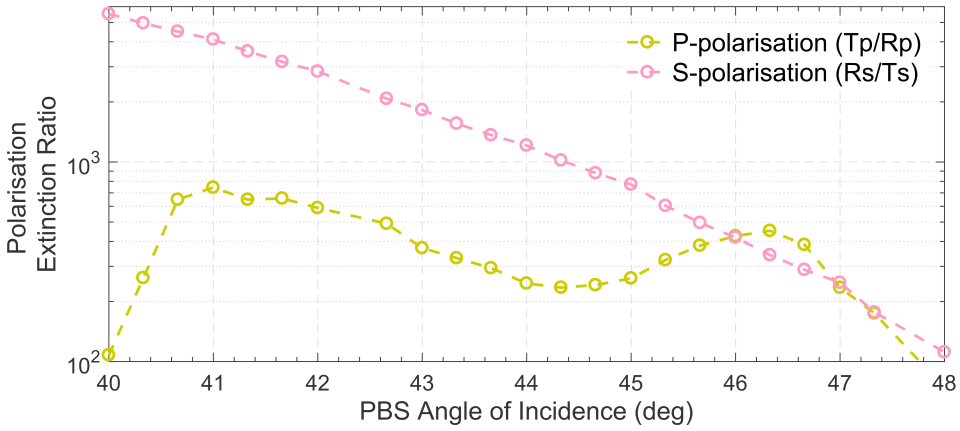


Figure 6.2: Measured polarisation extinction ratio of the PBS in transmission and reflection as a function of the angle of incidence on the PBS optic.

power to the readout than the Michelson leakage, as the relative amount of inner Sagnac signal only depends on the p-polarisation extinction ratio which is lower than the s-polarisation extinction ratio that determines the amount of Michelson signal at the readout.

The result shows that the level of optical power at the readout related to Michelson signals from the interferometer with the PBS at 45° angle of incidence is 0.294%, whereas the inner Sagnac contributes 1.94% of the power to the readout. The Michelson signal rises above 1% for an angle of incidence above 47° when the PBS s-polarisation extinction ratio is less than 250.

The three cavity path only contributes 2.16 ppm to the signal readout with the beam-splitter at 45° angle of incidence.

6.2.2 Effect of Imperfect Quarter Waveplate on the Sagnac Speedmeter

The birefringence of the QWP optic used in the linear cavity experiment (Chapter 4) was measured in Section 3.3.2 to be equivalent to $\frac{\lambda}{4}+\frac{\lambda}{324}$. This measurement was used in a model of a single arm of a polarisation Sagnac speedmeter to understand the implications of the actual birefringence of this optic that is part of each arm cavity of the full speedmeter scheme.

The result is shown in Figure 6.4, where the amount of Michelson leakage signal is modelled as a function of QWP fast axis alignment for the measured and ideal QWP and PBS optics. In this model the cavity was considered to have no birefringence. The PBS extinction ratio measured at the optimal angle of incidence is used in this simulation.

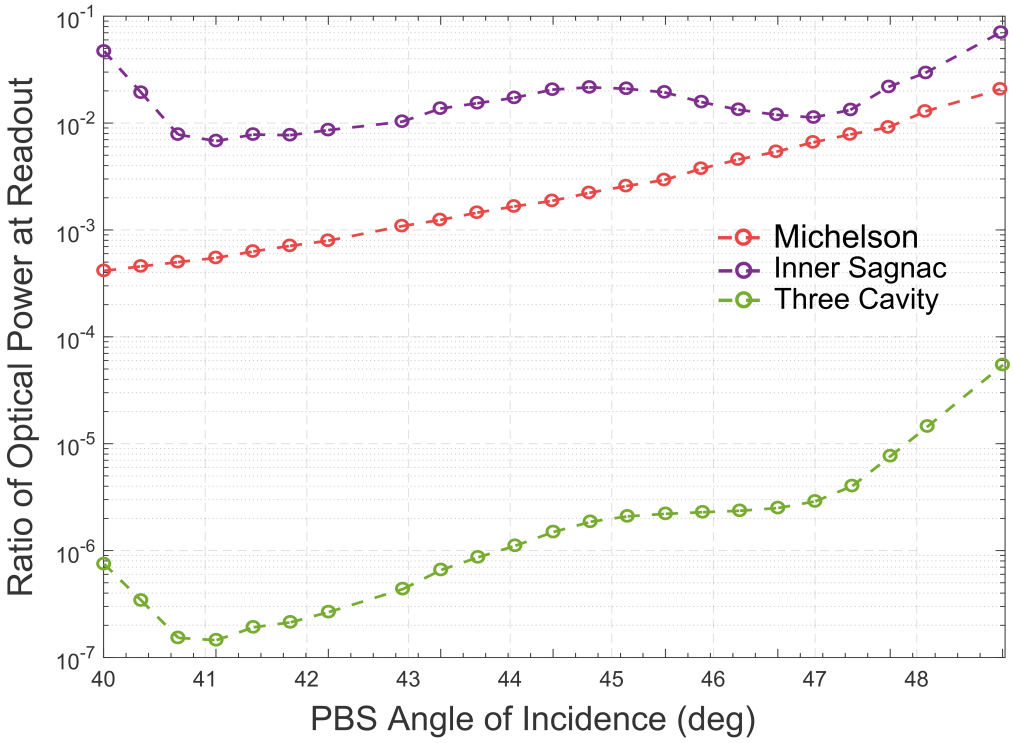


Figure 6.3: Modelled effect of PBS extinction ratio on the signal ratio in a speedmeter interferometer. The ratio of unwanted signal at the readout is shown as a function of PBS angle of incidence for the Michelson, inner Sagnac and three cavity paths through the interferometer.

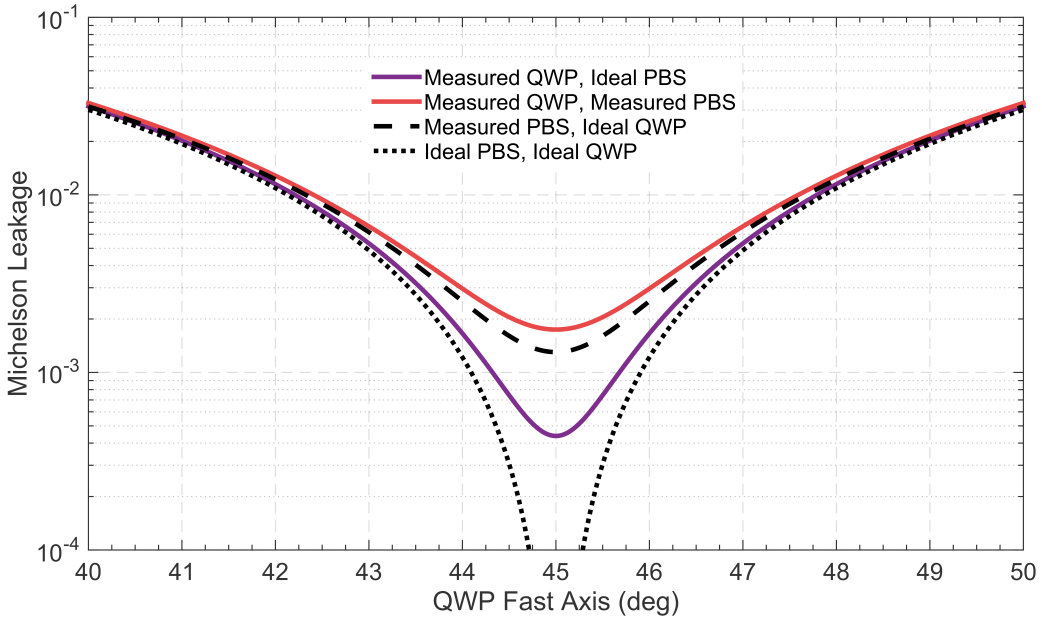


Figure 6.4: Modelled effect of QWP birefringence on the Michelson leakage from a single arm of a polarisation Sagnac interferometer. The ratio of unwanted signal at the readout is shown as a function of QWP fast axis angle. The model uses the measured polarising optics parameters in comparison to ideal polarising optics.

The full speedmeter configuration was also modelled, with the level of speedmeter and Michelson signals at the interferometer readout plotted as a function of QWP alignment in the north and east arm cavities (result in Figure 6.5).

6.2.3 Effect of Cavity Birefringence on the Sagnac Speedmeter

Coating birefringence in the cavity mirrors is expected to limit to the performance of the polarisation Sagnac speedmeter, as birefringence in the arm cavities will result in a change in the polarisation of light returning to the PBS from each cavity. This will result in an amount of light sent to the readout after probing only a single cavity because it is in the polarisation transmitted by the PBS.

In Chapter 5, results of an experiment to measure cavity birefringence are presented. These results put limits on the expected level of birefringence in the cavity optics of a polarisation Sagnac speedmeter. In this section the effect of cavity birefringence in a polarisation Sagnac speedmeter is modelled.

Figure 6.6 shows the result of modelling a single arm of a speedmeter interferometer with birefringent cavity mirrors, in this case ideal PBS and QWP optics were used in the model.

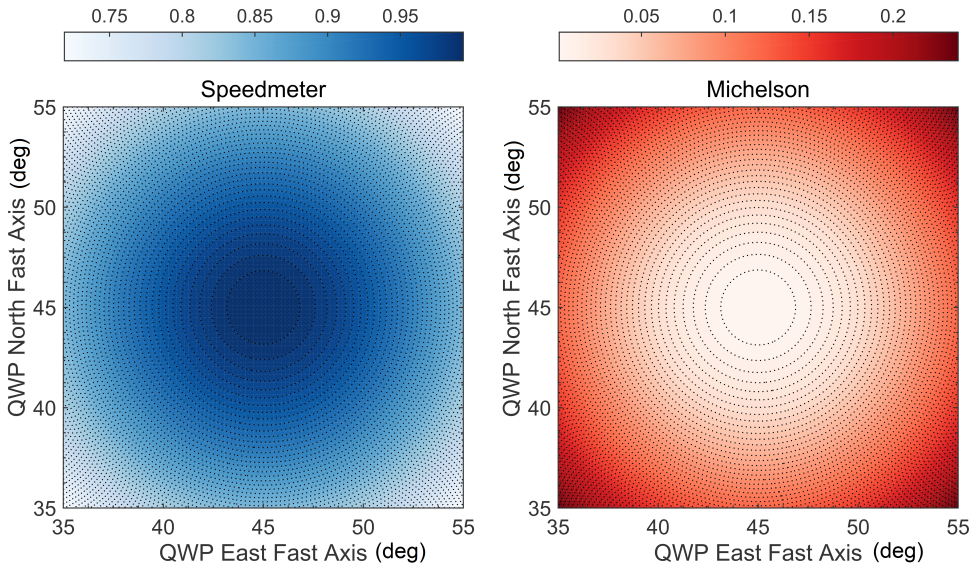


Figure 6.5: Relative levels of speedmeter and Michelson signals at the output of a polarisation Sagnac speedmeter as a function of the rotation of the QWP's in the two arms.

6.2.4 Combined Effect of Imperfect Polarising Optics on the Sagnac Speedmeter

In a real Sagnac speedmeter interferometer, all the polarising elements will, in combination, generate a level of Michelson signal at the interferometer readout. Using a model of the polarisation Sagnac speedmeter the combination of effects can be simulated.

Figure 6.7 shows the result of modelling the amount of Michelson signal leakage from a single arm of a polarisation Sagnac speedmeter as a function of cavity birefringence for difference level of imbalance between ITM and ETM birefringence. In this analysis the measured imperfections of the PBS and QWP optics were included. The individual limitations due to these optics as predicted by equations 6.1 and 6.2 are also plotted. For this analysis the measured PBS extinction ratios at the best angle of incidence were used.

The results from Chapter 5 suggested that the ITM is five time more birefringent than the ETM, but the parameter estimation showed that this imbalance could range from 1 to 50 in the extreme cases. This range of imbalance as modelled is included in the result in Figure 6.7.

This result shows that if the cavity mirror birefringence is less than 10^{-5}, the QWP birefringence will limit the performance of the speedmeter. Above this level the cavity birefringence will be the limiting factor. However, a point exists where the cavity birefringence can be compensated by that of the QWP. An initially

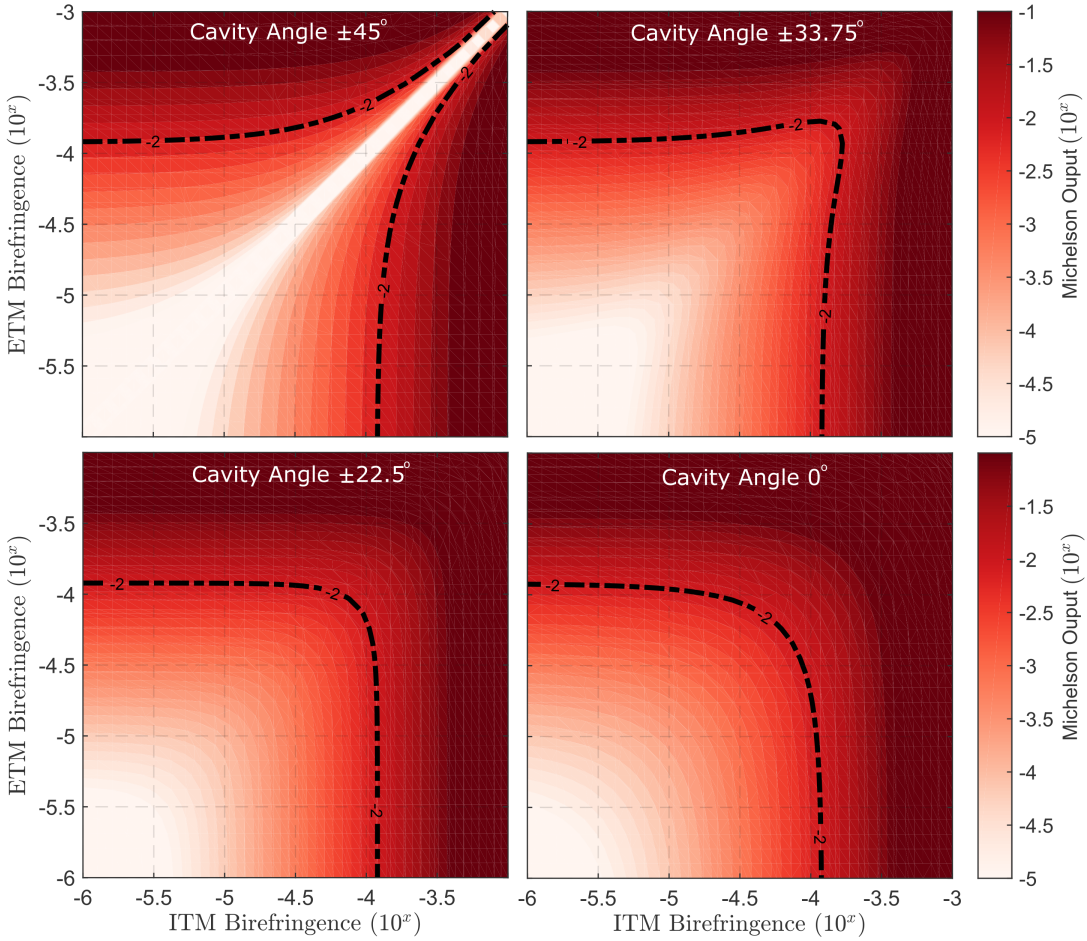


Figure 6.6: Contour plot of Michelson output signal for a single arm of a polarisation Sagnac speedmeter as a function of cavity mirror birefringence and relative alignment. For this result the two arm cavities were taken to have the same level of mirror birefringence and relative alignment. For each plot, the limit where the amount of Michelson signal is equivalent to 1% (10^{-2}) of the speedmeter signal is plotted. The cavity angle refers to the relative orientation of the ITM and ETM fast axes. In this model a perfect PBS and QWP were considered.

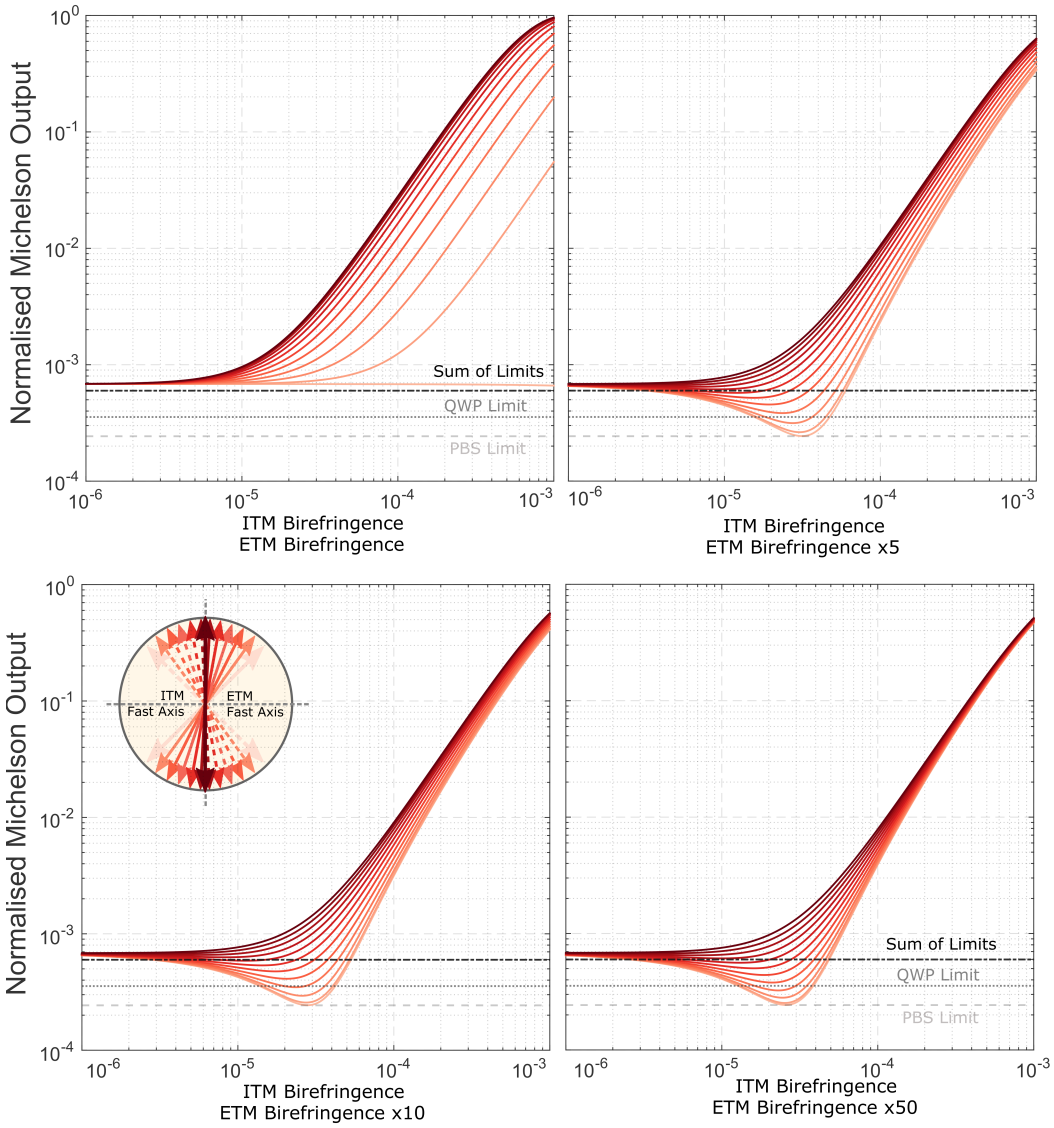


Figure 6.7: Effect of cavity birefringence on the amount of light leaking from a single cavity of the polarisation Sagnac speedmeter. To investigate the possible combinations of mirror parameters, the Michelson leakage is plotted as a function of mirror birefringence for four different imbalances in ITM to ETM birefringence ($\frac{\eta_{\mathrm{ITM}}}{\eta_{\mathrm{ETM}}} = 1, 5, 10, 50$). Each subplot also contains the modelled result for different orientations of the two cavity mirrors. The inset in the third subplot contains the key of mirror orientation versus line colour. The limitation due to the PBS and QWP imperfections are shown in grey dashed and dotted lines respectively.

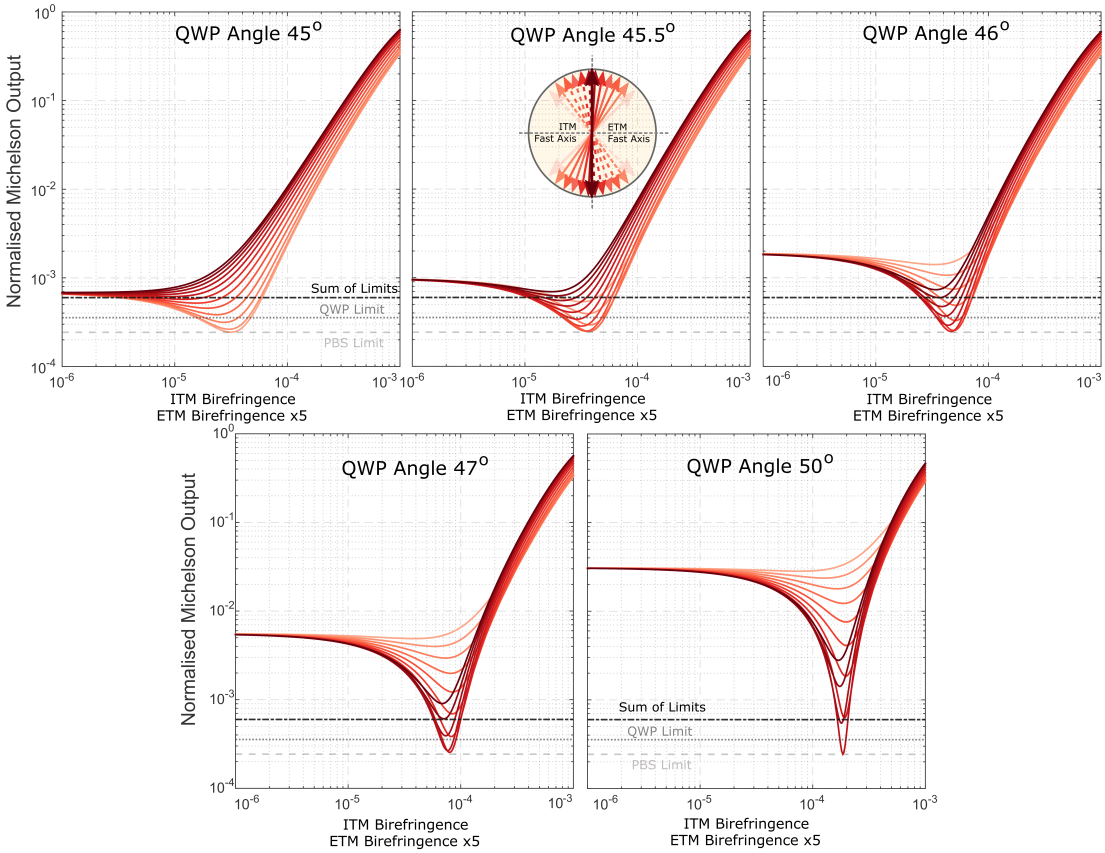


Figure 6.8: Combined effect of cavity and QWP birefringence on the amount of light leaking from a single cavity of the polarisation Sagnac speedmeter. In each subplot the cavity is modelled for a different rotation of the QWP. The insert in the second subplot contains the key of mirror orientation versus line colour. The minimum limitation due to the PBS and QWP imperfections are shown in grey dashed and dotted lines respectively.

undesirable higher level of birefringence of the cavity mirrors can result in a lower level of leakage with the correct relative orientation of the cavity mirrors. In this case the combination of QWP and cavity birefringence can cancel out and the interferometer is only limited by the extinction ratio of the PBS. For a cavity mirror birefringence less than 10^{-5}, there is insufficient birefringence in the cavity to compensate of the imperfect birefringence of the QWP.

In Figure 6.8 the Michelson leakage from a single arm of the polarisation Sagnac speedmeter is plotted verses cavity birefringence. This is plotted for five difference orientations of the QWP to show that by rotating the QWP different levels of cavity birefringence can be compensated to the level limited by the PBS. By misaligning the QWP, the limit due to QWP imperfections is increased if the combination of cavity birefringence and QWP birefringence do not cancel each other.

It was previously shown in Figure 6.7 that the QWP birefringence can be compensated with a given level of cavity birefringence. In Figure 6.8 it is shown that (for the range considered) any level of cavity birefringence can be compensated by rotating the QWP, reducing the amount of Michelson signal in the readout to the limitation set by the extinction ratio of the PBS optic.

In Figure 6.9 the combination of cavity and QWP birefringence is shown. This figure is similar to the plots shown in Figure 6.6 except here a realistic QWP in both arms with the measured parameters is considered, in combination with birefringent arm cavities and an ideal PBS.

This result shows that for a full speedmeter, an optimal combination of QWP and cavity birefringence is possible, but away from this solution the unwanted birefringence of the QWP and the cavity will limit the performance of the speedmeter. It should be noted that if a realistic PBS was considered, the optimised level would be increased to that imposed by the extinction ratio of the PBS.

6.3 Michelson Interferometer with Cavity Birefringence

In addition to the work modelling the polarisation Sagnac speedmeter based on measurements of the limitations of polarisation optics, a statement can also be made about the limitations of other speedmeter schemes that also rely on the optical separation of light in an interferometer based on polarisation.

In the polarisation circulation speedmeter schemes [76, 77] polarising optics are implemented at the output of a Michelson interferometer in conjunction with

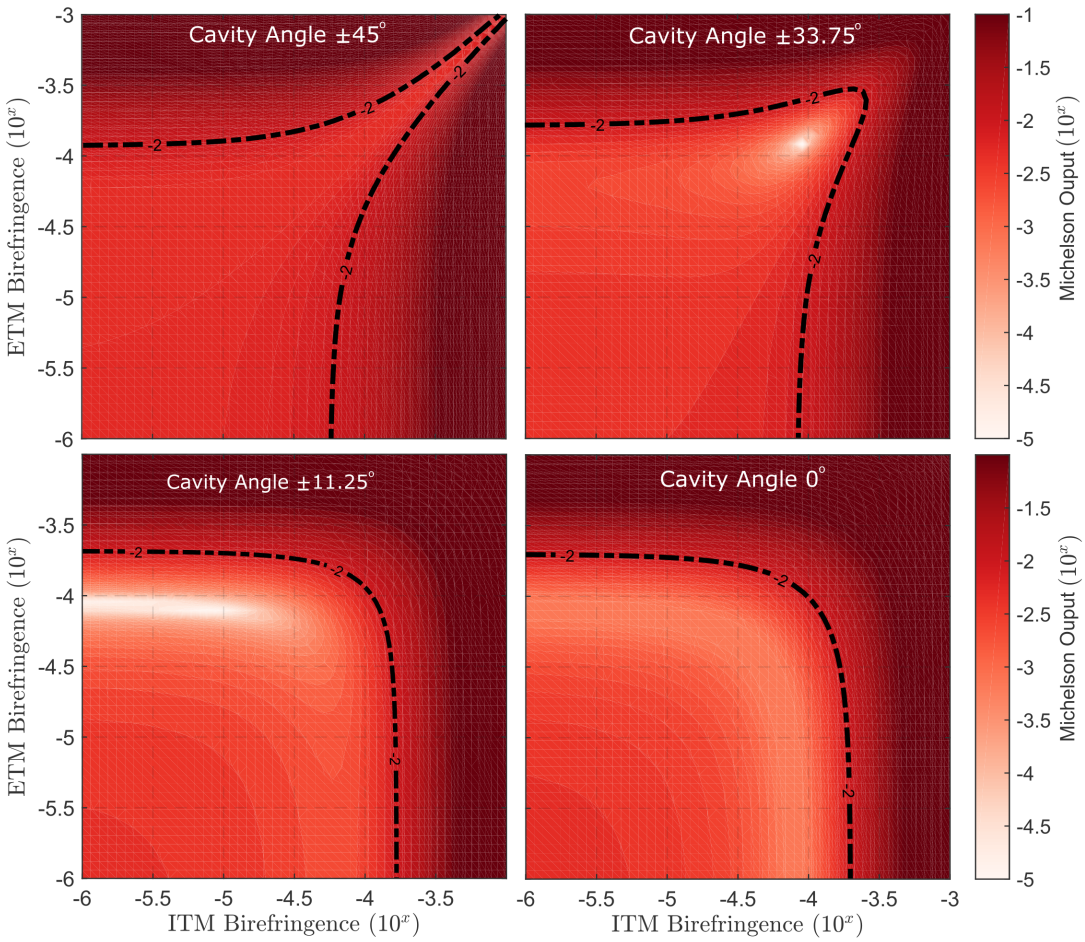


Figure 6.9: Contour plot of Michelson output signal for a single arm of a polarisation Sagnac speedmeter with measured QWP parameters as a function of cavity mirror birefringence and relative alignment. In each plot the limit where the amount of Michelson signal is equivalent to 1% (10^{-2}) of the speedmeter signal is plotted. For this result a perfect PBS was considered.

a recycling mirror or cavity to enforce a second interaction of the signal beam with the interferometer test masses. This approach relies on the preservation of the polarisation state of the light throughout the interferometer.

A Michelson interferometer was modelled with birefringent cavity mirrors in the two arms. In the model the interferometer was illuminated with p-polarised light and the ratio of s-polarised light to p-polarised light at the readout was modelled as a function of cavity birefringence with different orientations of the cavity mirrors and different offsets in the mirror birefringence. The result of this modelling work is shown in Figure 6.10 where it is shown that for a cavity birefringence greater than 10^{-4} radians and an offset between ITM and ETM birefringence greater than two, the level of s-polarised light at the readout can increase above 1 %.

From this result it is important to note the effect of birefringence offset between cavity mirrors. In the result shown in Figure 6.10, the level of s-polarised light at the readout increases to a significant level due to the mis-match in cavity mirror birefringence, not due to the absolute level of birefringence of each mirror. The level of polarisation effect is greater when the birefringence of the ETM is half that of the ITM, whereas the minimum level of change in the polarisation of the light in the interferometer is seen when the cavity mirror birefringence levels are equal.

6.4 Summary

In this chapter the polarisation optics characterisation results from Chapter 3 and the results of the linear cavity experiment from Chapters 4 and 5, are combined in a realistic model of a polarisation Sagnac speedmeter. These results are used to determine the limitations of this scheme in terms of its reliance on polarisation optics to define the propagation of the light within the interferometer. Imperfect polarisation optics cause light to couple to the interferometers readout along paths that will affect the sensitivity of the detector by degrading the speedmeter principle described in Section 2.4. The results of this work and the impact on future research into speedmeter interferometers for gravitational wave detection is summarised in Chapter 7.

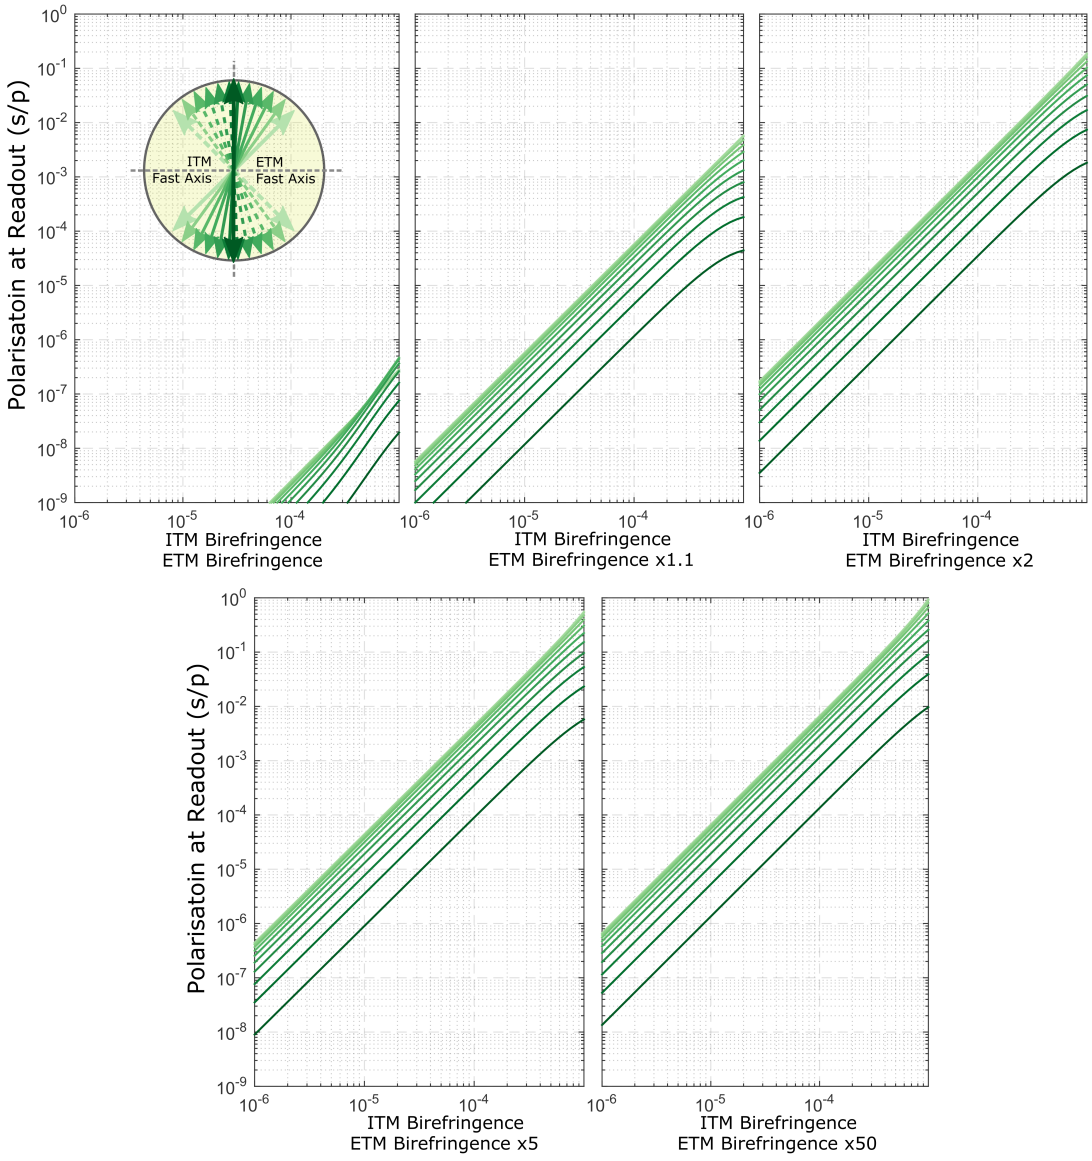


Figure 6.10: Modelled polarisation at the output of a Michelson interferometer with birefringent cavity mirrors. In each subplot the ratio of s-polarised light power to p-polarised light power at the interferometer readout is plotted as a function of mirror birefringence. For each subplot a different ratios of ETM to ITM birefringence is considered in each plot ($\frac{\eta_{\mathrm{ITM}}}{\eta_{\mathrm{ETM}}} = 1, 5, 10, 50$). The effect of mirror orientation is also shown as each subplot shows the maximum range of effect due to mirror rotation with the insert in the first subplot showing the colour key for mirror orientation.

Chapter 7

Conclusion

The polarisation Sagnac speedmeter interferometer is a design of interest for future gravitational wave detectors. This topology provides the attractive benefits of the speedmeter while maintaining linear arm cavities and with no requirement for additional optics or cavities in the readout port. The performance of such an interferometer will strongly depend on the quality of the polarisation optics which are core components of the interferometer.

Chapter 3 presented the results of table-top experiments to characterise the performance of polarisation optics required for the polarisation Sagnac speedmeter.

The extinction ratio of the polarising beam-splitter (PBS) optic was measured to be 770 ± 6.5 in transmission and 262 ± 1.8 in reflection for an angle of incidence of 45°, while a higher extinction ratios was measured at 42° angle of incidence. The quarter waveplate optic was measured to have a birefringence equivalent to $\frac{\lambda}{4} + \frac{\lambda}{324}$ or 1.590 ± 0.00184 radians.

In addition to the polarisation optics, the birefringence of non-polarising optics must also be considered. In Chapters 4 and 5 the birefringence of a linear optical cavity was measured.

Chapter 4 included a detailed description of an experimental system to investigate changes to the polarisation of light circulating in a 10 m linear cavity due to the birefringence of the cavity mirrors. The results of this experiment were analysed in Chapter 5.

This work showed that birefringence of the mirror coatings has a measurable effect on the state of the polarisation of the light resonating in the cavity. By rotating the ETM mirror the change in polarisation of the light could be measured and a parameter estimation of the magnitude and orientation of the birefringence of the two cavity mirrors was made. The measured data was found to be con-

sistent with an ITM of birefringence in the range from 1.29×10^{-4} radians to 1×10^{-3} radians, and an ETM of birefringence in the range from 2×10^{-5} radians to 9.12×10^{-5} radians.

This level of birefringence in the cavity mirrors is significant as the polarisation effect is amplified by the finesse of the cavity as the light interacts multiple times with the mirror coating. This cavity birefringence must be a consideration when discussing the limitations of the polarisation Sagnac speedmeter scheme.

In Chapter 6 a combination of polarisation effects was investigated in terms of the performance limitations of a polarisation Sagnac speedmeter. It was shown that the combination of QWP and cavity birefringence effects could cancel each other out by manipulation of the relative alignment of the birefringent optics. Such a cancellation scheme could be implemented in a full speedmeter scheme but would rely on knowledge of the polarisation parameters of the component optics which would require characterisation similar to that carried out here.

This is an advantage of the polarisation Sagnac scheme over other polarisation dependent speedmeter designs. Each arm contains a QWP meaning the birefringence of each resonant cavities can be independently cancelled. In a polarisation circulation scheme cavity birefringence would cause a similar problem but could not be compensated for in the same way.

In the case where the QWP and cavity birefringence effects are optimally cancelled, the performance of the polarisation Sagnac speedmeter scheme will still ultimately be limited by the extinction ratio of the PBS optic.

It was shown (in Figure 6.3) that with the PBS optic characterised in Chapter 3 in a polarisation Sagnac speedmeter the ratio of Michelson to speedmeter signals at the readout due to PBS extinction ratio will be 0.294%. This can be improved to be less than 0.1% by decreasing the angle of incidence of the this particular PBS optic to be below 43°.

7.1 Future Outlook For Speedmeter Interferometry

In this work it was shown how polarisation effects in a polarisation Sagnac speedmeter can introduce a source of optical loss and couple Michelson signal fields to the interferometer's readout. From our understanding of the speedmeter as a QND measurement scheme it is known that this will degrade the speedmeter mechanism. To fully understand the impact this has in the sensitivity of such a speedmeter a detailed quantum noise model is required; such work is out with the

scope of this thesis. Previous work has modelled the impact of losses [73, 74] and asymmetries [82] in speedmeters and a recent thesis [72] investigated backscattering in ring Sagnac schemes which has a similar effect as birefringence effects in a polarisation speedmeter.

The main issue of concern with the polarisation Sagnac speedmeter remains the quality of the polarisation optics. This work has shown that optics of a sufficiently high quality are available for prototype-scale experiments. Future work will need to show that polarising optics with high performance coatings can be fabricated on larger optical substrates for kilometer-scale detectors.

In a full-scale polarisation Sagnac speedmeter the PBS optic is the central beam-splitter between the arms of the interferometer which each contain a QWP. Both the PBS and QWP optics will be required to have large apertures and will be illuminated with a high level of laser power.

If PBS and QWP optics of similar quality to those characterised in this work can be fabricated, and if cavity birefringence can be cancelled in the way proposed here, then the polarisation Sagnac speedmeter will be a suitable interferometer topology for future gravitational wave detectors.

If large-scale polarisation optics of sufficient quality are not technically viable, an alternative speedmeter scheme will need to be considered. In the polarisation circulation speedmeter the PBS and QWP optics are part of the readout optics and as such the size requirement and optical power tolerance would be less demanding than in the polarisation Sagnac scheme. Cavity birefringence would still need to be compensated for either with alignment of the cavity mirrors or with the addition of birefringence compensation waveplates in the arms of the interferometer.

A detailed design study of both polarisation speedmeter schemes is needed to establish the requirements on the optics. Furthermore, a technical study of the viability of fabricating such optics is required as further work beyond the scope of this thesis before a decision can be made on the best speedmeter topology for future gravitational wave detectors.

Part III

Laser Development for Current and Future Gravitational Wave Detectors

Chapter 8

Introduction to Laser Technology for Gravitational Wave Detector Instrumentation

Laser interferometers for gravitational wave detection employ continuous wave, monochromatic, infra-red lasers as the probe light of a Michelson interferometer. The objective is to measure the relative length of two orthogonal arm cavities with enough precision to detect the length fluctuations coupled to the interferometer from a source of gravitational waves. This objective places very strict requirements on the optical power and noise performance of the laser light source.

This chapter is an introduction to the laser technology relevant to gravitational wave instrumentation, laser noise reduction, and the laser system employed in the Advanced LIGO detectors.

8.1 Laser Fundamentals

To generate a coherent beam of laser light, a source of photons produced by spontaneous emission, and a gain medium are required. In a beam of light, the energy (E) and frequency (f) of a photon are related by the Planck constant $E = hf$.

The internal energy levels of an atomic structure are discreet quantised levels. The structure and separation of energy levels is a unique characteristic of a given material. An atom can transition between energy levels by the absorption or emission of a photon of energy equivalent to the energy gap,

$$E_i - E_j = hf_{photon}. \tag{8.1}$$

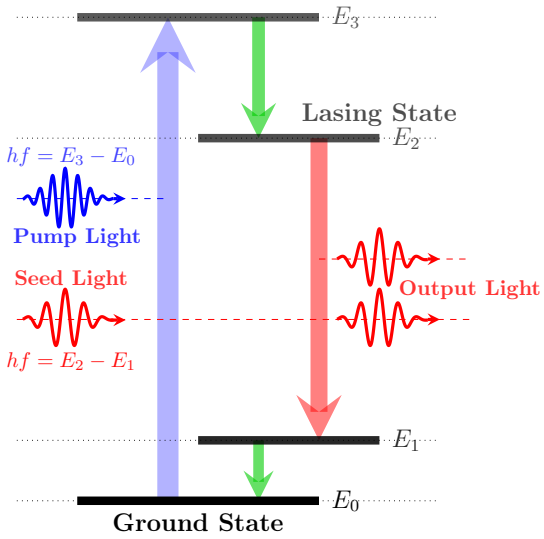


Figure 8.1: Basic operating principle for a four-level laser. The (blue) pump light photons are absorbed by the gain medium, raising the internal energy level from the ground state E_0, to the upper energy level E_3. From there, relaxation to the lasing state via a non-radiative process occurs (top green arrow). Forced relaxation from lasing state to the energy level E_1 is prompted by the interaction of a seed light (red) photon that causes the stimulated emission of a second photon. From the E_1 energy level a second non-radiative process lowers the energy back to the ground state, where another pump photon can be absorbed. A three-level laser process can also be implemented in which relaxation directly from the lasing state to the ground state occurs via stimulated emission. To achieve efficient laser emission population inversion is required where the population of the lasing state is greater that the ground state.

Stimulated emission can occur when a photon interacts with an atom in a raised energy level. The atom emits a photon of the same energy as the incident photon and relaxes into a lower energy state [101].

Spontaneous emission occurs when an atom relaxes to a lower energy level by emitting a photon, unprompted by any external process [102].

In a laser device a gain medium is excited to raise the energy level of the internal atoms. From this state, stimulated emission is triggered when a seed photon of the desired output frequency illuminates the gain medium.

To ensure a steady emission of laser light, the pump light needs to be of an intensity to ensure a steady population of energised atoms is maintained. In this way an equilibrium is achieved between emission and absorption.

The laser material must be carefully chosen such that the internal energy-level gaps match the desired frequencies of emission. Another important parameter is the lifetime of the energy level, this is the typical time an atom will remain in a

raised energy level before relaxing though spontaneous emission [84].

A multi-level system is typically implemented with pump and seed light of different wavelengths (see Figure 8.1).

There is a distinctintion between the terms laser *oscillator* and *amplifier*. A laser oscillator is system that contains a source of photons and a gain medium. Spontaneous emission initiates the lasing process and stimulated emission produces the light amplification process. A laser amplifier contains a gain medium but an external seed light field is required.

8.2 Laser Technology for Gravitational Wave Interferometry

8.2.1 Source Lasers

The laser system of a gravitational wave detector is required to deliver a constant level of power to the interferometer in a fundamental mode Gaussian beam. The first stage of the laser system is the known as the source or seed laser. Fluctuations in the laser intensity and frequency will couple to the output port of the interferometer, and the spatial properties of the laser beam determine the resonant coupling of the light to the optical cavities of the interferometer.

Solid-State Lasers

The current generation of gravitational wave detectors use solid-state laser devices with an emission wavelength of 1064 nm as the source of light for the interferometer. The term 'solid-state' refers to the gain medium being a solid, generally a crystalline or glass material.

More specifically, all current gravitational wave detectors use non-planar ring oscillator (NPRO) lasers as the first stage of the generation of the probe beam [103–105].

The gain medium is neodymium doped, yttrium aluminium garnet (Nd:YAG) cut into a coffin shape such that multiple internal reflections from the crystal faces form a ring cavity [106].

In such a device, the output frequency can be controlled via thermal and mechanical manipulation of the crystal to tune the optical path length in the resonator. By changing the temperature of the crystal with heater elements the length of the resonator can be controlled. Additionally, the crystal can be physically squeezed using a piezo-electric transducer.

This provision for control of the laser's output frequency is invaluable for interferometry applications, and along with the low inherent noise is the reason these laser devices were adopted for the initial gravitational wave detectors [107].

Diode Lasers

In a laser diode a current is passed through a semiconductor's p-n junction, this results in the generation of electron-hole pairs in the semiconductor material. The recombination of an electron-hole pair results in the emission of a photon, similar to the emission of photons due to relaxation of energy states (see Figure 8.1), this emission can be spontaneous or stimulated.

There is a subtle difference between the terms *laser diode* and *diode laser*. A diode laser is a class of laser device built around a laser diode. In some cases the two terms are used interchangeably but here the term diode laser is used to refer to the larger laser device that includes additional elements to enhance the output of the laser diode. For example, a diode laser will typically include some implementation of an optical resonator to increase the efficiency of the light amplification and to reduce variations in the frequency of the emitted photons.

An extended cavity diode laser features a laser diode inside a resonant linear cavity such that photons generated by the diode interact multiple times with the diode producing more photons by stimulated emission. As with other laser devices the frequency of emission is determined by the resonant frequency of the cavity.

Diode lasers in general have different inherent noise compared to solid-state lasers. In a diode laser the frequency noise is due to two effects. Spontaneous emission in the semiconductor material produces a white noise level in combination with a 1/f noise due to leakage current in the semiconductor material [108–110]. Different diode lasers exhibit different levels of frequency noise depening on the internal design of the laser cavity [111].

Many types of laser diode exist as they are a highly versatile class of laser device with many applications. In the field of gravitational wave interferometry, laser diodes are used to produce the pump light for the laser and amplifier devices, but in future gravitational wave detectors may be used directly as the laser light sources.

In this thesis an extended cavity diode laser with an emission wavelength of 1550 nm is considered for interferometry work relating to future gravitational wave detector research (see Chapter 11).

Fibre Laser

In a fibre laser system, the gain medium is a length of optical fibre doped with rare-earth elements. The gain fibre is typically pumped with light from one or more fibre coupled laser diodes. The seed light and pump light are coupled into the same fibre at a fibre splice. Fibre resonators can also be made to increase the optical amplification by making the fibre ends partially reflective.

Fibre lasers are relevant to future work towards cryogenic interferometry for gravitational detectors [112, 113] and may be used as the laser sources for future detectors as they offer longer wavelength emission which is required for detectors that will use mirror materials not compatible with 1064 nm lasers (see Section 8.2.4).

8.2.2 Laser Amplifiers

Seed lasers in isolation do not produce enough optical power for a gravitational wave interferometer to reach the sensitivity level required for detection, as the shot noise limited sensitivity of a detector scales with the square-root of the laser power. The required level of optical power is achieved with a combination of seed laser devices and laser amplifiers.

MOPA Amplifier

Also a solid-state laser device, a master-oscillator power-amplifier (MOPA) is a device that uses a crystalline gain medium to amplify the optical power of a seed beam.

Typically, multiple amplifier crystals are deployed in series, each crystal is illuminated with pump light to raise the atomic energy level into a lasing state. The laser beam is then passed though the gain media and produces stimulated emission of coherent photons of the same frequency as the seed light.

Ring Oscillator Laser

A ring oscillator laser amplifier is another class of solid-state laser amplifier, similar to a MOPA, with a pumped crystalline gain medium. The key difference is that the gain medium is placed inside a resonant cavity, meaning the light interacts multiple times with the gain medium [114].

This produces a greater level of optical amplification, but requires high levels of pump light to maintain the population of energised atoms in the gain medium [115].

The resonant cavity length can be controlled to ensure a single frequency and spatial mode is dominant. However, the fact that the cavity needs to be locked on resonance presents an additional degree of technical complexity and this locking loop can lead to noise being added onto the output laser beam.

Fibre Amplifiers

Similar to fibre lasers described above, laser fibre amplifiers use a doped fibre as gain medium to amplify the optical power of a fibre coupled input laser beam.

Fibre amplifiers have the potential to produce the high levels of optical power required of a laser system for gravitational wave detectors and such devices are currently in development for existing and future gravitational wave detectors [116].

A fibre amplifier is used in the 1550 nm laser system developed for the Glasgow 10 m prototype as described in Chapter 11.

8.2.3 Laser Modulators

Modulation of the laser beam is crucial for control of optical cavities and stabilisation of a laser noise for gravitational wave interferometry. Modulation is used to add sidebands to the main carrier field.

In a laser frequency stabilisation servo or arm cavity locking servo, phase modulation sidebands are used to generate a Pound-Drever-Hall error signal by beating the reflected sidebands with the carrier frequency. Modulators are used to actuate on the laser's frequency and intensity in order to stabilise the laser output to an external reference [117].

A mathematical description of the modulation of an optical field is contained in Section 1.4.4.

Electro-optic Modulators

To generate phase sidebands, electro-optic modulators (EOM) are typically used. This class of modulator uses materials that exhibit the Pockels effect, where the birefringence of the material is modified by an electric potential applied across the material.

An EOM device can be set up to change the phase of a linearly polarised light field passing through the optic by aligning the polarisation of the light to be parallel to an electric field applied across the crystal. The electric field induces a birefringence that changes the refractive index of the crystal while not changing the polarisation of the light if the two are well aligned to a common axis. This

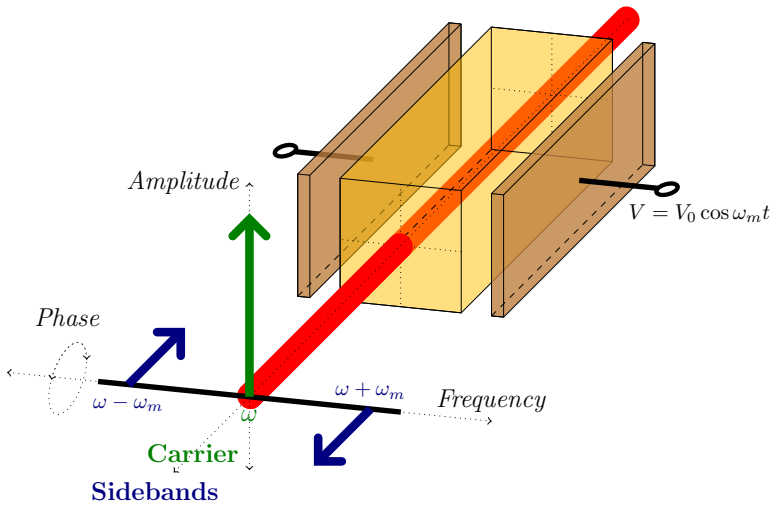


Figure 8.2: Operating principle of an electro-optic modulator as a phase compensator. In this example a sinusoidally varying voltage V is applied across the optic. This generates upper and lower phase modulation sidebands (blue arrows) on the carrier light (green arrow) as it propagates though the optic. The frequency separation of the sidebands from the carrier is defined by the frequency of the applied voltage signal. The amplitude of the generated sidebands is dependent on the amplitude of the applied voltage signal and the electro-optic susceptibility of the crystal. For simplicity only the first order modulation sidebands are represented in the diagram.

change to the refractive index in the polarisation direction of the light modifies the optical path length, and therefore the amount of phase accumulated by the beam as it propagates through the crystal.

An electric field can be applied to shift the frequency of a light field, as frequency is the integral of phase. If a sinusoidal voltage is applied to the EOM, the variation in the optical path length will cause some fraction of the optical field to be phase modulated, generating phase modulation sidebands. The amount of the optical field coupled into the modulation sidebands depends on the amplitude of the electric field. The frequency of the resulting sidebands is the sum and difference of the main carrier frequency of the optical field (ω_0) and the frequency of the applied electric field (ω_m) [118].

A carrier field with phase modulated sidebands is approximated for small modulation indices by the expression,

$$E = E_0 \exp(i\omega_0 t)\Big(1 + i\frac{m}{2}\big(\exp(-i\omega_m t) + \exp(i\omega_m t)\big)\Big). \tag{8.2}$$

where m is the modulation index.

In combination with an output polariser, an EOM can also be used as an amplitude modulator by setting the polarisation of the input light to be at 45°

relative to the applied electric field. This will result in polarisation modulation which in combination with a polariser at the output provides intensity modulation.

Modulators can also be used as frequency actuators. A signal applied to an EOM can be used to correct for fluctuations of the optical fields frequency and stabilise the output. As a modulator of phase, the frequency response of an EOM modulator increases linearly with frequency. The bandwidth of an EOM actuator is limited by the geometry of the crystal and the beam size as well as the speed of the modulation driver.

Acousto-optic Modulators

A different class of optical modulator that can also be used to change the intensity or frequency of a laser beam are acousto-optic modulators (AOM). An AOM is a crystal device that when mechanically driven with a piezoelectric transducer, generates acoustic waves that propagate through the crystal. The acoustic pressure waves in the crystal form a grating pattern from which the optical field is diffracted.

The amount of power diffracted from the input beam path can be modulated by changing the strength of the acoustic wave generation. The intensity of the transmitted beam can be controlled with the voltage drive of the piezo device generating the grating. As the strength of the acoustic wave generation is increased more power is coupled out of the main beam into higher diffraction orders.

By modulating the frequency of the acoustic drive, the frequency of the diffracted light is modulated. The diffracted beams are geometrically separated from the zeroth order input beam path by diffraction.

AOM actuators are used as both frequency and intensity control in gravitational wave detectors. The use in the LIGO laser stabilisation system is presented in Sections 8.4.3 and 8.4.2.

8.2.4 Wavelength choice for gravitational wave detectors

Currently all gravitational wave detectors use solid-state NPRO lasers with output at an infra-red wavelength of 1064 nm. This choice was motivated by the availability of lasers at this wavelength of appropriately low noise with good internal actuation, together with the availability of low-loss mirrors that can be fabricated to have low thermal noise properties and be highly reflective at this wavelength.

Future developments in the material science for gravitational wave detector optics require a move away from using 1064 nm laser light to a new dawn of longer wavelength lasers, either at 1550 nm or at 2 μm. The need to change the

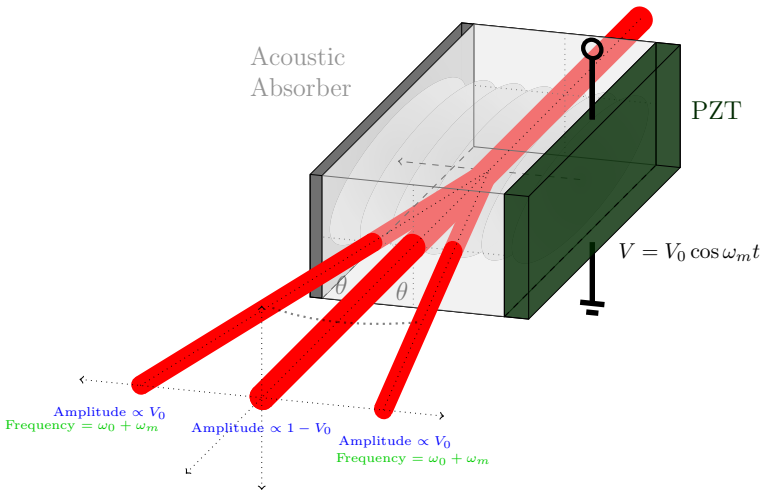


Figure 8.3: Operating principle of an acousto-optic modulator as an intensity or frequency modulator. A diffraction grating is formed by acoustic waves in a crystal medium. The amplitude of the beam in the zeroth diffraction order is reduced by the strength of the modulation signal as more power is diffracted into the higher order beams. The frequency of the light in the diffracted beams is also shifted by the applied signal.

laser source is a consequence of the proposal to use cryogenic cooling to reduce the thermal noise of the suspended test mass mirrors. The sensitivity limitation due to suspension thermal noise and coating thermal noise arising from Brownian motion in the optical materials was described in Chapter 1, with a modelled Advanced LIGO noise budget shown in Figure 1.3 of Section 1.3.

By cryogenically cooling the test mass mirrors of the interferometer thermal noise is reduced. However, this option is costly and requires a major infrastructure upgrade. The KAGRA detector [28], currently being commissioned in Japan, already uses such technology and a cryogenic upgrade of the LIGO detectors has been proposed, the so called LIGO Voyager project [18]. Third generation detectors such as the Einstein Telescope [29] and the proposed Cosmic Explorer both include cryogenic suspensions in their respective designs.

Fused silica is used as the mirror substrate and suspension material for detectors operating at room temperature with 1064 nm laser light [13, 119]. However, the noise properties of fused silica deteriorate at low temperatures, making it unsuitable for a cryogenically cooled detector.

Materials with more suitable thermal properties at low temperature have been investigated, with crystalline silicon the leading candidate [29, 120]. Silicon is however opaque for light with a wavelength of 1 μm, making it unsuitable as a mirror substrate material for use with a 1064 nm laser, but it is transparent at longer wavelengths.

The KAGRA detector design takes an alternative approach, and still uses a 1064 nm laser with sapphire as the mirror material [121, 122].

Lasers with an emission wavelength of 1550 nm are commonly used as telecommunications devices, so the development of laser technology for this wavelength is very mature. Lasers at this wavelength, as well as amplifiers, detectors and modulators are readily available and of high quality. In contrast, laser technology at 2 μm, an alternative solution for cryogenic interferometry under consideration, is not as well developed and requires additional research and development. The ultimate decision on what wavelength is to be used for future gravitational wave detectors has not been made and research continues on laser systems at both wavelengths. The choice was made to transition to using a 1550 nm, stabilised light source for future interferometry experiments at the Glasgow 10 m prototype laboratory. The commissioning of this new laser system is described in Chapter 11.

8.3 Laser Noise and Stabilisation Techniques

The stability of the laser that illuminates an optical interferometer is of critical importance to the achievable sensitivity of that instrument. The laser is the probe used to sense the relative length of the detectors arms. Noise in the probe beam intensity or frequency can directly couple to the readout of the interferometer and can mask gravitational wave signals. Asymmetries in the detectors arm cavities lead to multiple mechanisms by which laser noise can couple to noise in the interferometer readout [123].

A change in the frequency of the laser is equivalent to a change in the length of the arms of the interferometer if the cavities are locked to the laser. Since a gravitational wave detector aims to measure changes in the length of optical cavities caused by a incident gravitational wave, frequency noise of the laser can easily hide an astrophysical signal behind technical noise of a detector with unmatched arms.

Intensity noise of the laser can also create false signals at the interferometer readout, as a fluctuation in the intensity of light detected in the readout if the interferometer is operating away from the dark fringe condition [124].

In addition, intensity noise of the laser contributes length noise to the cavity via the radiation pressure force exerted on a suspended mirror by an incident beam: if the intensity of the laser varies then so too does the strength of the radiation force pushing the mirror [15]. Difference in the optical absorption of

the cavity mirrors in each arm lead to a difference in the power build up factor inside the cavity. This asymmetry of the arm cavities leads to a coupling of laser intensity noise to differential length noise [125]. This issue scales with injected laser power and is a significant problem for the operation of gravitational wave detectors with tens of watts of optical power at the input.

Frequency and intensity noise are inherent to any laser system but can be suppressed with feedback to appropriate actuators provided adequate sensors are available. Due to the weak nature of the gravitational wave signal, the noise requirements of a laser system for interferometer gravitational wave detector are extremely demanding.

8.3.1 Passive Noise Reduction

A method for reducing the noise of a laser beam is to use the passive filtering provided by a resonant optical cavity.

The frequency response of an optical cavity can be approximated as that of a first-order low-pass filter [123]. For frequencies above the linewidth of the cavity the response of the cavity is less than one and so noise fluctuations at frequencies higher than the linewidth are attenuated. In this way the frequency and intensity noise of a laser beam in transmission of the cavity will be reduced. If the length noise of the cavity is greater than the free running laser noise, then this additional noise will be added to the laser beam in transmission of the cavity.

In this way an optical cavity can act as a filter of a noisy laser beam, reducing the intensity and frequency fluctuations of the incident beam above the linewidth of the cavity. The corner frequency f_c of a Fabry-Perot cavity is defined as

$$f_c = \frac{c}{4L\mathcal{F}}, \tag{8.3}$$

where L is the length of the cavity c is the speed of light, and $\mathcal{F}$ is the finesse (see Equation 1.42 in Section 1.5.2).

In a passive filter the level of noise reduction is a fixed quantity dependent on the linewidth of the optical cavity. Passive filters are best at reducing the laser noise at higher frequencies, and the lower the linewidth of the cavity the greater the noise suppression.

Passive filtering of a laser beam with an optical cavity can also reduce the higher order mode content and pointing noise of a laser beam. Additionally, passive noise filter schemes can be combined with active noise stabilisation methods to produce high levels of noise suppression.

8.3.2 Laser Intensity Stabilisation

Any laser system will exhibit intensity fluctuations of its output light due to technical noise internal to the laser device. The source of this technical noise depends on the laser device. In laser diodes, electrical current noise leads to intensity noise of the emitted light. Technical noise can also arise from statistical fluctuations in the population of the lasing state or internal beam jitter depending on the laser device [126].

Whenever a laser beam is detected or interacts with an optic, there is an additional level of quantum shot noise due the quantisation of the photons that constitute a beam of light, see Section 2.2 for a full description of quantum noise in an interferometer.

Technical intensity noise can be reduced with active feedback, intensity fluctuations can be detected as power fluctuations with a photo-detector and an error signal can be generated to add the measured noise back onto the laser exactly out of phase to cancel out the free-running fluctuations down to the shot noise limit of the sensed power [123].

The shot noise limitation of a laser interferometer can be reduced by increasing the optical power illuminating the interferometer. Such an attempt to improve the sensitivity of Advanced LIGO is covered in Chapter 9, while the active intensity stabilisation system for Advanced LIGO is discussed in Chapter 10.

8.3.3 Laser Frequency Stabilisation

The frequency of a laser beam can be defined as a carrier frequency $f_0 = c/\lambda$ plus some frequency dependent noise fluctuation δf. For NPRO lasers and external cavity diode lasers the free running laser frequency noise is in part as a result of length fluctuations in the laser resonator. In diode lasers frequency noise can also be generated by fluctuations in the current through the diode junction.

Frequency fluctuations cannot be measured as simply as power fluctuations. It is necessary to employ a frequency discriminator to convert frequency fluctuations to intensity fluctuations that can be sensed with a photo-detector. In the field of gravitational wave interferometry, the ubiquitous method of laser frequency stabilisation is the Pound-Drever-Hall method using an optical cavity as a frequency reference.

The Advanced LIGO frequency stabilisation system is discussed in Chapter 10, and the frequency stabilisation of a 1550 nm laser is presented in Chapter 11.

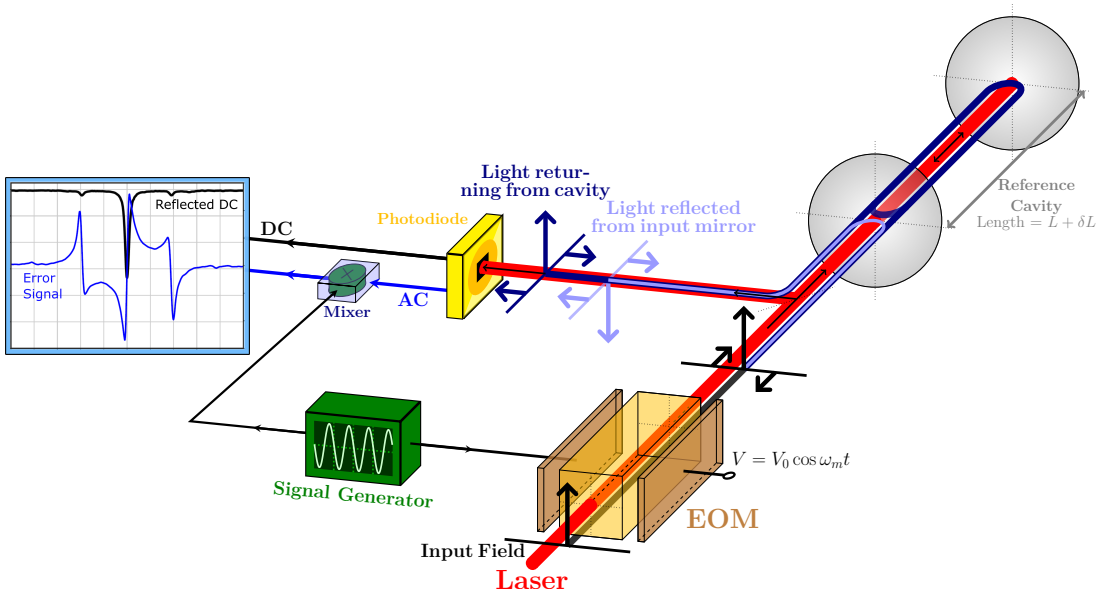


Figure 8.4: Visualisation of the generation of the Pound-Drever-Hall error signal from the phase modulated field reflected from a linear reference cavity. A signal generator produces a local oscillator signal that is used to add phase modulated sidebands to a laser beam and to demodulate the signal measured with a photodiode collecting the light reflected from a resonant cavity. The detected DC signal and the demodulated 'error' signal are plotted on the left against slow cavity detuning. On resonance the DC signal is at a minimum and the error signal is zero, either side of the resonance the DC voltage increases, and the error signal is a bipolar linear function of the cavity detuning.

Pound-Drever-Hall Laser Stabilisation

The Pound-Drever-Hall (PDH) technique is a method of laser frequency stabilisation that is widely used in gravitational wave interferometry [100]. In this scheme an optical cavity is used as a frequency reference and an error signal is fedback to one or more frequency modulators to correct the free running laser frequency noise [127]. A linear Fabry-Perot cavity is often used as a reference cavity but ring cavities can also be used.

An error signal is generated by looking at the signals reflected from a resonant cavity. It was shown in Chapter 1 that light will be resonant in a Fabry-Perot cavity if the length of the cavity is equal to an integer number of half-wavelengths. A cavity can be detuned from resonance by changing the length of the cavity L or the frequency f of the light. A fractional change in length or frequency are equivalent,

$$\frac{\delta f}{f} = \frac{\delta L}{L}. \tag{8.4}$$

For an optical cavity on resonance the amount of light returning from the cavity is at a minimum. When the cavity is detuned the amount of return light increases, however the change is symmetric about the resonance point.

The intensity of light returning from the cavity passes through a minimum on resonance, in addition, the phase of the return light changes by 180° across the resonance.

In the PDH scheme the light is phase modulated before reaching the cavity to add sidebands to the carrier light to acts as frequency references. The frequency of the sidebands is defined relative to the carrier frequency $\omega_0 = 2\pi f_0$. The sideband frequencies are then defined as $\omega_{SB} = \omega_0 \pm \omega_m$ where ω_m is the frequency of the modulation signal.

The sidebands can also resonate in the cavity if the resonance condition is met for either of the frequencies ω_{SB}. The phase of the reflected light has a similar phase difference above and below resonance.

In the PDH scheme the light returning from the reference cavity is detected with a photodiode. The optical field at this detection point has multiple components. The measured field is a mix of light that has resonated inside the cavity and light that has been directly reflected from the input mirror. The field is also composed of carrier and sideband light.

For simplicity, this explanation is restricted to the case where the sidebands and carrier light are far enough separated in frequency that the linewidths of the carrier and sideband resonances do not overlap so carrier and sidebands do not resonate at the same time.

When an optical field is detected on a photodiode the resulting signal is proportional to the multiplication of the incident field with its own complex conjugate. When an optical field containing multiple frequency components is detected, beats between all frequency components are measured. A single frequency can be isolated by demodulating and averaging the photodiode signal with a local oscillator signal of that frequency. The beat between the carrier and the sidebands is down-converted to a DC signal voltage and up-converted to a signal at twice the local oscillator frequency. Either side of the resonance the phase of the carrier light rotates relative to the two sidebands which are unchanged. This changes the amplitude of the demodulated signal in a way that is useful.

The demodulated signal is filtered to remove the higher frequency components and to average the signal. What is left is a signal, that in the region close to the cavity resonance, is bipolar, linear and proportional to the length detuning of the cavity. This error signal is a measure of the difference between the laser frequency and the cavity resonance frequency.

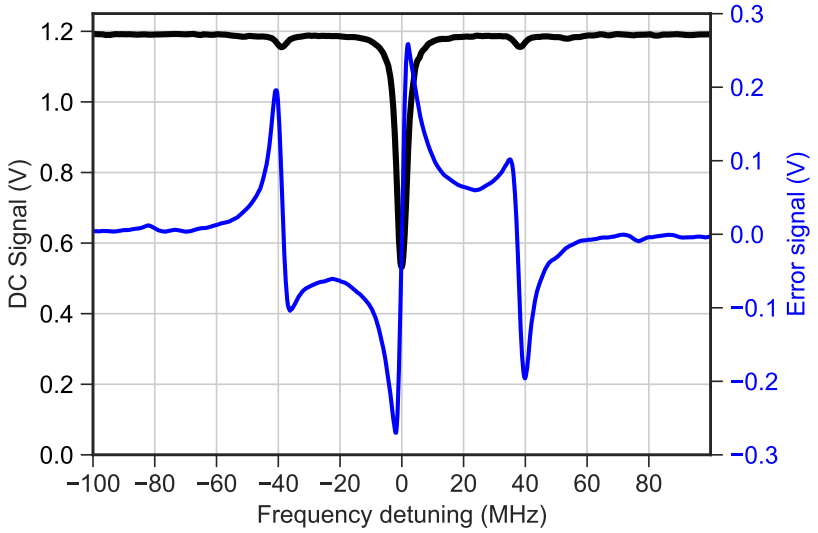


Figure 8.5: Example of Pound-Drever-Hall error signal (blue) the demodulated RF reflected signal is shown as a function of slow cavity detuning along with the reflected DC signal (black). The error signal is a bipolar linear function of the cavity detuning for small detunings about the resonance. The PDH error signal also provides a larger capture range as the the error signal between carrier and sideband is of the correct sign.

The PDH error signal is then filtered by a servo system to shape the frequency response of the feedback signal. The resulting signal is then applied to a frequency actuator to hold the laser locked to the cavity so the frequency of the laser follows the length of the cavity. If the length noise of the cavity is less than the equivalent frequency fluctuations of the free running laser then the laser noise will be reduced.

Consideration of the Limitations of Laser Stabilisation Schemes

The noise reduction provided by a PDH scheme is limited by the systems sensing noise. The shot noise at the detection photodiode can limit the sensing; this will be improved with increased optical power. In frequency stabilisation schemes with a rigid cavity the limiting factor will be the length noise of the cavity that acts as the frequency reference. For a rigid reference cavity, the stability limit is determined by the level of thermal length fluctuations in the spacer material. It is typical to use ultra-low expansion materials to form the rigid spacer to which the cavity mirrors are bonded. If this is the case, then thermally driven expansion will not significantly limit the stability of the cavity. Instead the limitation will come from thermal noise of the cavity mirrors. This thermal noise will either be dominated by the Brownian noise in the bond between mirror and spacer or the

Brownian motion in the cavity mirror coating causing the effective length of the cavity to fluctuate. The dominant effect will depend on the specification of the cavity mirrors and the construction of the rigid cavity.

For a suspended reference cavity, the stability will be limited by the residual motion of the cavity mirrors. This can be caused by the coupling of external seismic and acoustic disturbances to the motion of the mirrors. Due to the pendulum nature of a suspended cavity, the strength of this coupling will drop off rapidly above the pendulum resonance, so this coupling will mostly affect the motion of the mirrors at low frequencies. For a well isolated suspended cavity the dominant length noise will then be related to thermal noise in the suspension or mirror material.

A rigid cavity will typically provide a more stable frequency reference compared to a suspended cavity of similar length, however there is a technological limitation to the length of suitable spacers for rigid cavities meaning long baseline rigid cavities are not feasible. The achievable frequency stability scales linearly with the length of the cavity, so long suspended cavities can provide a better frequency reference depending on the application.

The Pound-Drever-Hall scheme can equally be used to stabilise the length of an optical cavity to the frequency of a laser. The error signal is generated in the same way as described above but the feedback is applied to an actuator of the cavity length. In this way the cavity is locked to the laser beam which remains resonant with the optical cavity. This method of cavity locking is employed in gravitational wave detectors to hold the multiple optical cavities of the interferometer on resonance with the input beam, which is pre-stabilised to another reference cavity to reduce the free-running laser frequency noise.

8.4 The Advanced LIGO PSL System

The laser system that provides the light for the Advanced LIGO interferometer is called the pre-stabilised laser (PSL) sub-system. This consists of the seed laser, two stages of optical power amplification, stabilisation loops to lower the noise of the laser's intensity and frequency, and a mode cleaner cavity to filter out temporal and spatial defects in the laser beam. A schematic of the PSL scheme is shown in Figure 8.6.

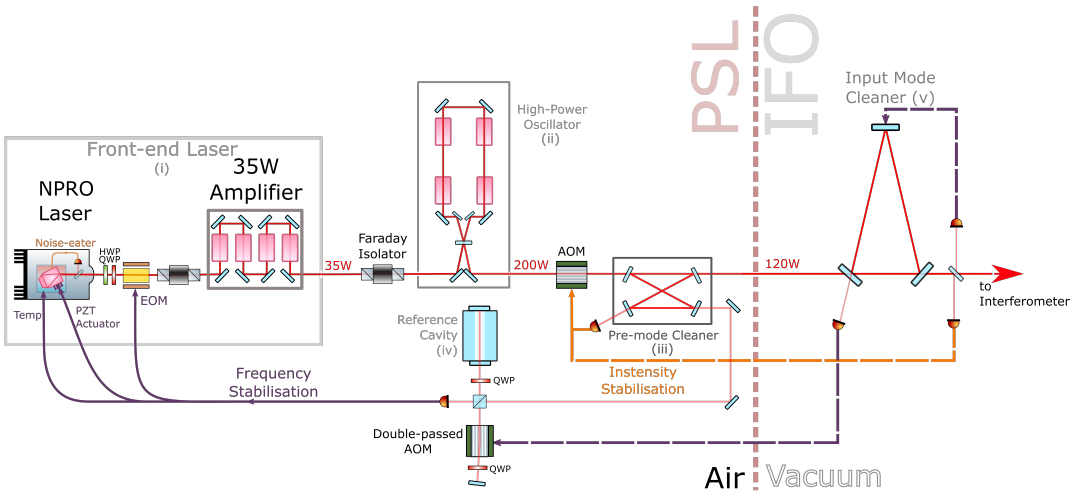


Figure 8.6: Overview of the Advanced LIGO PSL laser amplification and stabilisation scheme. The front-end laser (i) contains the NPRO and a MOPA amplifier producing 35 W of laser power. The next stage of power amplification comes from the high-power oscillator (ii). The PSL also contains the premode cleaner (iii), reference cavity (iv) and photodiode arrays for laser noise stabilisation. The output laser beam of the PSL is coupled into the main interferometer via the suspended input mode mleaner (IMC) inside the vacuum system.

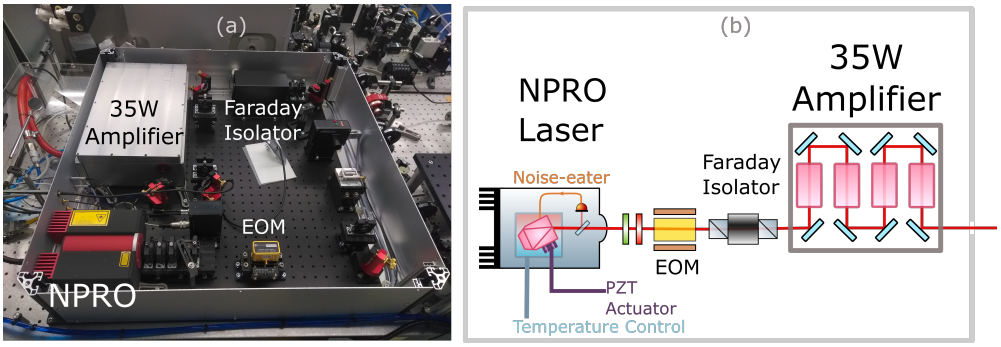


Figure 8.7: LIGO front-end laser. Inside of the front-end laser box (a). Schematic of LIGO front-end laser (b). The operation of the temperature control and PZT actuator is discussed in the text.

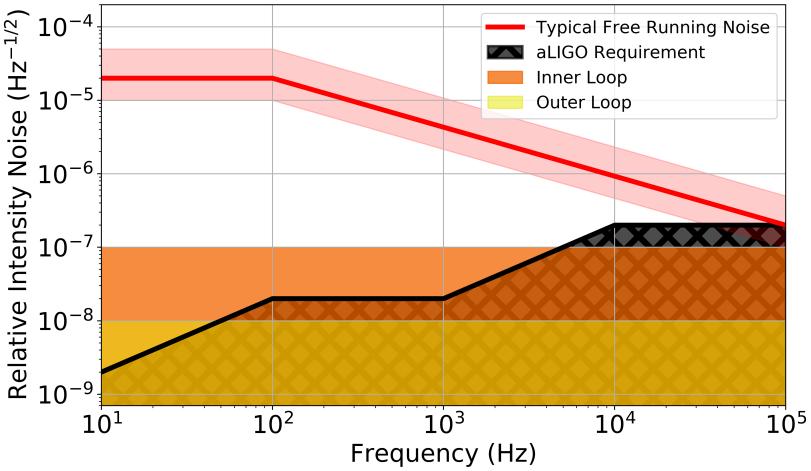


Figure 8.8: Stabilisation requirement of the ISS for Advanced LIGO. The typical free running intensity noise of the laser is shown in red. This noise has to be reduced down to the orange area by the inner loop of stabilisation feedback and further reduced into the yellow area by the outer loop to be below the requirement at 10 Hz. The dark grey area represents the Advanced LIGO design requirement and includes the passive filtering of the laser noise.

8.4.1 Laser Amplification

The first stage of the PSL system is an NPRO laser with an output of 1 W of optical power at a wavelength of 1064 nm. This laser has an internal intensity noise stabilisation loop called the 'noise eater' that suppresses the relaxation oscillation [107].

The output of this laser passes through a medium power amplifier, called the front-end amplifier, which is a solid-state MOPA consisting of four crystals, each crystal acts a laser amplifier as it is pumped by an external diode laser [128]. The NPRO laser and first stage of power amplification are commonly referred to as the front-end laser. The output of the front-end is a beam of 35 W of coherent light.

The front-end output beam is sent to the high power oscillator (HPO). The beam again passes through four optically pumped gain crystals, but in this instance a resonant cavity is formed so the beam passes through each amplification stage multiple times so the power amplification is greater. From this second power amplification stage it is possible to produce a beam of 200 W optical power. The HPO is described in more detail in Section 9.1.1 and in [129].

8.4.2 Intensity Stabilisation

For the laser to be a suitable source for a gravitational wave detector, the relative intensity noise of that laser has to less than $10^{-7}\,\mathrm{Hz}^{-1/2}$ at 1 kHz and $2 \times 10^{-9}\,\mathrm{Hz}^{-1/2}$ at 10 Hz [130].

The power fluctuations of the laser are passively filtered above the linewidth of the power recycling cavity of the main interferometer (described in Section 1.5.3).

The design requirement is met with the inclusion of active stabilisation in the form of the intensity stabilisation scheme (ISS). The level of intensity noise reduction required for the Advanced LIGO ISS is shown in Figure 8.8. The ISS is a double stabilisation loop, the intensity noise is sensed at two points in the beam path and an error signal is derived and fed back to a common intensity actuator.

The inner loop of the ISS measures the power of a beam picked off from the premode cleaner (PMC). The second loop takes a beam of higher power picked off after the in-vacuum suspended input mode cleaner.

The intensity stabilisation servo is limited by the sensing noise of the photodetectors. In the absence of technical limitations this is determined by the shot noise of the laser beam detected, which is proportional to the square root of the laser power. For the outer ISS loop the sensed beam is split over multiple photodiodes so that a higher power level can be sensed.

Both loops of the intensity stabilisation servo feedback to the same actuator, an acousto-optic modulator in the path of the high power laser beam. This couples some fraction of light from the main beam path to higher order diffraction modes. These diffracted beams diverge from the main beam at some angle and are dumped or measured with a power meter as a monitor of the performance of the ISS.

8.4.3 Frequency Stabilisation

The frequency noise of the laser at the input to the interferometer is required to be less than $10^{-6}\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 Hz, compared to the free running frequency noise of an NPRO laser which is of the order $100\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at 100 Hz [131].

The Advanced LIGO frequency stabilisation scheme (FSS) consists of three nested feedback loops, all three of which are Pound-Drever-Hall (PDH) schemes that use different actuators and different optical cavities as frequency references.

Each loop works to reduce the frequency noise of the laser light over different bandwidths with feedback to actuators before the sensing point of the proceeding loop.

The innermost loop senses the light reflected from a rigid reference cavity between the PSL laser and the PMC, and feeds back to two actuators internal

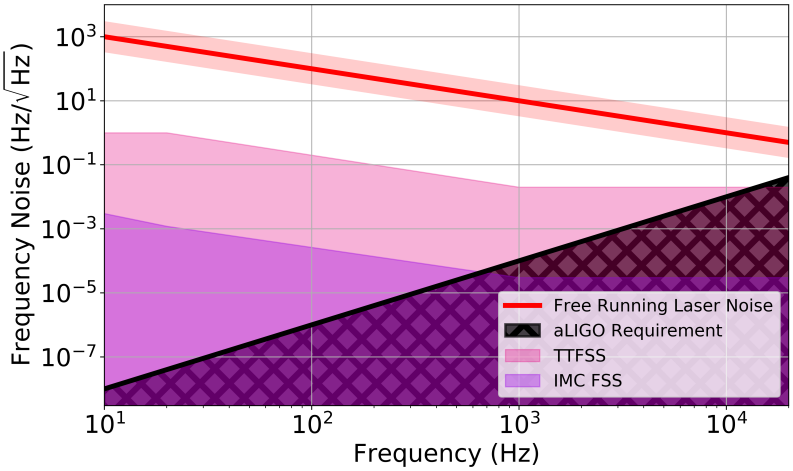


Figure 8.9: Stabilisation requirement of the FSS for Advanced LIGO. The free running frequency noise of the laser is shown in red. This noise has to be reduced down into the pink area by the table-top stabilisation system (TTFSS) and further reduced by the IMC frequency stabilisation into the purple area. The noise is further reduced into the dark purple area which represents the Advanced LIGO design requirement with a combination of the 4 km cavity stabilisation and the passive filtering of the laser noise by the power recycling cavity.

to the NPRO seed laser and a third EOM inside the front-end laser enclosure, as shown in Figure 8.7.

The second loop senses the light reflected from the input mode cleaner (IMC) and feeds back to an AOM actuator before the reference cavity on the PSL table. The AOM is implemented as a frequency actuator by adding a folding mirror after the AOM that reflects the first-order diffracted beam back though the actuator. This ensures the first order beam, which is frequency modulated, is reflected back along the input path. This modulates the frequency of the light before the reference cavity which is the sensing point for the innermost loop of frequency stabilisation system.

The outermost loop senses the light reflected from the main 4 km interferometer in the symmetric port (see Section 1.5.1), and feeds back to the length of the IMC to stabilise the length of the triangular cavity before the second loop sensing point. The feedback signal from this loop is also added to the error point of the second loop to actuate on the laser at the double passed AOM before the reference cavity.

This nested loop system work by adding frequency deviations sensed by the IMC and the 4 km arm cavities onto the laser beam before the frequency descriminator of the inner loop so that they are added to the sensing point of this loop.

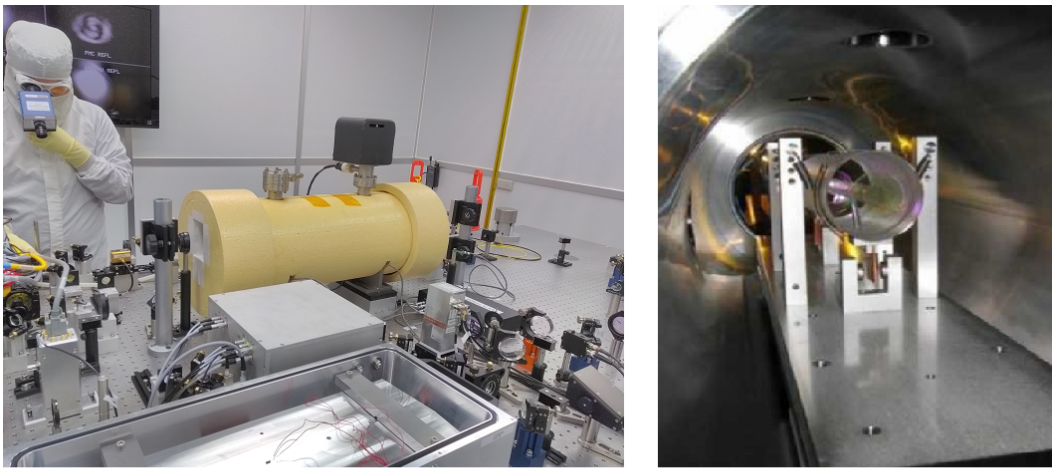


Figure 8.10: Frequency reference cavity. Vacuum chamber containing reference cavity on PSL optical bench (a). Inside the vacuum chamber showing suspended rigid cavity (b) [133].

This means that the signal fedback to the actuators (NPRO temperature, PZT and EOM) stabilises the frequency of the laser to all three cavities. The 4 km arm cavities provides the best frequency reference in the stabilisation scheme but the other two loops are still recuired for the lock acquisition process.

The reference cavity in the PSL system used for the inner loop frequency stabilisation system is a 20.3 cm long linear Fabry-Perot cavity formed by two mirrors bonded to a spacer made out of ultra-low expansion silica. This spacer is suspended and housed in a small vacuum chamber that sits on the PSL optical bench. This isolates the cavity from external seismic, acoustic and other environmental noise [132]. The reference cavity can be seen in Figure 8.10.

The stability of this loop is limited by the thermal noise of the cavity which is about 6×10^{-3} Hz/$\sqrt{\text{Hz}}$ at 100 Hz [130, 131].

The second loop of the FSS uses the IMC as a frequency reference. The IMC is a suspended 16.5 m triangular cavity inside the main LIGO vacuum system and as such is not part of the PSL. The error signal from the IMC is fed back to the double-passed acousto-optic modulator (AOM) that sits on the PSL optical bench before the reference cavity. The stability of this loop is limited to the order of 10^{-5} Hz/$\sqrt{\text{Hz}}$.

The outer most loop of the FSS uses the 4 km long Fabry-Perot arm cavities of the main detector as frequency reference. The signal for this is measured in the symmetric port of the main Michelson, and is fed back to the AOM and to the length control of the IMC. The stability of this loop allows the noise of the

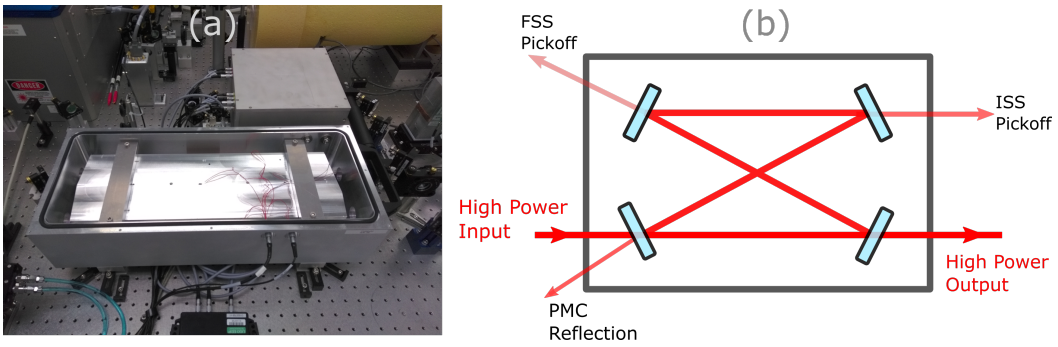


Figure 8.11: Premode cleaner cavity. Air-sealed isolation tank containing PMC on PSL optical bench (a). Diagram of the optical cavity (b).

laser to be stabilised to below 10^{-8} Hz/$\sqrt{\text{Hz}}$ at 100 Hz which is well below the requirement of the PSL.

The inner loop of the FSS is necessary to reduce the free running laser noise so that the IMC can be locked. The principal purpose of the outer two loops is to provide the noise suppression by a combination of the gain of these loops.

8.4.4 Premode Cleaner

The final key element of the PSL is the premode cleaner (PMC). The PMC acts as an interface between the PSL table and the main interferometer in vacuum. The PMC also acts to filter the spatial profile of the laser beam to deliver a beam of fundamental TEM00 mode to the interferometer. The PMC also acts as a passive filter of laser intensity and frequency noise and beam pointing noise [134].

The cavity is formed of four mirrors bonded to an aluminium spacer in a bow-tie configuration. The spacer is encased in a sealed container on the PSL table to reduce the coupling of acoustic and seismic noise. The cavity round trip length is 2.01 m and the finesse is 130 resulting in a cavity line width of 1.15 MHz.

The beam transmitted through the PMC is coupled to the main interferometer via the PSL output as described in the next section. The power reflected from the input port is used to lock the PMC and also as a monitor of the laser output power. The cavity is designed to leak a small amount of power through the other mirrors of the cavity, these beams are used by the ISS and FSS, with one beam coupled to the on-bench reference cavity and the other to the first ISS photodiode array.

8.4.5 PSL output

The output beam from the PSL is coupled into the vacuum system via a periscope on the PSL table and a pipe between the PSL clean room and the interferometer vacuum system.

The amount of laser power delivered from the PSL chamber to the vacuum system is set by a PBS/HWP power setting combination. The PBS reflects light of one linear polarisation and transmits light of the orthogonal linear polarisation. By rotating the HWP relative to the PBS, the orientation of polarisation of the light is rotated, meaning the ratio of power reflected to power transmitted also varies. In this way the power transmitted by the PBS can be controlled by rotating the HWP, the reflected power is dumped.

The power setting HWP on the PSL table is controlled by a stepper motor. In this way the optical power delivered from the PSL to the interferometer can be controlled and set to a desired level.

Chapter 9

A 70 W Laser Amplifier for the Third Observing Run of Advanced LIGO

Plans were made during the second observing run (O2) of the Advanced LIGO detectors to install a new laser amplifier at both detector sites. This amplifier would increase the laser power available to the interferometer, lowering the shot noise of the gravitational wave detector. The O2 observing run started in November 2016 and was completed in August 2017. The amplifier was installed in the LIGO Livingston laser facility as the first completed upgrade of the LIGO detectors between the second and third observing runs in September 2017.

This chapter presents the motivation for this upgrade to the LIGO laser system, and covers the work carried out by the author on-site to test and install the laser amplifier. This work was completed over two visits to the LIGO Livingston site as part of the LSC Fellows programme in collaboration with on-site scientists. The scope of work carried out by the author and included in this thesis includes the testing of the amplifier before installation, and the integration into the LIGO Livingston PSL which was completed in support of the LIGO Livingston laser team.

9.1 High Power Laser Options for Advanced LIGO

The design sensitivity of Advanced LIGO is a projection of the best operating sensitivity of the instrument and is contingent on a pre-stabilised laser (PSL) system that can deliver a laser beam of 120 W of optical power in the fundamental Gaussian mode to the interferometer.

Increasing the power injected into the interferometer increases the sensitivity of the instrument as the magnitude of the detected signal scales linearly with optical power and shot noise scales as the square root of power, so doubling the injected laser power improves the signal to noise ratio by a factor of $\sqrt{2}$.

The light sent to the interferometer from the PSL is required to have spatial mode profile of a fundamental order Gaussian. A description of a laser beam's spatial properties can be found in Section 1.4.3. The optical cavities that form the main interferometer of Advanced LIGO are designed so that at the operating point only light in the fundamental Gaussian mode is resonant. Light in higher-order modes is not resonant and therefore does not contribute to signal generation.

To filter the spatial beam properties, the seed beam is passed through a suspended triangular ring cavity called the input mode cleaner (IMC), in addition to the premode cleaner (PMC) which is part of the PSL. A mode cleaner cavity is designed such that the required mode is transmitted from a given output port with high efficiency, and higher-order modes are rejected. This reduces the power available to the interferometer but ensures maximal coupling of light to the optical cavities.

9.1.1 High Power Oscillator

The original design of the PSL included the high power oscillator (HPO) [114], an injection locked ring oscillator laser amplifier. This laser device was commissioned and installed during the upgrades of the initial LIGO instrument to the Advanced LIGO specifications.

During the running of Advanced LIGO issues with the HPO were uncovered relating to the water cooling system for the HPO and the resulting injection of beam-jitter noise [135].

As a resonant cavity with high circulating power, a significant amount of heat is dissipated in the HPO around the pump light injection and the amplifier crystals. To prevent damage the device is water cooled. This requires a constant circulation of water between the HPO box on the PSL table and an external chiller [129].

The flow of water though the piping system couples motion to the PSL table, introducing jitter, or pointing noise, to the laser beam. This jitter can couple to the output of the LIGO detector limiting the sensitivity. It was shown during testing of the HPO that the level of beam-jitter noise on the output of the HPO met the design requirement [103]. However, while the HPO at LIGO Hanford (LHO) was in use for the second observing (O2) run, the coupling of beam-jitter noise from the PSL table to the differential arm length readout was stronger

than expected. This produced significant noise in the gravitational wave signal readout channel that limited the sensitivity of the LHO detector during the second observing run.

9.1.2 Laser Amplifier for O3

For the first two observing runs of the advanced detector network both LIGO detectors were operated with around 25 W of laser power injected into the interferometer from the front-end laser, for O3 the plan was to use 50 W of laser power. This is more power than the front-end laser alone can produce, therefore an alternative to the HPO is needed.

The solution implemented for the third observing run (O3) of Advanced LIGO is a second solid-state laser amplifier in series with the front-end amplifier. This doubles the available laser power for O3, increasing the sensitivity of the detector at high frequencies by a factor of $\sqrt{2}$.

The modelled effect of doubling the laser power on the detector sensitivity is shown in Figure 9.1 along with the design sensitivity with 120 W of laser power available. Increasing the laser power reduces the quantum noise of the detector at the high frequency end of the detection bandwidth with a corresponding increase of quantum noise at frequencies below 50 Hz. However, at these frequencies the LIGO detectors are not currently limited by quantum noise, so for the scale of laser power increase discussed here the corresponding increase of lower-frequency quantum noise does not impact the detector sensitivity.

9.2 The neoVAN Laser Amplifier

The new high power laser amplifier for Advanced LIGO to meet the target of doubling the available laser power for O3, is a neoVAN-4S solid-state laser amplifier [136], pictured in Figure 9.2(a). The operating principle of this type of laser amplifier is shown in Figure 9.2(b) and described in the Section 8.2.2.

The neoVAN-4S includes four Nd:YVO_4 crystals in series, with each crystal pumped with light from an external diode laser. In the amplifier unit, the steering mirrors to direct the seed light, the focusing lenses for the pump light, and the laser crystals themselves, are all bonded to a common breadboard to form a quasi-monolithic system with fixed internal alignment.

The disadvantage of using multiple MOPAs is a degradation in the spatial mode quality of the output beam. The amplifier is capable of increasing the 35 W

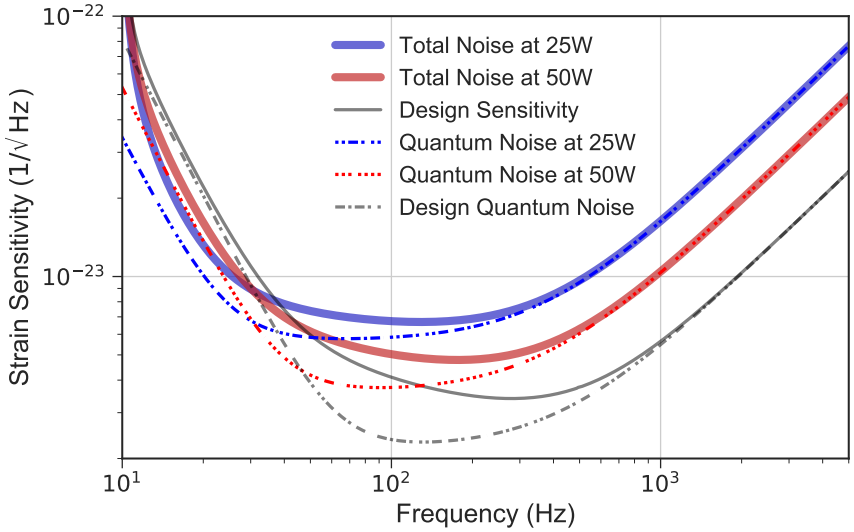


Figure 9.1: Predicted strain sensitivity of Advanced LIGO for different input laser powers. The strain sensitivity of the detectors is shown for 25 W and 50 W of laser power available to the interferometer. Also shown is the design sensitivity for a laser power of 120 W. The plot also shows the contribution of quantum noise to the projected total noise of the detector.

front-end beam up to an optical power of 80 W but typically only 50 W of this will be in the fundamental mode.

The laser amplifier was tested at the AEI in Hanover on a replica PSL table [137]. The amplifier was then installed in a facility called the 'Testing and Training' (TnT) lab at the LIGO Livingston site, to rehearse the integration of the amplifier into the optical table layout and prepare the infrastructure including water cooling, electrical wiring and control software.

9.3 Integration into the PSL System

The neoVAN amplifier unit has a small footprint and low technical complexity compared to the HPO and can be integrated into PSL system with only minor changes. The schematic of the PSL table, detailing the integration of the amplifier, is shown in Figure 9.3.

To seed the laser amplifier with the 35 W front-end laser beam, the beam path needs to be re-routed to bypass the HPO. This is done by replacing the mirror that directed the beam into the HPO box with a polarising beam-splitter (PBS). A half-wave plate (HWP) placed before this PBS is used to set the amount of power sent to HPO or to the new amplifier. This leaves open the possibility to

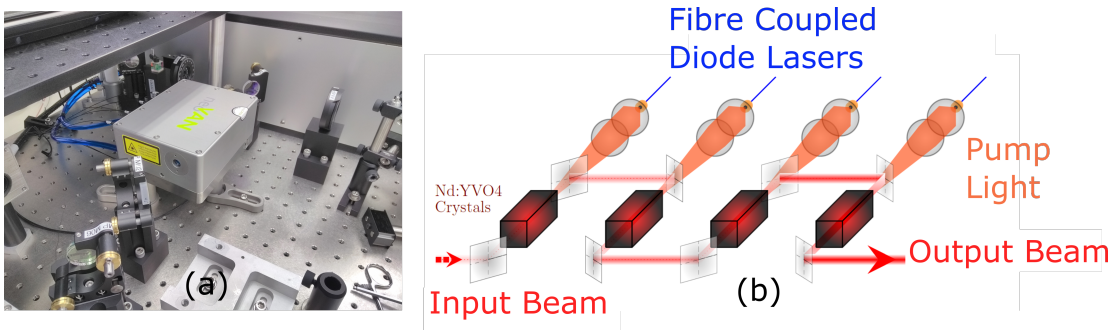


Figure 9.2: Laser amplifier for Advanced LIGO. Picture of the neoVAN laser amplifier on the PSL table (a). Diagram showing the operating principle (b). The four amplifier crystals are illuminated by four external diode lasers, fibre coupled to the rear of the amplifier box. The seed laser is directed though the four crystals with fixed steering mirrors and the power is amplified at each stage.

switch between high power laser options.

An additional high-power Faraday isolator is placed between the neoVAN amplifier and the front-end box, to prevent back-scattered light from propagating back into the front-end amplifier. This also ensures that the light entering the amplifier is of a high-purity linear polarisation, increasing the amplifier efficiency.

A shutter was installed after the neoVAN amplifier. This shutter can be controlled remotely to block off the laser output of the PSL, this is also linked to an automated safety watchdog. The shutter consists of a mirror on a fast-release motorised arm that directs the beam onto a water cooled, high power beam dump. A short focal length concave lens is used to spread the beam across a larger area of the dump.

9.4 Testing and Installation

9.4.1 Alignment to the Amplifier

The steering mirrors inside the neoVAN amplifier are fixed in place during manufacture, defining an internal optical path that the seed beam needs to be aligned to. The amplifier was initially aligned on the table so that the beam passed clearly though all four crystals, this was done with the amplifier off to prevent damage due to stray light being amplified.

To maximise the optical amplification, the seed light beam needs to be aligned to the volume of the gain medium that is illuminated by the pump light. If the seed beam is misaligned slightly relative to the pump light, the resulting amplifier beam will have a distorted beam shape and the amplifier output beam will no

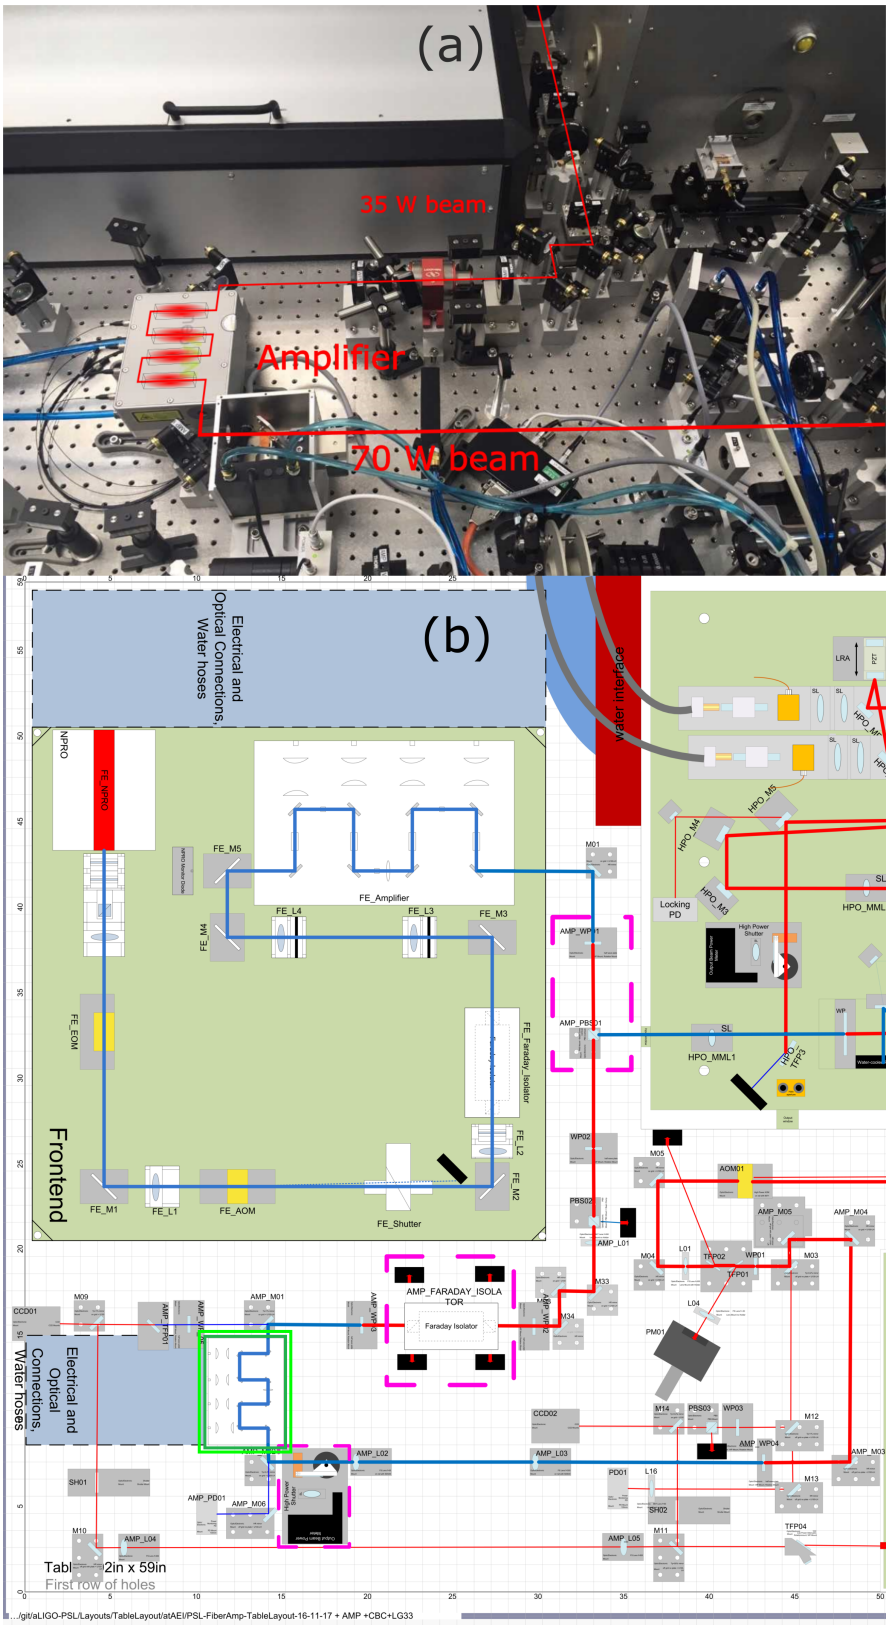


Figure 9.3: Photograph of amplifier unit in place on the LLO PSL table with the beam path overlaid (a). Schematic of the integration plan (b); this is a segment of the full detailed schematic meant to show the information relevant to this topic; additional detail that may not be legible can be found in the full schematic [138]. The amplifier is in the lower left, outlined in green. Highlighted in pink are the additional auxiliary elements also installed. From top to bottom these are the power setting HWP and PBS, Faraday isolator and high power safety shutter.

longer have a Gaussian profile.

There are five HR steering mirrors between the front-end output and the amplifier input window. However, the final two steering mirrors before the amplifier are either side of the Faraday isolator, the input and output apertures of which restricts the amount of steering freedom available without beam clipping.

The fine alignment of the seed beam to the amplifier was done with the amplifier powered and the input beam steering adjusted to maximise the optical amplification and beam shape.

9.4.2 Mode-matching to the Amplifier

To enable optimal amplification, the beam size of the seed light should match the volume of the gain medium that is illuminated by the pump light. This spatial overlap is known as the mode-matching of the seed beam to the amplifier.

The seed light and pump light are both Gaussian beams (see Section 1.4.3) and as such for each case the transverse beam profile is described by Equation 1.16, and the beam width evolves over the propagation direction as described by Equation 1.17.

To achieve good mode-matching of the seed beam, a mode-matching telescope consisting of two lenses is used in the path between the front-end amplifier and the neoVAN amplifier.

The initial placement and focal length values of these lenses can be modelled prior to installation from an accurate measurement of the beam profile out of the front-end laser and knowledge of the beam specifications of the amplifier, but to achieve good mode-matching this solution must be optimised with careful movement of the lenses on the table.

This was done initially while observing the transverse profile of the output beam on a camera, later this was redone with the beam aligned and locked to the premode cleaner (PMC) cavity. In this case the optical cavity acts as a mode reference as a beam in the fundamental Gaussian mode will be optimally resonant in the PMC. So the output of the amplifier can be jointly optimised for power amplification and beam shape by maximising the transmission through the PMC cavity.

9.4.3 Beam Quality of the Amplifier Output

When amplifying a laser beam with a MOPA laser amplifier, it is inevitable that some of the optical power is coupled into higher-order modes. With multiple

amplifier stages implemented in series, a significant percentage of the output power will be in higher-order modes.

The level of amplification and the mode content of the output beam depends on both the alignment and mode-matching of the seed light to the pump fields. In addition, if the seed light beam has some higher-order mode content, this will lead to an even greater higher-order mode content of the amplified beam. It is possible to have a strong optical amplification, with a poor beam quality.

If the seed and pump beams are not co-linear, the amplified beam will contain a fraction of first-order Hermite-Gaussian mode (see Figure 1.7). This is the effect of misalignment on beam quality.

If the seed beam size is larger than the pump beam in the gain medium, then the amplification will not be equal across the beam profile resulting in a distortion of the beam shape due to poor mode-matching.

In an ideal case with optical alignment and a seed beam of only fundamental mode content, the amplification process should result in a fundamental mode Gaussian beam output, resulting from the multiplication of two Gaussian profiles. However, there are reasons why the volume of the gain medium effectively illuminated by the pump light in not described by a Gaussian profile.

At the tails of the pump light Gaussian profile the pump light power can be lower than the threshold needed for laser amplification. This results in a clipping of the profile of the gain volume relative to the seed light profile and can result in a non-Gaussian amplified beam.

Similarly, at the centre of the pump light profile where the power density is highest, the gain medium may become saturated. In this case the amplified beam will also not have a perfect Gaussian profile.

These mode-matching effects can result in the presence of Laguerre-Gaussian modes in the amplified beam (see Figure 1.8).

The coupling of light from the fundamental mode to the higher-order modes can be minimised by careful tuning of the pump light power to the four amplifier stages to prevent distortion effects in the gain media. In the neoVAN-4S system the current that drives the pump light diodes and the temperature of the diodes can be set in the control software. The pump diodes are controlled in pairs allowing tuning freedom to optimise the amplifier output.

9.4.4 Alignment to the PMC

The PMC acts as an interface between the main interferometer and the PSL beam and can be used as an alignment reference. The PMC is the last major element

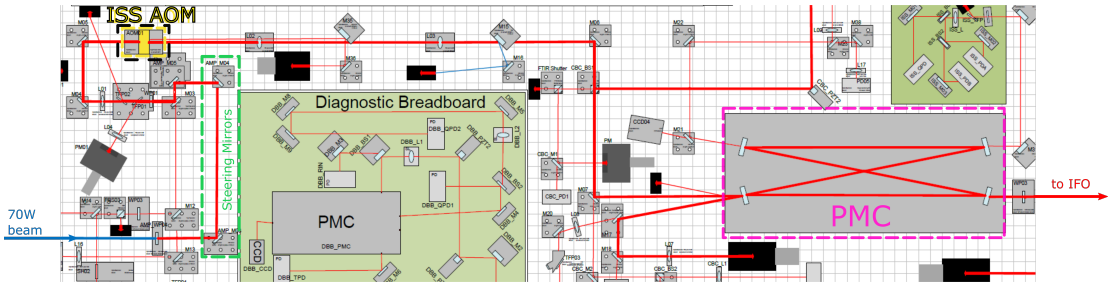


Figure 9.4: Section of the PSL table layout schematic showing the alignment path of the 70 W beam to the premode cleaner (PMC) outlined in pink. Also outlined is the ISS AOM (yellow) and the two steering mirrors (green) used to align the beam.

of the PSL as the output is injected into the main vacuum system. During the installation of the laser amplifier nothing 'downstream' of the PMC was altered. So if the beam out of the amplifier is aligned and mode-matched to the PMC, then by extension, the alignment to the main interferometer will have also been restored.

The optical layout between the neoVAN amplifier and PMC is shown in Figure 9.4.

The output of the laser amplifier must pass through an acousto-optic modulator (AOM) used as the actuator for the intensity stabilisation servo before being coupled to the PMC.

Mode-matching to an optical cavity is important to maximally couple the optical field into the cavity and ensure efficient resonance. Mode-matching was previously described in the case of the laser amplifier as the spatial overlap of two beam profiles. In the case of an optical cavity the curvature of the cavity mirrors defines a cavity eigenmode, an ideal beam shape that would be maximally resonant as the curvature of the cavity mirrors matches the phase front curvature of the eigenmode. Any difference between the input beam profile and the cavity eigenmode will result in a reduced coupling of the beam to the cavity on resonance.

The mode-matching of the existing beam path to the PMC is achieved with a combination of two lenses after the AOM. In the new scheme the amplified beam had to be mode-matched to this beam with an additional two lens telescope that must also focus the beam though the AOM.

With the amplified beam coupled to the PMC with sufficient fringe depth to lock the cavity, the mode-matching and fine alignment could be improved by maximising the transmitted power. With this method a transmitted power percentage of 80% for a 12 W input beam was achieved. Later, the power to the PMC was turned up to the full output power of the amplifier and a slightly lower

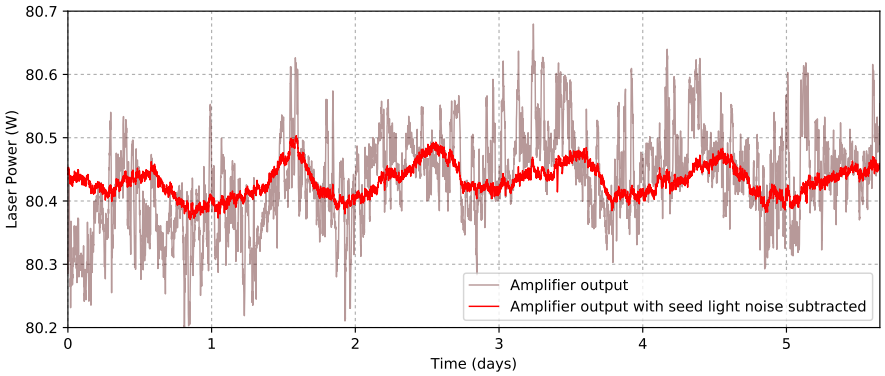


Figure 9.5: Output power of the laser amplifier as tested in the TnT lab. The output of the amplifier was measured by a power meter. The raw data output over the 6 days is shown in the brown trace. The seed laser power was also monitored and this data was scaled to the level of the amplifier output and subtracted from the measured amplifier output; the result is shown in the red trace.

transmission through the PMC was achieved to deliver 50 W of optical power to the main interferometer.

9.4.5 Power Stability

The laser module was tested in a replica PSL laser lab prior to installation in the main LLO PSL. This testing included a measurement of the long-term power stability of the amplifier. The result is shown in Figure 9.5 for a 6-day period when the amplifier was run without interruption.

In this result, the coupling of power fluctuations from the seed NPRO laser to the amplifier output could be clearly be seen. By scaling the measured power output of the NPRO and subtracting from the amplifier output the power drifts introduced by the amplifier alone can be seen. The NPRO in the TnT lab exhibited higher power fluctuations over long timescales when compared to the NPRO laser in the main PSL.

After Installation into LIGO Livingston PSL

The neoLASE amplifier was installed in the PSL over the course of three weeks in September 2017. During the installation process, the power stability was measured over a weekend between the second and third week of installation work. The output of the laser amplifier was deflected onto a water-cooled power meter and recorded over 60 hours. The result is shown in Figure 9.6.

The output power changed repeatedly by about ± 0.1 W over an hourly cycle,

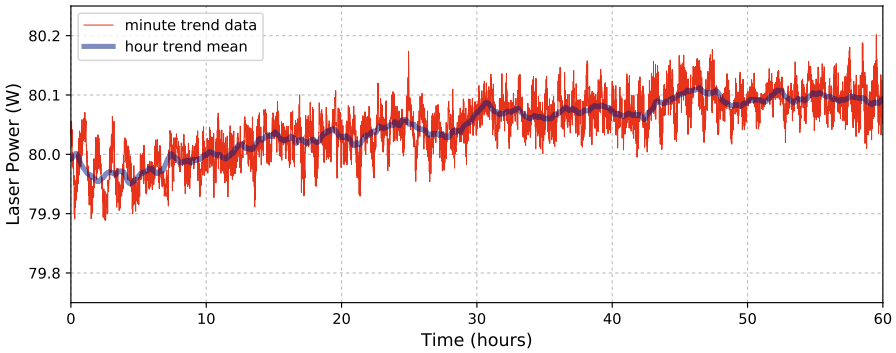


Figure 9.6: Output power of the laser amplifier after installation in the LLO PSL. The power was measured over a weekend with no disturbance to the PSL room. The mean of the minute trend of the power recorded by a power meter is shown in red, with the mean hour trend in blue. The amplifier output rose slightly over the first day and then remained steady just below 80.1 W.

this is related to the thermal cycle of the PSL enclosure due to the environmental controls. There was also a longer term upward drift in the power of about 0.15 W over the 60 hour period. This was with no active intensity stabilisation applied to the laser and the magnitude of the power drifts was within acceptable limits.

9.5 Detector Sensitivity at Higher Laser Power

Figure 9.7 shows the improvement in the sensitivity of the LIGO Livingston detector between the operating sensitivity during O2 and the sensitivity measured before the start of O3 with 50 W injected into the interferometer.

Increasing the power inside the detector can make the locking of the arm cavities more difficult due to the opto-mechanical interaction between the laser and the suspended mirrors. When a suspended cavity is moved onto resonance the increase in optical power circulating inside the cavity provides a radiation pressure 'kick' to the cavity mirrors that pushes the cavity away from resonance. If the cavity locking control is too slow or lacks the dynamic range, it will not be able to hold the cavity on resonance.

During the lock acquisition process the input laser power is slowly ramped up to the full power to reduce the magnitude of the radiation pressure kicks and thus allow time for the cavity control loops to damp the opto-mechanically driven motion of the cavity mirrors.

The optical power inside the interferometer heats the mirrors due to optical absorption in the mirror coating. The heating has the effect of distorting the

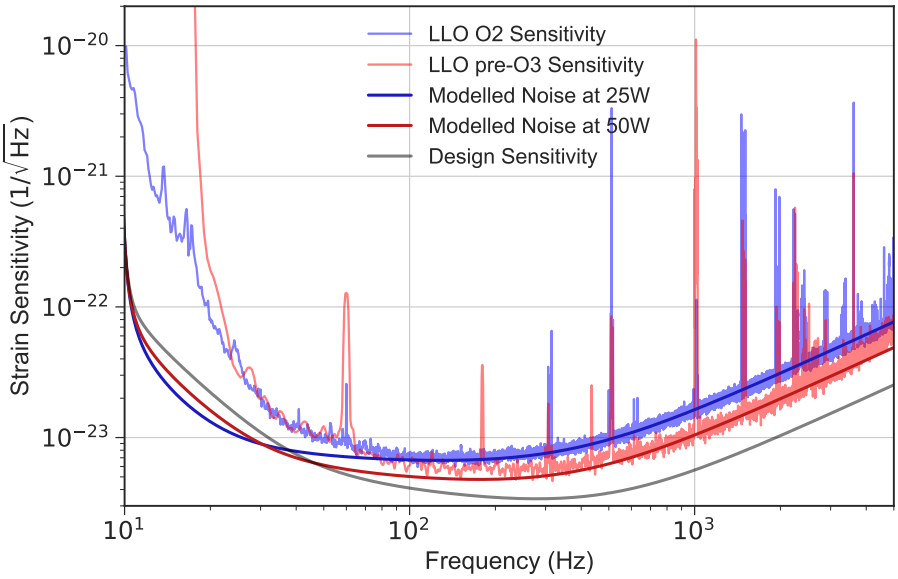


Figure 9.7: Sensitivity of the Advanced LIGO detectors with different injected laser powers. The best sensitivity of the LLO detector during the second observing run is shown in the faint blue trace. At this time the laser power injected into the interferometer was 25 W, the modelled noise limit to the strain sensitivity for this amount of laser power is shown in the darker blue trace. Similarly the strain sensitivity of LLO before the start of O3 when the laser system was providing 50 W to the detector and the modelled noise are shown in red. The design sensitivity of Advanced LIGO is shown in grey, this is projected noise limit with 125 W of laser power.

physical shape of the mirrors. This can change the mode-matching of the cavities as the radius of curvature of the mirrors is altered and changes the internal modes of the mirror which can lead to instability in the cavity locking.

The target for O3 was for the PSL to provide 50 W of optical power to the main interferometer. The inclusion of the neoVAN amplifier into the PSL system made this possible. However, during operation of the full interferometer it was found to be easier to lock and achieve long-term low-noise operation with only 40 W of laser power injected into the interferometer and it was in this state that the LIGO detectors operated for the start of the third observing run of the advanced detector network.

Chapter 10

Laser Stabilisation for Advanced LIGO during the Second Observing Run

An overview of the LIGO laser stabilisation scheme was presented in Section 8.4. In this chapter the first stages of the intensity and frequency stabilisation systems are described in more detail as background for two investigations of stability issues during the second observing run (O2) of the Advanced LIGO Livingston detector.

Sections 10.1 - 10.3 cover an investigation into the high frequency performance of the frequency stabilisation system electronics. This includes a study of the existing hardware and circuit modelling.

In Sections 10.4 -10.7 an investigation into lock losses during the operation of the LIGO Livingston detector during O2 relating to the laser intensity stabilisation system is presented.

The work in this chapter was carried out during a four month visit to the LIGO Livingston site under the LSC Fellows programme with the assistance of on-site scientists.

10.1 Advanced LIGO Table-top Frequency Stabilisation Scheme

The innermost loop of the frequency stabilisation loop of the Advanced LIGO laser system is contained fully within the PSL laser enclosure and provides suppression of laser frequency noise. This loop is referred to as the 'table-top' frequency stabilisation servo (TTFSS). A sketch of this system is shown in Figure 10.1.

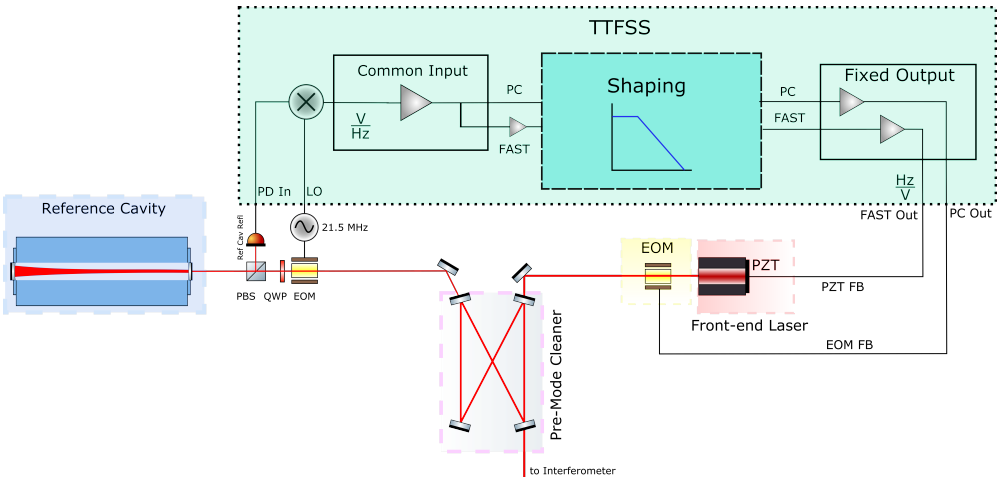


Figure 10.1: Schematic of the Advanced LIGO table-top frequency stabilisation scheme. The main laser beam to the interferometer passes though the premode cleaner (pink box) cavity, a small fraction of the light is transmitted by one of the cavity mirrors and is directed to a linear reference cavity (blue box). The returning beam is measured with a photodiode, the AC coupled photodiode signal is mixed with the local oscillator signal also used to generate phase sidebands on the light entering the cavity. The demodulation is part of the table-top frequency stabilisation servo box (green box). The error signal is split between two feedback paths with different filtering. The 'FAST' feedback is applied to a piezo-electric, PZT, inside the seed laser. The phase compensation (PC) feedback is applied to an electro-optic modulator (EOM).

The plot in Figure 8.9 shows the noise suppression requirement of the TTFSS, this shows the required frequency noise level of the laser at the input to the interferometer. The laser noise is passively filtered by the main interferometer. The noise requirement at the input to the interferometer rises with frequency as shown in the plot in Figure 8.9. This suggests that above 20 kHz the free running noise will cross with the noise requirement level and no further stabilisation is needed at higher frequencies. While this is the case to negate the coupling of laser frequency noise to the readout of the detector, the free running noise of the laser needs to be reduced at frequencies up to and above 20 kHz to enable the locking of the input mode cleaner and the main interferometer cavities. To achieve a low-noise full interferometer lock a TTFSS bandwidth between 500 kHz and 1 MHz is needed.

To achieve a wide control bandwidth the TTFSS servo uses analogue electronic filtering. The TTFSS servo box contains the signal processing electronics for the locking loop, including the RF mixer for beating the signal from the reference cavity with the modulation sidebands to produce the Pound-Drever-Hall (PDH) error signal. The error signal then passes through a common path gain stage with a variable gain amplifier that can be controlled externally. The error signal is then split between the feedback paths to the laser temperature (SLOW path), the laser PZT (FAST path) and the external phase compensation EOM (PC path) [132, 139].

The three feedback paths have different filtering to shape the frequency response of the error signal to produce stable crossings between the feedback paths and a stable unity gain point. The servo also must have enough gain at lower frequencies to suppress the free-running noise and stabilise the laser to the required level.

The fast feedback path contains a second externally controlled variable gain amplifier, with which the crossover between fast and PC feedback paths can be controlled.

10.1.1 Modifications to the TTFSS

The current version of the TTFSS servo box has been in use since the days of initial LIGO. The circuit was modified as part of the Advanced LIGO upgrade to compensate for changes made to the PSL.

In the initial LIGO PSL the premode cleaner (PMC) cavity was outside of the TTFSS loop; in the Advanced LIGO scheme the sensing point of the TTFSS was changed to be in transmission of the PMC. The linewidth of the PMC cavity is

570 kHz and the addition of this cavity pole to the optical portion of the control loop means the loop gain drops off as $1/f$ above the cavity pole and the phase decreases by 90°.

To compensate for this change, a daughter circuit board was added to the output stage of the PC path inside the TTFSS servo box (a simplified circuit diagram of the PC path is shown in Figure 10.3). On this daughter board was a single phase-lead filter. The circuit design of this daughter board as part of the PC path is shown in the white box in Figure 10.3. The purpose of this board was to compensate for the PMC pole in the new optical path.

10.2 Investigating High Frequency Stability of the TTFSS

The measured open loop response of the table-top frequency stabilisation servo has a 'bulge' in the gain between 5-7 MHz. The presence of this feature, which is characteristic of a unity gain instability, was not greatly affecting the performance of the TTFSS in relation to the operation of Advanced LIGO but the presence of the bulge reduces the gain margin of the open loop response to less than 3 dB, in turn meaning the unity gain frequency cannot be increased to be more than 500 kHz. This concern initiated an investigation of the TTFSS that adopted two strategies. The first step was to use existing measurements of the TTFSS open loop gain to investigate the presence of this feature.

The plot in Figure 10.2 shows transfer function measurements collected over 4 years of Advanced LIGO operations. The gain bulge is a clear feature of all measurements dating back to 2014.

The PC path was the main element investigated as the feature of interest was above 5 MHz, where only the high frequency PC circuit could be contributing to loop stability. A simplified circuit diagram of the PC path of the TTFSS is shown in Figure 10.3.

10.2.1 Modelling the TTFSS

A model of the TTFSS circuitry was produced in LISO [140,141], a linear element analysis program that can simulate the response of a circuit using a library of known op-amp performance parameters.

The modelling of the TTFSS suggested that the gain bulge above 5 MHz is related to the output stage of the PC path. The output stage includes two parallel

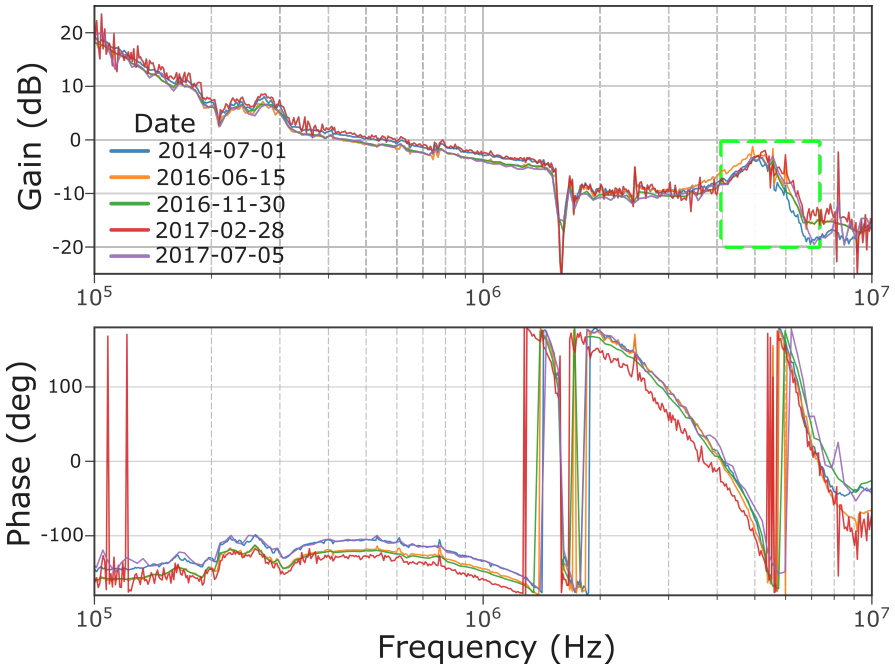


Figure 10.2: Open loop gain measurements of LIGO table-top frequency stabilisation servo. Collection of measurements taken between 2014 and 2017. The bulge in the open loop gain above 5 MHz was the focus of this investigation.

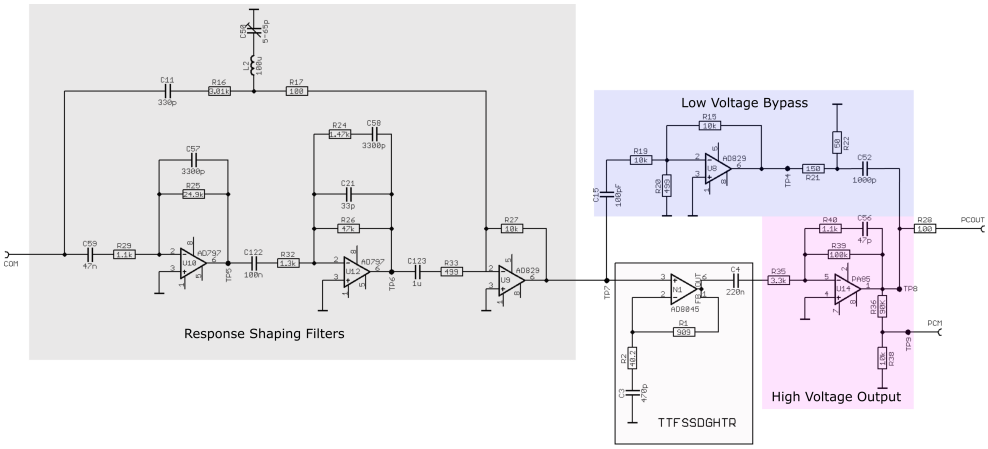


Figure 10.3: TTFSS PC path. Dark grey block contains the shaping stage before the output stage consisting of daughter board (light grey) high voltage driver (pink) low voltage bypass (purple).

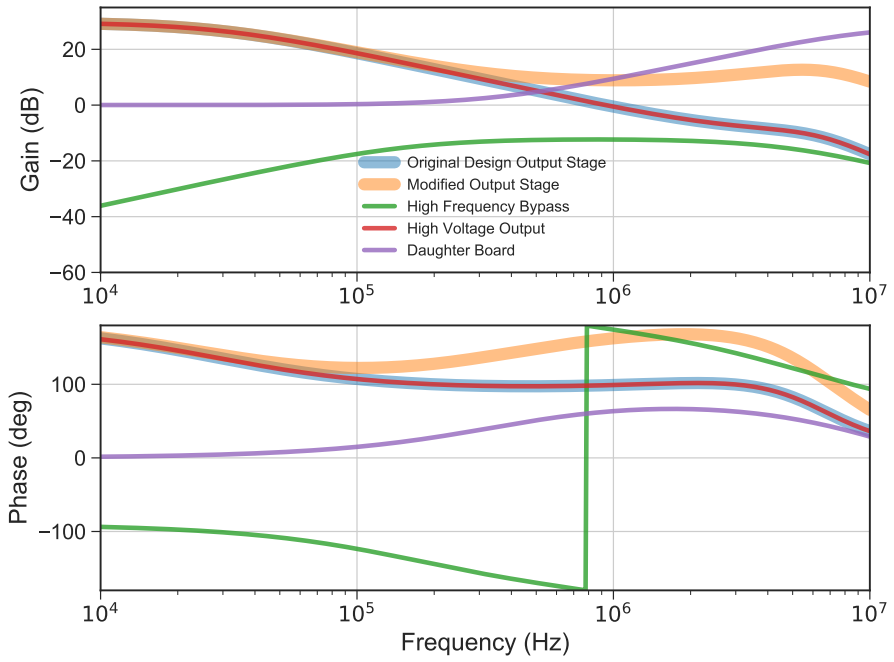


Figure 10.4: Transfer function breakdown of PC path output stage. Solid blue and orange trace show the modelled response of the PC path with and without the low voltage bypass respectively. The dashed traces show the component parts. Low voltage bypass (green), PA85 high voltage driver (purple) and the combination of PA85 amplifier and daughter board (red).

drivers, the so called *high voltage driver* and the *high frequency bypass*.

The high voltage driver path, highlighted in pink in Figure 10.3, uses an Apex PA85 [142] op-amp capable of driving large voltages (± 200 V) onto the EOM actuator. The high voltage range of this op-amp comes at the cost of limited high frequency performance. To compensate for this, a parallel stage called the high frequency bypass, highlighted in blue in Figure 10.3, makes use of the better high frequency performance of an Analog Devices AD829 [143] op-amp to still allow stable feedback up to high frequencies.

Every op-amp has a finite bandwidth over which they can operate. Above some frequency the gain and phase both drop off and will eventually pass through a unity gain point. In a control loop if the phase is $-180°$ at a frequency where the gain of an op-amp crosses unity, then the op-amp response will become unstable and the op-amp output will oscillate. If the phase is close to $-180°$ but not yet unstable, a characteristic gain bulge may be seen in the response of the amplifier. This can lead to loop instabilities as the gain margin is reduced and the response of the op-amp close to instability becomes non-linear.

The feature observed in the TTFSS open loop gain resembles the characteristic unity gain instability of an op-amp in a feedback loop close to instability. Modelling of the TTFSS circuit predicts that the high voltage drive stage of the PC path will exhibit such an instability above 5 MHz. The results of modelling the output stage is shown in Figure 10.4.

The high frequency bypass response modelled from the circuit design is shown in the green trace in Figure 10.4, The gain of the high frequency bypass never crosses with the high voltage drive stage, suggesting the gain of the bypass stage was originally set too low. If the gain of the bypass stage was at least 10 dB higher the AD829 would provide feedback gain above 1 MHz meaning the instability of the PA85 above 5 MHz would be compensated for by the alternative driver.

The use of the AD829 op-amp in this circuit is also problematic. The AD829 chip is a good high frequency op-amp but is not unity gain stable and is not recommended to be used in circuits with low gain without external compensation capacitors.

As can be seen from the circuit diagram in Figure 10.3 the daughter board is placed in the path of the high voltage driver as part of the output stage of the PC path. The addition of this board unbalances the two parallel output stages. In terms of the servo response the effect of this modification can be seen in Figure 10.4. In this plot the purple transfer function shows the response of the daughter board to compensate for the PMC cavity pole. The thick traces show the transfer function of the overall output stage of the PC feedback path before (blue) and after (orange) the daughter board is added. The orange trace, which include the daughter board as part of the output stage, shows that high frequency bypass could not crossover with the high voltage drive stage to compensate for the instability above 5 MHz.

10.2.2 Previous Investigation of TTFSS Instability

An instability in the PC output was first noticed during commissioning of the LLO PSL for Advanced LIGO in 2013. The AD829 high frequency bypass stage was found to be oscillating. To address this, the bypass was removed, this prevented the output stage from oscillating but meant that the unity gain bulge of the PA-85 op-amp was not compensated at all by the bypass.

It is only since the low-voltage bypass was disconnected that the bulge in the FSS loop gain has been seen. Measurements of the FSS open loop gain before and after the bypass was removed show that the gain bulge consistently shows up in all measurements since the change as can be seen in Figure 10.2.

10.3 Conclusion of Investigation into TTFSS Instability

The conclusion of the investigation into the TTFSS was that a gain bulge in the loop response was a result of the low voltage bypass from the EOM driver stage of the servo electronics not properly compensating for the instability of the high voltage driver.

The low voltage bypass was disconnected due to a more obvious instability arising from the inclusion of the daughter board to the TTFSS, a modification made necessary by the inclusion of the PMC in the FSS loop as part of the Advanced LIGO PSL.

The removal of the bypass allowed the TTFSS to be used to stabilise the frequency of the laser to below the noise requirement for Advanced LIGO and was used in this configuration for the highly successful first and second observing runs of Advanced LIGO. However, this issue did limit the bandwidth of the TTFSS loop and could be the cause of some instability in the frequency stabilisation loop.

A problem relating to the stability of one of the many control loops of the LIGO detector can reduce the operating up-time of the detector. For the detector to be in low-noise observing mode, multiple control loops have to be locked to hold both of the 4 km arm cavities, the power recycling, signal recycling and inner Michelson cavities at the control set points. This also requires the three mode-cleaning and reference cavities to be locked. The coupling of signals between the multiple nested control loops of the instrument is very complicated and not completely understood. An instability in one control loop can affect the locking of the whole interferometer and if one loop falls out of lock it causes the full interferometer to fall out of lock.

The simplest solution for the TTFSS issue would be to reconnect the high frequency bypass with the correct gain, and to find a way to prevent the additional daughter board from making the output stage unstable. The daughter board is necessary so cannot be removed, but it could be moved to a stage earlier in the PC path so that it does not adversely affect the output stage. As a shaping filter it should be located before the output driver stage.

The conclusion of this investigation has been presented to the scientists at the LLO detector site and has been considered as a future solution to be implemented.

10.4 Advanced LIGO Inner-Loop Intensity Stabilisation Servo

The scheme for stabilising the intensity of the laser light that seeds the main Advanced LIGO interferometer is described in Section 8.4.2. The following sections focus on the inner loop of the intensity stabilisation system (ISS) and an investigation into how the stability of this system affected the performance of the Livingston detector during the second observing run of the advanced gravitation wave detector network.

The inner loop of the ISS senses a small fraction of the laser power picked off from a partially transmissive cavity mirror of the premode cleaner. The power fluctuation of the light field is sensed and an error signal generated and processed by an analogue servo with unity gain around 50 kHz. The feedback signal actuates the intensity of laser beam with an AOM actuator, to correct out for the detected fluctuations in the optical power and thus stabilise the intensity of the laser. The actuation signal is added to a reference RF oscillation signal that is applied to the AOM. The reference signal diffracts power out of the frundamental mode and the actuation signal modulates this RF signal to modulate the amount of power diffracted.

To operate as a feedback actuator, the AOM must have enough range to compensate for the detected fluctuations. An AOM can only decrease the power in the main beam path by diffracting power into geometrically separated beams, so the ISS must operate from a set point were an offset voltage is applied to the AOM so that servo has the range to compensate for the size of the fluctuations. This offset voltage is typically set to couple between 3% and 10% of the optical power into the first diffraction order. This reduces the amount of power available to the interferometer. The diffracted beam is dumped onto a power meter that records the power diffracted by the AOM.

10.5 Lock Losses Due to ISS Instability

The detectors are run in the low noise observation states for long stretches of time to maximise the likelihood of observing a gravitational wave event. The aim is to be observing for 24 hours a day with only scheduled down time for maintenance and commissioning.

In the low noise operation state the full dual-recycled Michelson interferometer is locked and left in an undisturbed state. Many things can occur that disturb

the state of the interferometer and can cause the detector to lose lock. The cause of these losses of the locked state can often be easily attributed to some external factor such as excess environmental noise, high wind or seismic activity. However instabilities in the control loops or glitches in the data recording system are also common causes of lock loss.

Often it is less clear what caused the interferometer to fall out of the low noise observing state. One way to investigate the cause of a lock loss is to look at the recorded control and auxiliary data channels for the period of time immediately prior to a lock loss and search for any unusual oscillations or saturations. However, when the detector falls out of lock, saturations are seen in many channels, so precise timing is required to determine if a saturation could be a potential cause or an obvious consequence of a lock loss event.

The Advanced LIGO detector infrastructure includes over 200,000 channels of recorded data. The coupling between these channels is complex and not fully understood. But this method of investigating the lock losses can often point towards the problem or the problem area.

10.5.1 ISS Dropouts and Lock Loss

During the early stages of O2, a number of lock losses were recorded at LIGO Livingston that where coincident with unusual behaviour of the intensity stabilisation servo (ISS) subsystem of the pre-stabilised laser (PSL). It was noticed that immediately before a lock loss the power diffracted by the ISS AOM dropped suddenly. As the power diffracted by the AOM drops towards zero, the ISS loop will run out of actuation range and the loop will saturate.

Throughout January and February 2017, several lock losses occurred and were coincident with dropouts in the AOM diffracted power. The rate of these dropouts also seemed to be increasing. The dropouts were thought to be caused by either an instability in the ISS feedback loop or a lack of actuator range.

As a solution, the gain of the stabilisation loop was decreased to increase the phase margin and therefore the stability of the loop, and the diffracted power was increased to provide more range for the loop actuator. The downside of decreasing the gain is a reduction of the noise stabilisation level and the cost of increasing the diffracted power is a slight reduction in the power coupled into the interferometer.

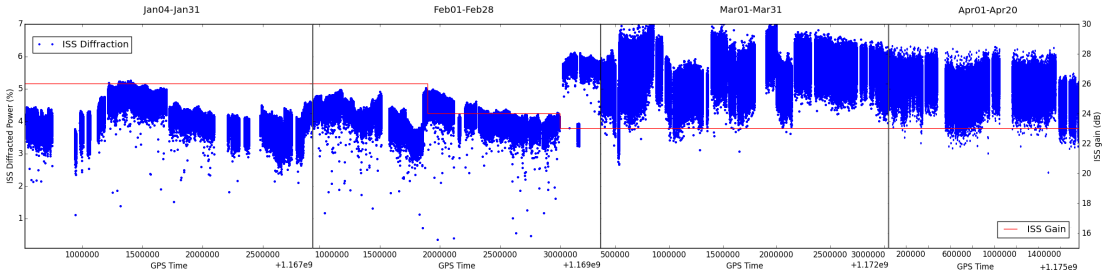


Figure 10.5: ISS diffracted power during observation times between 04-01-2017 and 20-04-18, showing the ISS loop gain on the right axis (red). The ISS loop gain was lowered from 26 dB to 24 dB on the 11-02-2018 and again to 23 dB on the 24-02-218 when the diffracted power was then also increased to around 6%. On the 21-04-2018 the diffracted power was lowered again to around 4%.

10.6 Investigation into ISS Dropouts

To better understand the issue, it was necessary to look back at the ISS performance and more closely characterise the timing of the dropouts. It was important to understand how many of the dropouts actually occurred before the interferometer fell out of lock and may therefore have caused the lock loss.

10.6.1 Timing of ISS Dropouts

The ISS diffracted power between the 4th of January and the 20th of April was trended and plotted as the ISS diffracted power during observation shown in Figure 10.5. From this it can be seen how the changes made to the diffracted power level and the ISS loop gain reduced the number of dropouts of the ISS.

After the loop gain was reduced there was a significant reduction of sudden changes in the ISS diffracted power even during lock. This suggest that the problem was not due to a lack of range in the diffracted power but was solely related the stability of the loop. If the issue was due to a lack of range for the AOM diffraction, then increasing the diffracted power should decrease the number of dropouts causing lock loss but not reduce the total number of drops in the diffracted power. This suggested strongly that the diffracted power could be reduced again to allow more optical power to the interferometer.

The power diffracted by the AOM was reduced from roughly 6% to 4% on the 21st of April. The initial result seemed promising, as there was no return of the dropouts in diffracted power causing lock loss in the following weeks and this continued through April and May. In June a number of ISS dropouts were again found to occur immediately before the interferometer fell out of the locked state.

The plot in Figure 10.6 shows the cumulative sum of lock losses during O2

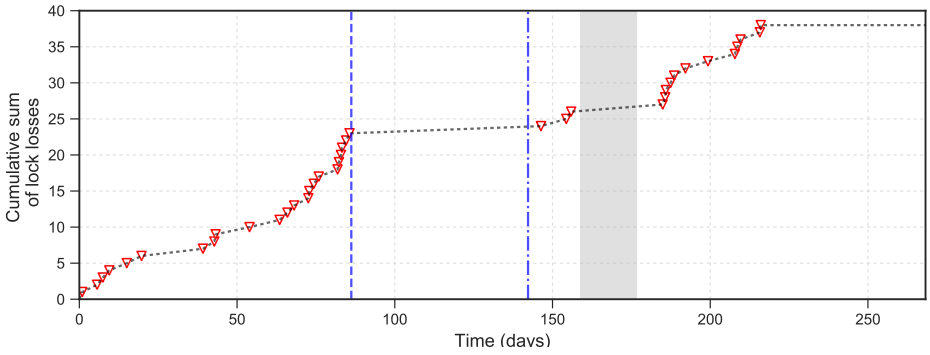


Figure 10.6: Cumulative sum of lock losses that could have been caused by a dropout in the ISS diffracted power for the duration of O2. The time of each lock loss is indicated with a red diamond. The grey shaded section marks the time of the vacuum intervention during the observing run when the Livingston detector was out of observing mode for an extended length of time. The blue vertical lines indicate changes made to the ISS system. The dashed line on the 24-02-2018 is when the diffracted power was increased to 6% and the loop gain was decreased to 23 dB. The second, dot-dash, line shows when the diffracted power was decreased back down to 4%.

that can be attributed to dropouts in the power diffracted from the ISS AOM. The grey region on this plot marks a break in the observing run that was taken to preform detector maintenance. The changes to the diffracted power and loop gain are marked by the first vertical blue line, after this time no lock losses were seen in the following 50 days. The second blue line marks when the AOM diffraction offset was reduced, this is proceeded by the return of the problem albeit not immediately following the change.

It is not clear what changes made to the system could have resulted in the sudden reduction in the number of lock losses after the 5th of July. No changes were made the ISS around this time. Although changes were made to the PSL table but only to the alignment of the FSS reference cavity.

10.6.2 ISS Loop Gain

The dropouts in the diffracted power of the AOM are related to the stability of the stabilisation loop. If the loop becomes unstable, large signals will be fed back onto the AOM driver, beyond the range where it can linearly actuate on the laser, causing a saturation in the control loop. This in-loop saturation or the sudden increase in laser noise, is a likely cause of an interferometer lock loss.

The servo filter for the inner loop of the ISS is an analogue filter with a overall gain that can be set using the LIGO digital control system, CDS. The digital

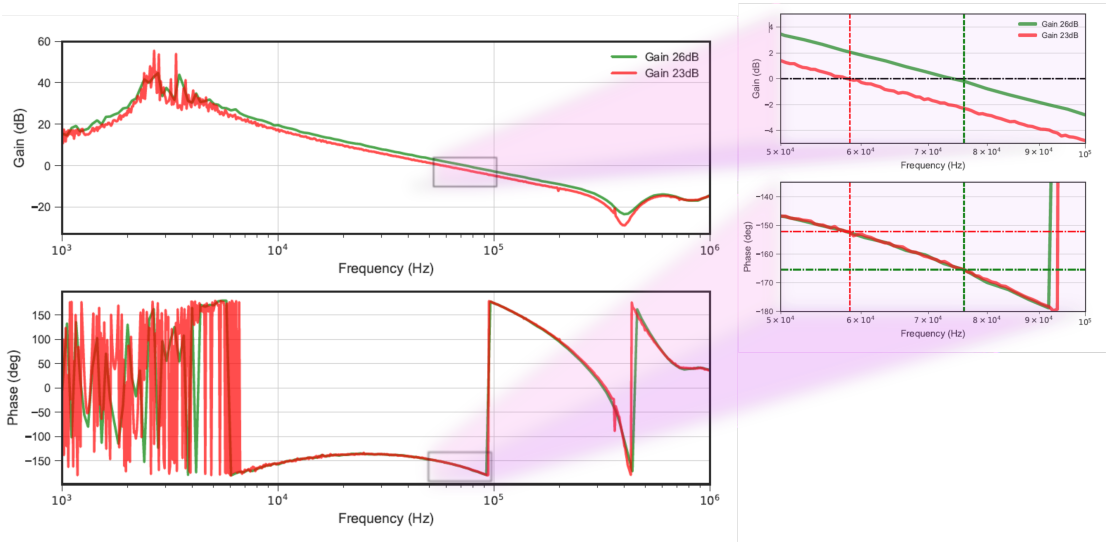


Figure 10.7: Open loop gain of the inner ISS loop measured with gain settings of 23 dB and 26 dB, the insert shows the unity gain crossing at 58.5 kHz and 75.9 kHz for the two gain setting respectively. The phase margin of the loop with the gain at 26 dB is only 14.5° which is close to instability, with the lower gain the phase margin is improved to 27.8°. These measurements were taken by commissioners at the LLO site.

control system interacts with the analogue servo via a variable gain amplifier, the gain is set by an external voltage from CDS. The overall gain factor sets the unity gain point of the servo loop and determines the amount of noise suppression implemented.

Examples of the response of the inner ISS loop are shown in Figure 10.7. The two transfer function measurements shown in this figure were measured in January and February of 2017 during the investigation of the ISS lock losses. The insert shows a zoomed in plot around the unity gain point. With the loop gain set to 26 dB, as it was for the start of O2, the phase margin is only 14.5°, this is close enough to instability to cause problems for the long term operation of the system as a change of -2 dB in the gain of the loop would cause the loop to become unstable.

Reducing the unity gain from 76 kHz to 59 kHz by dropping the loop gain to 23 dB increases the phase margin to a more stable 27.8° and doubled the gain margin to 4 dB. This is at the expense of noise suppression bandwidth.

It should be noted that the values of total loop gain quoted before as 23 dB and 26 dB are the values set with the amplifier stage but are not well calibrated as changing the servo gain nominally by 3 dB only changed the open loop gain by a factor closer to 2 dB

The investigation of the ISS loop points to stability of the unity gain crossing

being an important issue with regards to the impact of the ISS on the operation of the interferometer.

The loop stability was found to change over the course of an extended observing run as the gain of the loop drifts. The alignment of the beam to the AOM crystal affects the diffraction efficiency, and therefore the strength of the actuator, which in turn changes the gain of the loop. Alignment of the beam on the PSL table does drift very slightly over time and so the alignment to the AOM needs to be tuned up every few months during operation.

At the end of O2 during commissioning of the new laser amplifier, the ISS loop gain had to be set to 27 dB to recover a unity gain of 35 kHz further suggesting that there is a slowly changing gain element affecting the inner ISS loop.

10.6.3 Lock Loss Tool

To monitor and to help understand this issue, a data monitoring tool was developed to log the occurrence of dropouts in the ISS inner loop and the relation to interferometer lock loss. This tool was used to compile the full list of lock losses shown in Figure 10.6.

The ISS diffracted power is measured by a power-meter on the PSL table and the read-out is recorded in the CDS system at a rate of 16 kHz.

To efficiently search this data the minute trends for a block of time are analysed to identify potential dropouts. Once identified, the raw data of the interferometer state and the ISS diffracted power is pulled for these times segments. The lock losses occur rapidly so only in the raw data can it be seen if an ISS dropout occurred before or after a lock loss.

From the raw data, plots are generated of each dropout and occasions when the dropouts could be the cause of a full interferometer lock loss are flagged. Example of the plots produced can be seen in Figures 10.8 and 10.9. In these plots each subplot records the time of a dropout; these times are also written to a text file for reference. The occasions where it is likely that the ISS dropout caused the lock loss are highlighted in red. The time series of the diffracted power from the AOM are plotted in green while the interferometer is locked and in observing mode and plotted in red otherwise.

This code can be used as a monitor of ISS dropouts and to build up a record of occurrences as a diagnostic tool for the intensity stabilisation system to prevent this issue affecting the operation of the LIGO detector. These plots were generated for the full length of the second observing run of the LIGO Livingston detector.

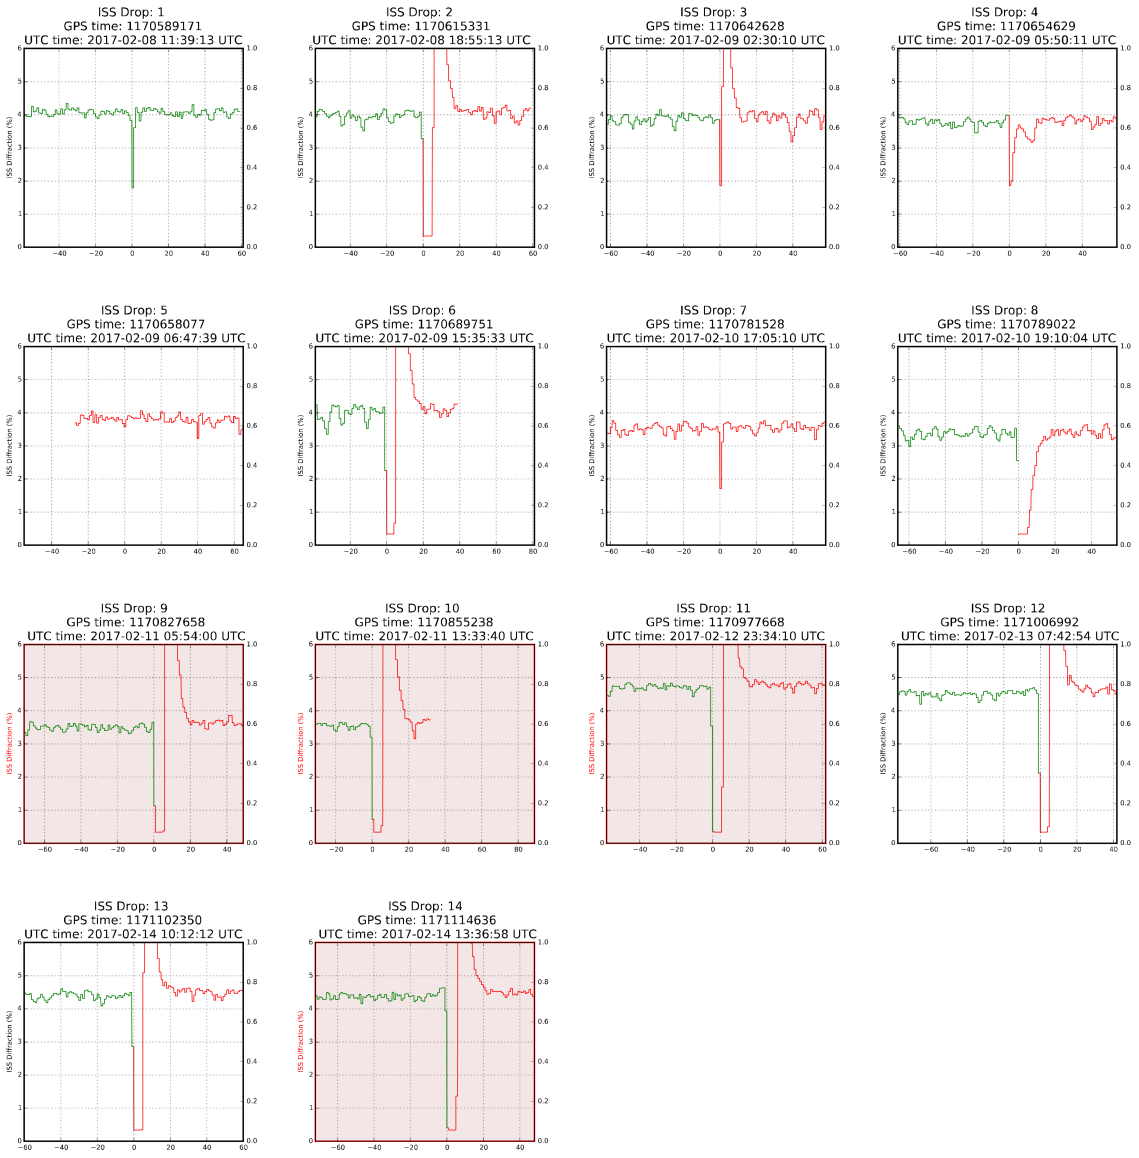


Figure 10.8: Catalogue of ISS dropouts during 7 days in February 2017. During this time 15 dropouts were recorded, of these four occurred milliseconds before the interferometer fell out of lock. On another 7 occasions the diffracted power dropped during a lock loss and on one occasion the interferometer remained locked during the dropout.

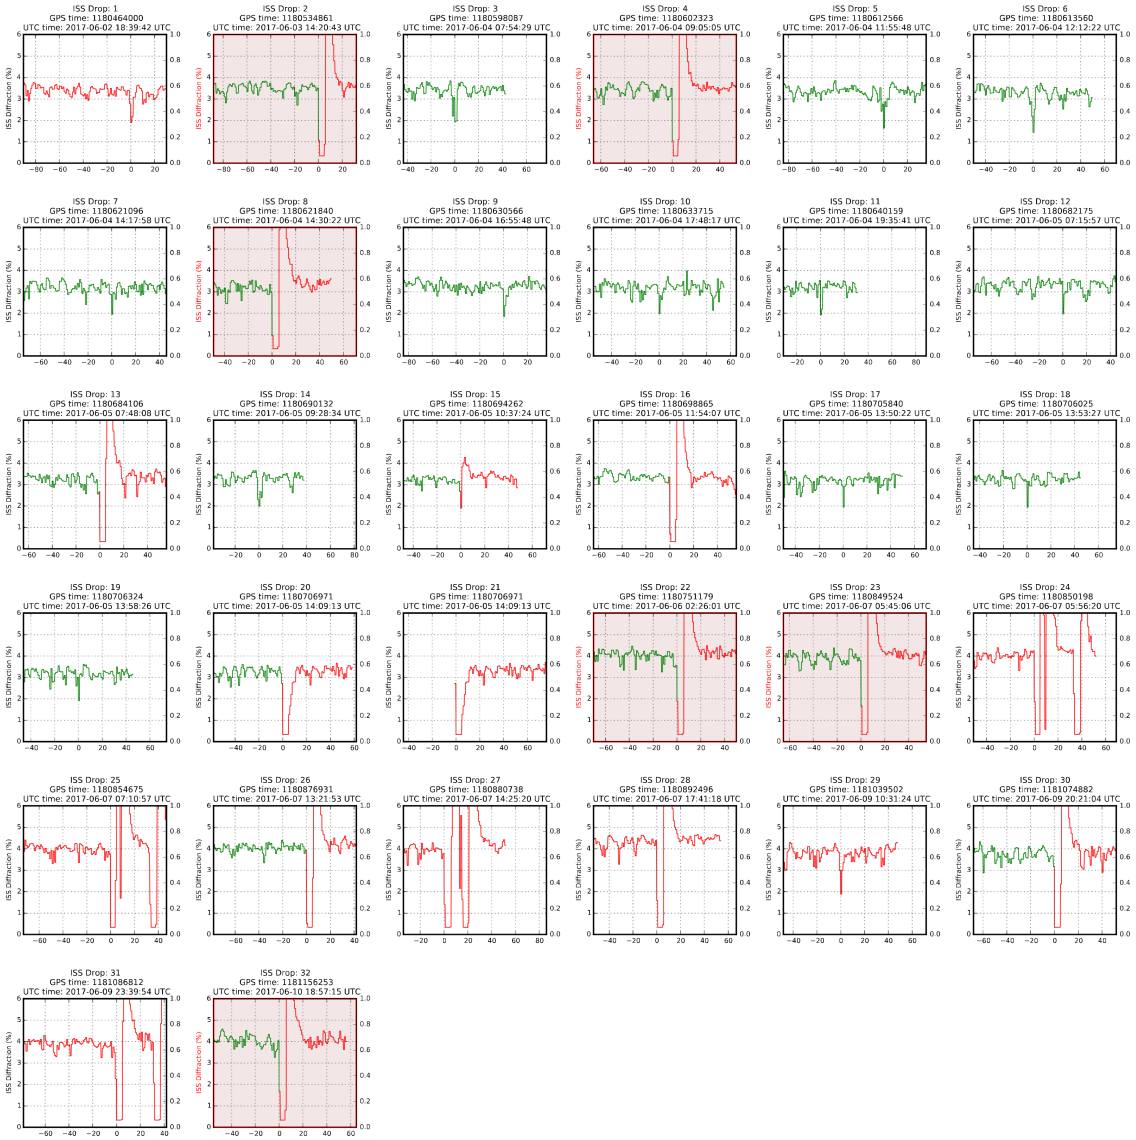


Figure 10.9: Catalogue of ISS dropouts during 10 days in June 2017. In this time 6 dropouts were identified that occurred before a lock loss and could have caused the interferometer to fall out of the locked state. Of the remaining dropouts in the ISS diffracted power, 12 were most likely caused by the lock loss as they occurred during but not before the lock loss. Another 14 dropouts did not affect the lock state of the instrument.

10.7 Conclusion of Investigation into ISS Instability

The laser system of the Advanced LIGO detectors is designed so that laser noise does not limit the performance of the detector. The PSL system should provide the beam that seeds the interferometer and in no other way affect the running of the gravitational wave detector. However, it was been shown how small issues with the PSL sub-systems can become problematic and that monitoring of the PSL system is important to prevent these issues from interfering with the long-term operation of the detectors.

From the investigation of ISS dropouts it was initially thought that the rate of dropouts was not strongly related to the diffraction set point of the AOM and as such the AOM drive offset could be lowered to slightly increase the optical power available to the interferometer. Sometime after reducing the diffracted power the ISS dropouts returned and again interrupted the locking of the main interferometer.

The finding of the ISS investigation was that the loop gain of the inner loop ISS was the main factor affecting the rate of dropouts in the ISS AOM diffracted power. It was also shown that there were a number of ISS dropouts that occurred immediately before lock losses of the full interferometer and a tool was developed and deployed to track these occurrences. It remains important to monitor this issue and when necessary make changes to the loop gain or diffraction strength to reduce the occurrence of ISS dropouts that can affect the full interferometer.

Chapter 11

A Frequency Stabilised 1550 nm Laser System for the Glasgow 10 m Prototype Facility

The current Glasgow 10 m Prototype [144] has been a test bed for gravitational wave detector development since the early 1990's. Housed in a temperature controlled clean room, the prototype consists of nine individually accessible vacuum chambers and two 10 m long beam tubes in the arrangement shown in Figure 11.1.

Previous experiments in the Glasgow 10 m prototype (e.g. [145,146]) have used 1064 nm NPRO lasers as a light source (see Section 8.2.1). While experiments with 1064 nm lasers are ongoing in the prototype lab, future experiments will use a longer wavelength of light inline with the development of technology for the next generation of gravitational wave detectors.

The choice of laser wavelength for gravitational wave detectors was discussed in Section 8.2.4. The decision has not yet been made on the wavelength and the type of laser to be in future experiments. For the Glasgow prototype it was decided to install a 1550 nm laser system. This choice was motivated by the availability of telecommunication lasers and the existing knowledge and experience with this type of laser and the availability of optics and photodetectors for this wavelength.

The first experiment to use 1550 nm light is the Polarisation Sagnac Speedmeter experiment described in Part II of this thesis.

11.1 The 1550 nm Laser System

As a light source for future interferometry experiments, the 1550 nm laser for the Glasgow 10 m prototype is required to be a low noise laser source with an output

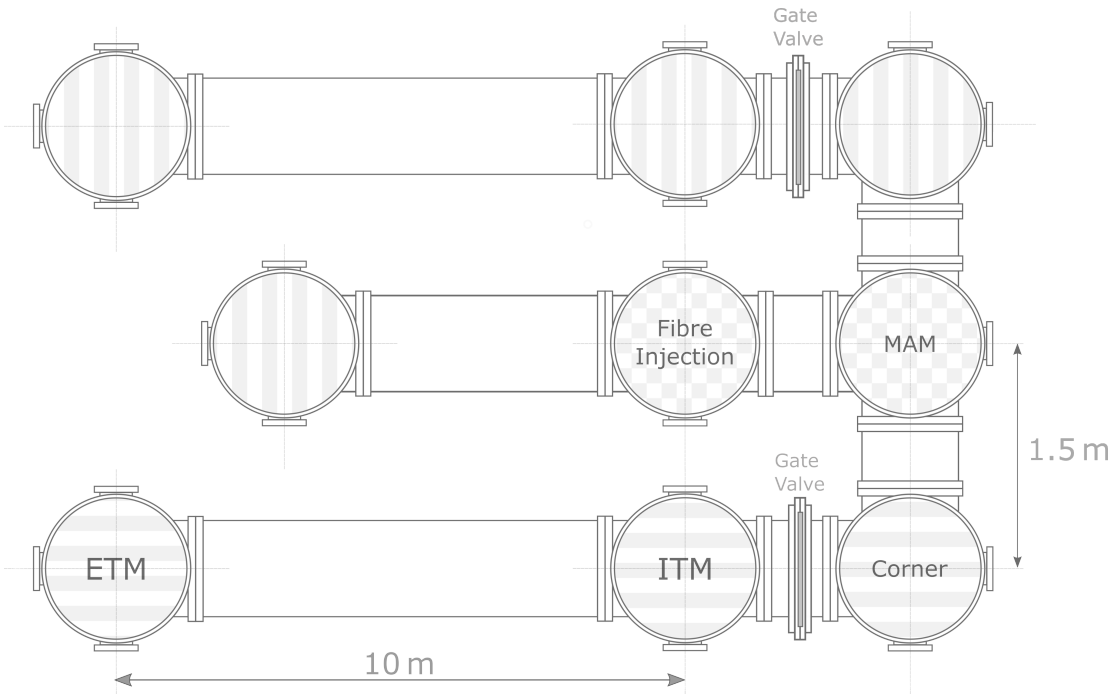


Figure 11.1: Overview of Glasgow 10 m prototype vacuum system. The vacuum system features two 10 m arms, separated from a central volume with gate valves, meaning there are three vacuum volumes that can be vented separately. Tanks are shaded with vertical lines to represent the presence of 1064 nm laser light and optics, horizontal shaded lines show the tanks used for the experiments with 1550 nm light, checked shading indicates tanks with both wavelengths of light.

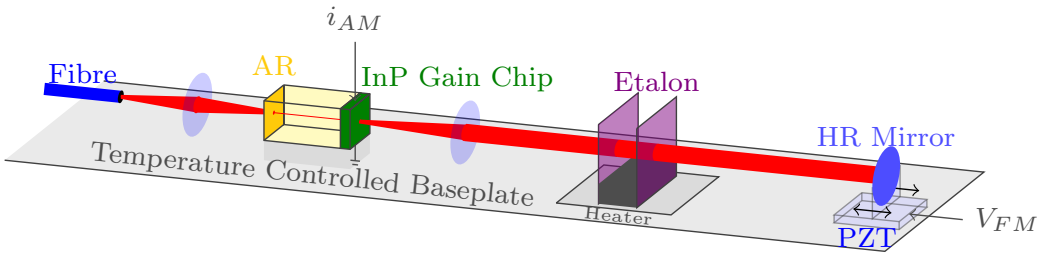


Figure 11.2: External cavity diode laser operating principle. Photons are produced when a current i_{AM} is applied through a diode consisting of an InP gain chip. A cavity is formed by a reflective surface of the gain chip (AR) and an external mirror. The resonant frequency of the cavity can be set with a PZT length actuator on the external cavity mirror or with temperature control of an intra-cavity etalon filter. Photons of the selected wavelength will resonate in the cavity and interact again with the gain chip to produce stimulated emission at the selected wavelength. The output power is set by the current through the gain chip. The laser frequency can be modulated by driving a voltage V_{FM} across the cavity PZT, and the intensity of the laser can be modulated by injecting a varying current through the diode. Current modulation also generates a level of frequency modulation. The cavity and diode are mounted on a common temperature controlled base and the output of the cavity is focused into a polarisation maintaining fibre.

power of a few watts, enough to support several in-vacuum and table-top projects run in parallel. The complete system includes a diode laser and a fibre amplifier capable of producing 5 W of optical power.

11.1.1 Diode Laser

The seed laser used for the 1550 nm laser system is an external cavity diode laser sold by Pure Photonics [147]. This is a telecommunications laser that has been re-purposed with access to frequency and intensity actuators internal to the laser module. A detailed description of the operating principle of this laser can be found in Figure 11.2.

The output power and frequency can be set through digital communication with the laser module. The current applied to the diode can be set by the module to give the requested output power which is measured with an internal photodiode pick-off.

The communication to the laser module includes both an amplitude modulation (AM) and a frequency modulation (FM) channel. The voltage provided to the AM channel changes the amount of current applied through the diode. Similarly, a voltage on the FM channel will drive the PZT mounted mirror on the laser cavity and change the resonant frequency of the cavity and therefore the

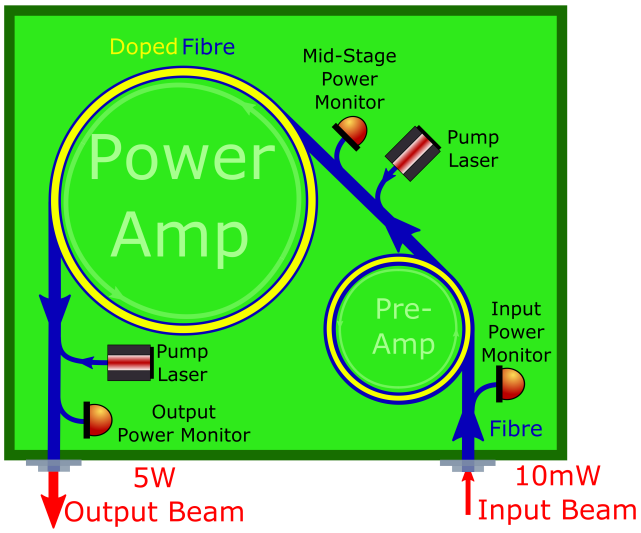


Figure 11.3: NuAmp fibre laser amplifier internal workings, showing the two stages of amplification and power monitoring stages.

frequency of the emitted light.

The diode laser is capable of emitting 50 mW of laser light from the fibre coupled output.

11.1.2 Fibre Laser Amplifier

The diode laser described above is used as the seed for a fibre amplifier capable of producing 5 W of 1550 nm laser power. For the Glasgow prototype a NuAMP single-frequency fibre amplifier was purchased from the company NuFern [148].

The general operating principle of a fibre laser amplifier is described in Section 8.2.2. The design of the NuAmp device is represented in the sketch in Figure 11.3.

The amplifier consists of two stages of amplification. The first stage (named the 'pre-amp' fibre) is a variable gain stage that ensures that a constant power seeds the second stage 'power amplifier' fibre. This ensures that the amplifier produces an output laser beam of steady power for a varying input power from the seed laser so long as the input to the pre-amp stage is within the operating range of 1-15 mW.

Each fibre gain stage is end-pumped with two laser diodes combined at a fibre tap. The current delivered to these pump diodes is set by the amplifiers internal control system. The more current drawn by the diode the more pump light illuminates the gain media leading to greater amplification. The gain of the first stage amplifier is controlled internally to provide a constant power at the mid-stage between gain fibres. The gain of the second stage can be set with the

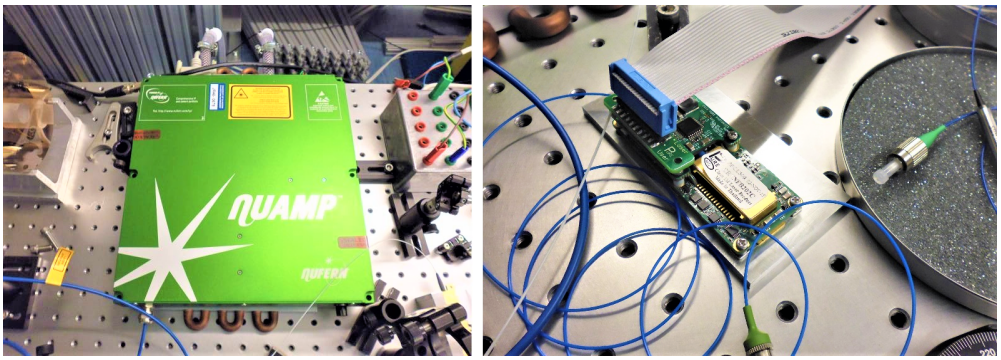


Figure 11.4: NuFern fibre amplifier (a) and Pure Photonics diode laser (b).

external GUI via USB connection to the amplifier control module. The output power can be set in the range 0.5 W to 5 W.

Scattered light propagating back into the gain fibre can cause damage to the amplifier, so the high-power output beam is coupled into air via a free-space isolator in the fibre coupled output of the amplifier. The laser amplifier also features an internal fibre tap monitoring the level of back-scattered light and the internal control includes an automatic shut off if a high level of back-scatter is detected.

11.1.3 Auxiliary Optics

Optical Fibres

Many elements in the 1550 nm laser system are fibre based. For this system there are two important requirements on the type of fibre optic used: the fibres need to be polarisation maintaining (PM) and use only angle point contact (APC) fibre connections (FC).

A polarisation maintaining fibre can be used to ensure that the polarisation of the light ejected from the output of the fibre is of the same orientation as the input light. These fibres are manufactured with stressing rods built into the cladding either side of the core (as shown in Figure 11.5). The rods stress the fibre in one axis, with no stress in the orthogonal axis, inducing birefringence in the fibre core. This stress induced birefringence defines a fast and slow axis for the light inside the fibre.

One of the birefringence axes is indicated by a key on the external fibre connectors so when multiple PM fibre components are used together they share a common polarisation alignment.

By aligning the polarisation of the light coupled into the fibre to one of the axes defined by the stress rods, the polarisation of the light though the fibre is

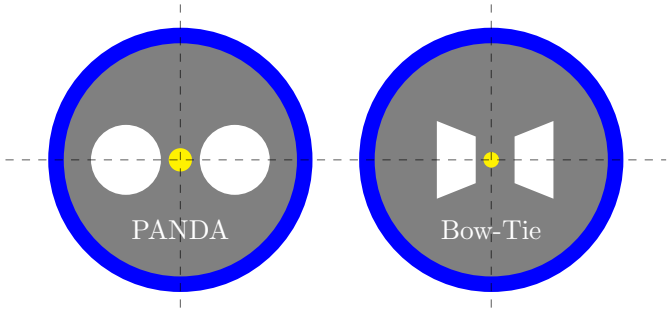


Figure 11.5: Cross-section view of the two types of polarisation maintaining fibre. The blue outer sleeve is an acrylic coating. The stress rods in white define polarisation axes through the grey inner cladding. Light propagates thought the yellow inner core.

maintained. The output light will be polarised in the direction indicated by the alignment key of the output coupler. Light polarised in any direction other than the fibre birefringence axes will also be transmitted by the fibre but the output polarisation will vary depending on the bending of the fibre and is susceptible to drifts due to temperature changes and vibration of the fibre.

Two designs of PM fibres are available, so called 'Panda' and 'Bow-tie' type fibres, named for the cross-sectional shape of the stress rods within the cladding. Bow-tie fibres can be made to generate higher levels of birefringence and are used in polarising fibres that actively attenuate one polarisation relative to the other. For the purposes of the 1550 nm laser system either type of PM fibre is sufficient, cross-sections of PM fibre are shown in Figure 11.5.

When two fibre elements are connected together a mating sleeve is used, this hold the tips of the two fibre against a shared fibre interface making a physical point contact between fibres. The manufactuing tolerance of these connections determines the rotational alignment tolerance of two fibres at the connection. For this experiment it was necessary to use angled point contract (APC) fibre connectors, to reduce back-scatter from fibre interfaces, preventing the formation of in-fibre parasitic cavities and back-scatter into the laser. The industry standard is to polish the fibre tips to have an 8° surface angle.

For the Glasgow 1550 nm laser system Thorlabs angled point contract polarisation maintaining fibres were used with narrow key mating connections (P3-1550PM-FC) [149].

Faraday Isolators

Faraday isolators are used in optical systems to prevent light traveling back along the input path. When set up correctly the isolator acts as an 'optical diode' where

light can only travel in one direction through the device. To achieve this a Faraday isolator consists of a pair of crossed polarisers on either side of a Faraday rotator.

Input light not of the same polarisation orientation as the polariser is rejected at the first stage. The then linearly polarised input beam is transmitted through the Faraday rotator that changes the orientation of the polarisation by 45°. The light then interacts with the second polariser which is orientated at 45° relative to the first polariser, allowing maximum transmission of the linearly polarised light.

Light of any arbitrary polarisation propagating in the reverse direction will first interact for the polariser at 45° which will transmit light polarised in axis of the polariser and reject all other polarisations. This light will then be rotated by 45° meaning it is in the linear polarisation that is maximally rejected by the first polariser. In the ideal case no light is transmitted though the isolator in the reverse direction.

The isolation factor of a Faraday isolator is limited by the polarisation extinction ratio and relative alignment of the two polarisers, and the homogeneity of the rotation of the Faraday rotator.

For the 1550 nm laser system, an in-fibre Faraday isolator is used between the seed laser and the amplifier. A second in-air Faraday isolator is used after the laser amplifier launch fibre.

If light was scattered from the amplifier back through the fibre to the seed laser it could cause damage to the laser diode (if it was high enough power) or destabilise the external cavity leading to increased frequency noise of the laser output.

Modulators

Modulators for 1550 nm laser light are available as in-fibre packages or as larger crystal optics for free-space beams. For the frequency stabilisation of the 1550 nm laser system a broadband in-fibre EOM is used, and for the generation of phase modulated sidebands a free-space resonant EOM is used (see section 8.2.3 for a description of the operating principle of an EOM).

An electro-optic modulator is characterised by the half-wave voltage (V_π), the voltage necessary to change the phase of the light by π radians. The half-wave voltage scales with the size of the crystal and depends on the electro-optic susceptibility of the material and the wavelength of the light.

The half-wave voltage of the in-fibre EOM is 4.5 V for 1550 nm light. The free-space electro-optic modulators with larger crystals require larger voltages to drive the same modulation, typically hundreds of volts. The free-space EOM used

to add the control sidebands is driven by an LC circuit resonant at 12.08 MHz (where the laser is assumed to be shot noise limited).

Vacuum Fibre Feed-through

For the 1550 nm laser system a method of injecting the beam into the vacuum system is required. Previous experiments with 1064 nm light in the Glasgow 10 m prototype have used a periscope to direct the laser beam from the optical bench into the vacuum system through a anti-reflection (AR) coated view-port. To avoid space conflict with the ongoing experiments, a fibre feed-through was used instead to directly inject the 1550 nm laser light into the vacuum system.

This feed-through consists of a polarisation maintaining fibre captured in a threaded cylinder of steel with a 40mm diameter KF vacuum flange.

The feed-through has angled point contact fibre connectors (FC/APC) on either side. On the air side, the beam in free-space on the optics bench is aligned and mode matched into a fibre collimator and a 1 m fibre patch cable couples the light to the feed-through. On the vacuum side, a bare fibre couples the light from the feed-through to a fibre coupled collimator that launches the beam in the vacuum system. A bare fibre is required as the fibre cladding material would out-gas under vacuum.

11.1.4 Performance of the Laser System

Frequency Noise

The frequency noise of the PurePhotonics seed laser is shown in Figure 11.6 compared to the frequency noise of a typical NPRO laser at 1064 nm.

The noise of the 1550 nm seed laser is a factor of 10 higher when compared to the NPRO below 100 Hz where the noise falls as $1/f$. Above 100 Hz the NPRO frequency noise continues to decrease as $1/f$ but the noise of the 1550 nm laser does not decrease as rapidly, meaning the noise at 100 kHz is three orders of magnitude higher than the NPRO noise.

Figure 11.6 also shows the noise target for the laser stabilisation, as limited by the thermal stability of the reference cavity and the bandwidth achievable with analogue feedback. Given the target noise level, the noise profile of the 1550 nm laser places a demanding requirement on the frequency stabilisation system used to reduce the free-running frequency noise of this laser.

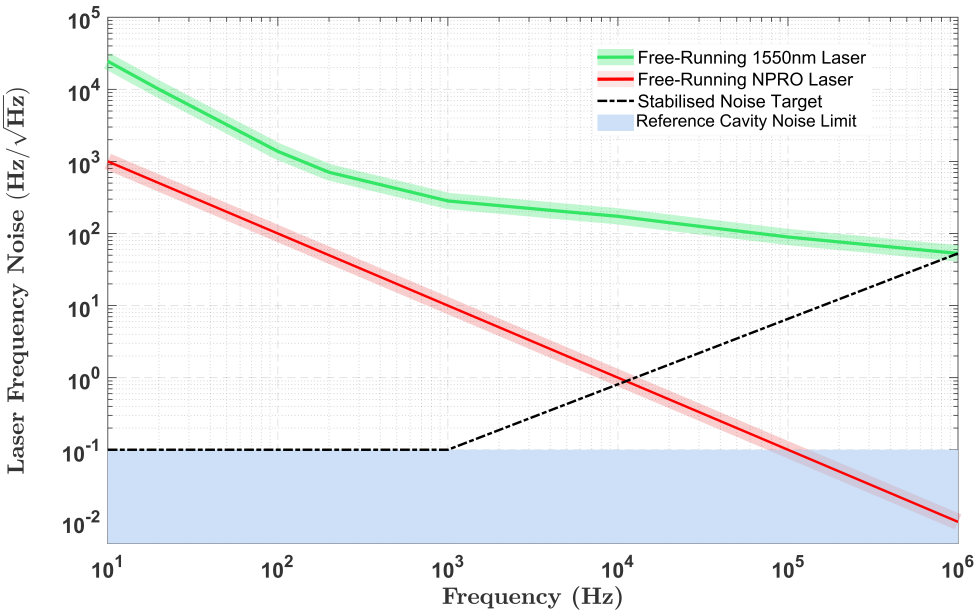


Figure 11.6: Laser frequency noise of the 1550 nm laser compared to a NPRO seed laser. The blue region represents the expected stability of the reference cavity. The black dashed line represents the target noise level for the frequency stabilisation system.

Intensity Noise

The intensity noise of the laser output was measured on the optical bench, with 3.5 mW of light measured on a photodiode. The amplitude spectral density of the relative intensity noise (RIN) and the time series of the laser power are shown in Figure 11.7.

It was found that the output of the laser amplifier as measured on the bench drifted with the lab temperature. The power stability of the laser amplifier output was measured over a 12 hour period with the laser in constant operation. The laser amplifier output was split from the main beam and a fraction of the optical power was measured on a photodiode; the trend is shown in Figure 11.8.

Polarisation Stability

The polarisation of the light injected into the system should be a high purity linear polarisation. The polarisation of the laser beam in vacuum is set with a HWP after the fibre launcher.

The polarisation of the light injected into the vacuum system was seen to be correlated to the temperature of the lab. Drifts in the polarisation of the laser beam injected into the vacuum system were seen to follow the 40 minute thermal cycle of the lab.

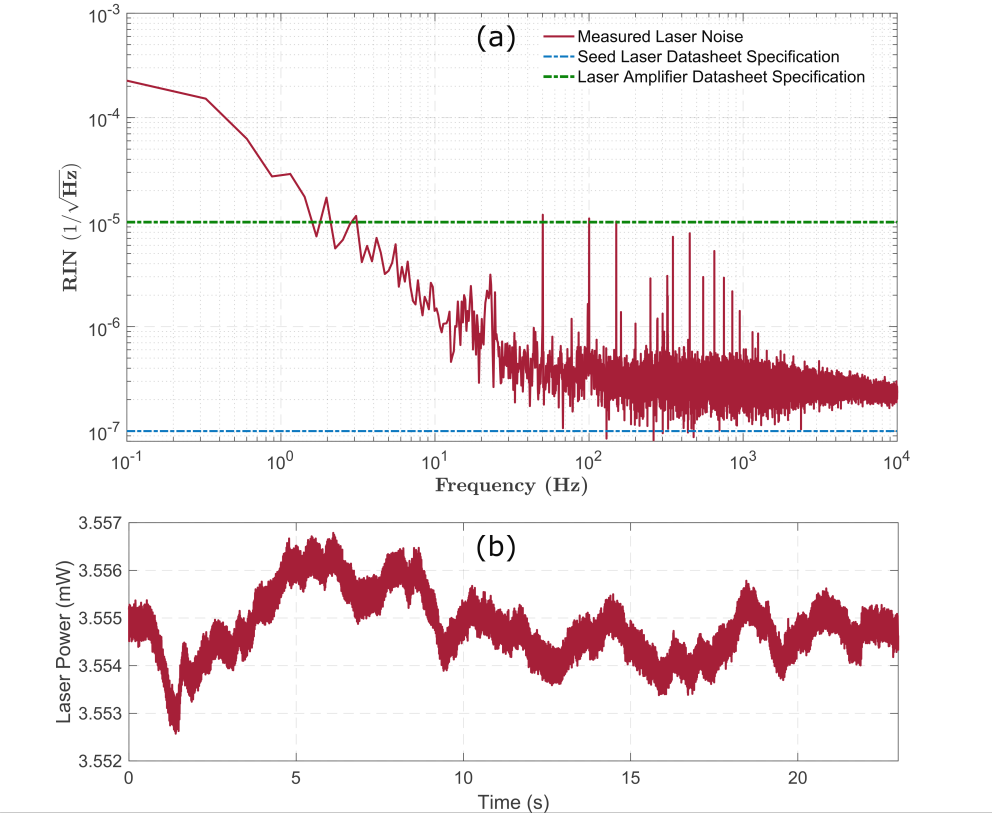


Figure 11.7: Laser intensity noise of 1550 nm laser system. The relative intensity noise (RIN) of the output of the laser amplifier is plotted in (a) as an amplitude spectral density from the time series of the measured laser power in (b). The expected level of RIN is plotted in (a) based on the specifications of the laser and amplifier.

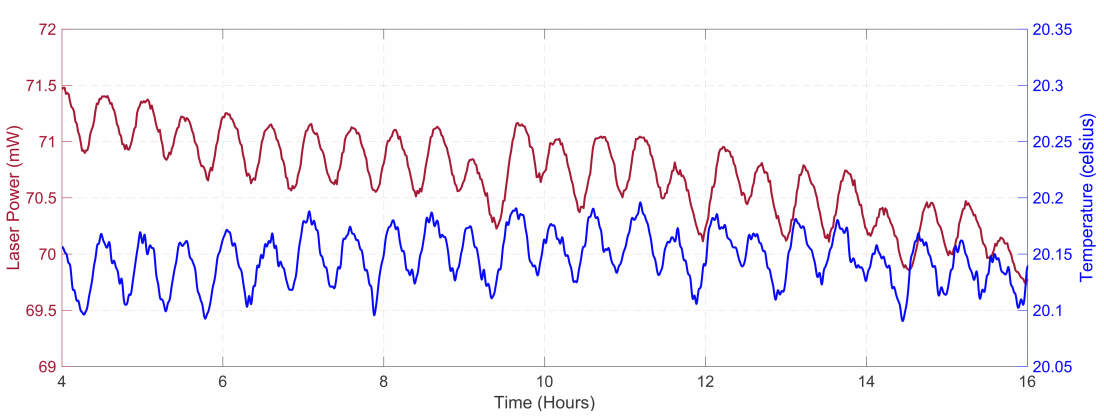


Figure 11.8: Long term power output of 1550 nm laser system. The red trace (left axis) shows the scaled power measured on the photodiode, a neutral density filter was used to reduce the 71 mW beam down to roughly 3.5 mW. The blue trace (right axis) shows the temperature of the lab during the measurement.

It was found that the in-air fibre delivering the beam from the optical bench to the vacuum feed-through was the main source of temperature-to-polarisation coupling. Thermal shielding was added to the fibre and the polarisation at the input was well matched to the birefringence axis, but drifts in the output polarisation were still evident. It it thought to be due to a stress induced depolarising effect at the fibre feed-through. The fibre captured within the vacuum flange is required to be polarisation maintaining, any stress applied to the fibre by the flange will introduce additional birefringence in the fibre core, changing the polarisation axis, resulting in a depolarisation of the light output and temperature driven drifts in the output polarisation.

11.2 Frequency Stabilisation of a 1550 nm Laser

The 1550 nm laser system was developed as a new laser system for the Glasgow 10 m prototype, and required active frequency stabilisation for use as a seed laser for interferometry experiments.

The first experiment to use the 1550 nm laser system was the polarisation Sagnac experiment (see Part II). For this experiment the free-running noise of the laser had to be reduced to a level where the frequency noise of the laser was less than the linewidth of the 10 m cavity to enable the cavity to be locked on resonance with the laser.

11.2.1 Reference Cavity

In order to stabilise the free-running noise of the laser, a frequency reference more stable than the laser is required. As described previously in Section 8.3.3, this can be achieved by use an optical cavity.

For this experiment a linear, rigid reference cavity formed by two mirrors bonded to a spacer was used. The spacer is made of Zerodur [150], a glass ceramic material with a low coefficient of thermal expansion.

The two cavity mirrors are bonded to flat silica plates that are themselves bonded to the spacer, forming a cavity of length 21.4 cm. The input mirror has a flat face with a reflectivity of 98.3 %, and is bonded with a piezo-electric (PZT) length actuator between the mirror and the silica plate. The end mirror has a concave face with a radius of curvature of 1 m and a reflectivity of 99.3 %.

The length stability of this reference cavity is expected to be limited by thermal noise, with the largest source of thermally driven length noise expected to result from the bond between the cavity mirrors and the silica plates. This will limit

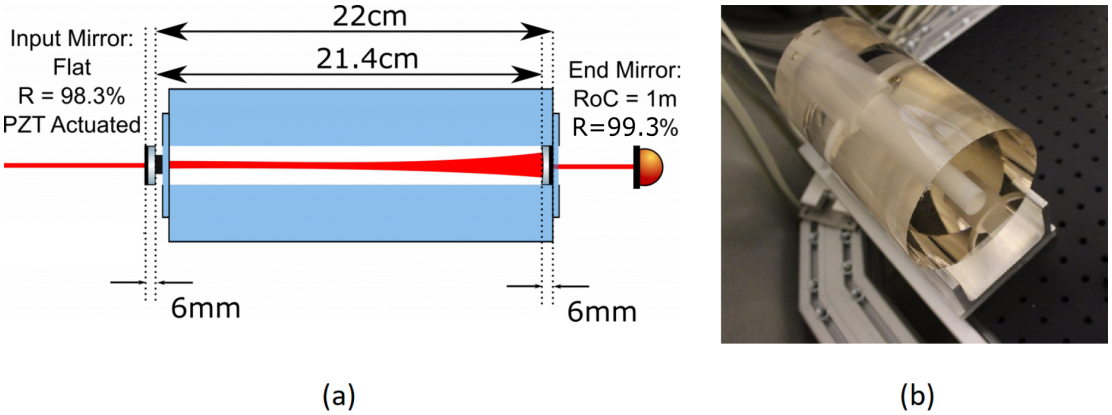


Figure 11.9: Rigid reference cavity for the frequency stabilisation system. Side-view of the reference cavity construction (a). Reference cavity spacer in the vacuum system (b).

the length noise of the cavity to of order $10^{-16}\,\mathrm{m}/\sqrt{\mathrm{Hz}}$ and set a lower limit of $0.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ on the stabilised frequency noise.

11.2.2 Stabilisation Scheme Design

To stabilise the frequency of the laser a Pound-Drever-Hall scheme is used to stabilise the laser to the length of the rigid cavity. An overview of the scheme developed to stabilise the 1550 nm laser is shown in Figure 11.10.

To reduce the free running frequency noise of the 1550 nm laser to the level of $0.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ a gain of 80 dB at 100 Hz is required.

The design requirement was to have a unity gain point between 500 kHz and 1 MHz. To achieve the high loop bandwidth two actuators are used to stabilise the laser at different frequencies. For stabilisation of the laser at high frequencies a fibre EOM (between the seed laser and amplifier) is used as a phase actuator. For frequencies between 1 Hz and 30 kHz the internal PZT actuator of the seed laser is used. This PZT modulates the length of the resonant cavity where the lasing process occurs.

The beam reflected from the reference cavity is ejected from the vacuum tank via a view-port onto a readout bench where the light is detected with a broadband photodiode (see Figure B.4 in Appendix B). The AC coupled signal from the photodiode is amplified by a HELA [151] RF amplifier with a gain of 11.4 dB. This signal is then demodulated with the 12.08 MHz local oscillator signal. The local oscillator frequency was set to match the resonant frequency of the EOM used to generate the RF sidebands.

A notch filter tuned to 12.08MHz is used after the mixer so that any local

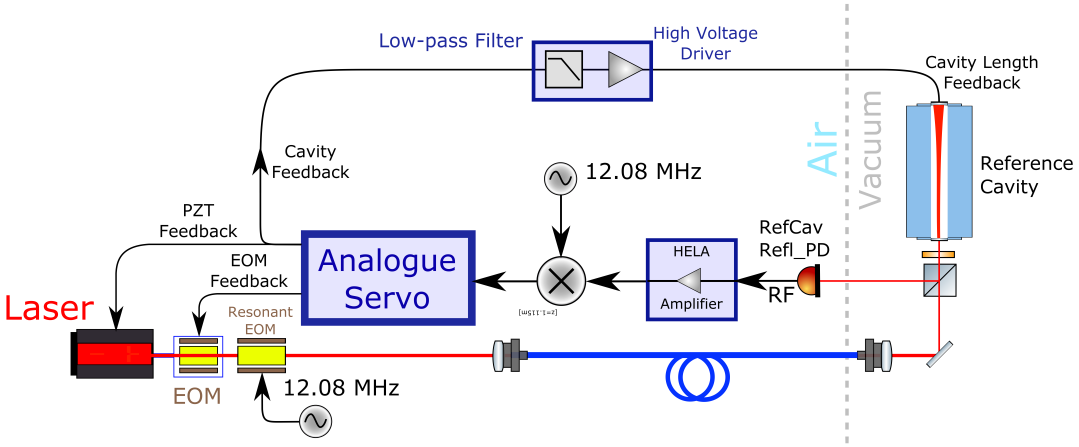


Figure 11.10: Frequency stabilisation scheme for the 1550 nm laser system. The laser output is phase modulated by a 12.08 MHz local oscillator signal applied to a resonant EOM. The beam is fibre coupled into the vacuum system and aligned to a rigid reference cavity. The return beam is detected with a broadband photodiode, the AC coupled output of which is amplified before being demodulated with the local oscillator to produce the error signal that is filtered and fed back to multiple frequency actuators. For wide bandwidth control, both an in-fibre EOM and an internal PZT actuator are used for laser stabilisation, while feedback to the cavity length at low frequencies is used to hold the cavity on resonance.

oscillator RF signal leaking though the mixer is removed from the demodulated error signal (see Appendix B.4 for circuit diagram of notch filter). In addition a five-pole low-pass butterworth filter with cut-off frequency at 5 MHz is used to further filter out any RF leakage.

The DC component of the demodulated signal is used as the PDH error signal, with the RF signals heavily filtered. This signal is then passed through a series of analogue filters and split between the different feedback paths.

The PZT signal is injected onto the laser PZT via a distribution box. This handles the electronic communication with the laser module.

Low Frequency Drifts

It was found during the early stages of characterising the laser system that the diode laser's output frequency drifted over long timescales. This drift could not be compensated for with the internal PZT as it lacked the dynamic range and could only maintain a lock to the reference cavity for about 10 s.

A scheme was devised to have the cavity follow the laser frequency drifts at low frequency while the laser is stabilised to the cavity length. This means that the free running laser noise is not reduced below the frequency of the crossover between cavity feedback and laser feedback, but the stabilisation loop can be

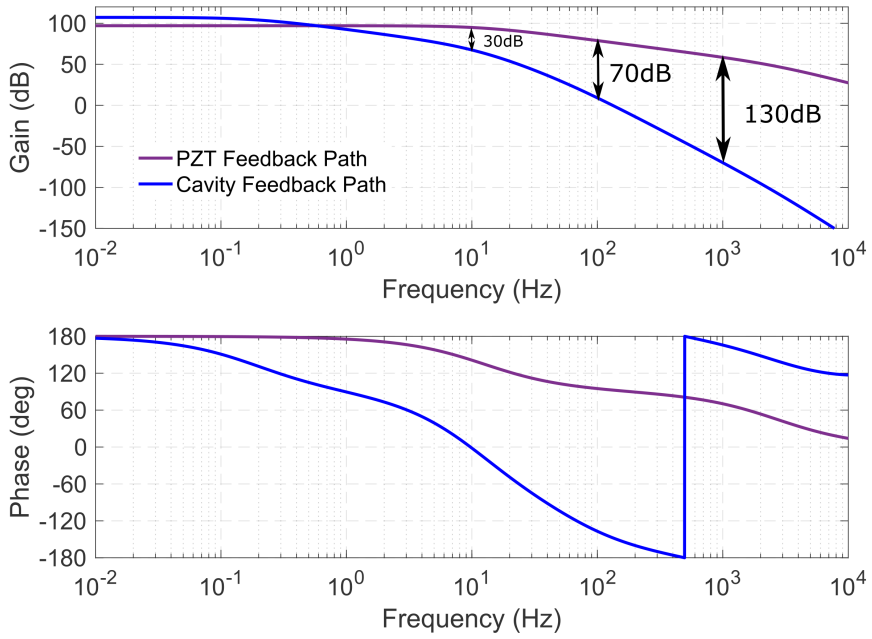


Figure 11.11: Simulated open loop gain of the feedback paths to the laser PZT (purple) and to the reference cavity PZT (blue). The shows the case with a crossover just above 0.5 Hz which gives a gain separation of 70 dB at 100 Hz between the two feedback paths.

locked indefinitely, and the noise of the laser is still stabilised above the low frequency crossover.

A feedback filter was designed so that the crossover between cavity feedback and laser feedback was between 0.3-0.7 Hz and the gain of the cavity feedback loop dropped off above 1 Hz so the high level of laser frequency stabilisation above 1 Hz could be achieved. This is shown in Figure 11.11.

The cavity locking signal is picked off from the PZT feedback signal and passed through a digital controller before being applied to the cavity PZT length actuator. This digital controller can also be used to scan the length of the cavity and includes the implementation of an 'auto-locker' for the stabilisation scheme.

The additional low frequency filtering of the cavity feedback signal is implemented in the digital controller and the output of the controller is amplified by a high voltage driver to the range needed to drive the cavity PZT. The voltage driver increases the signal from a maximum range of ±10 V to a range of -10 V to +90 V. The drive signal is then passed into the vacuum system via an electronic feed-through and applied to the PZT mounted to the input mirror of the reference cavity (see Section B.1).

11.2.3 Stabilisation Loop Response

When designing a filter to be part of a feedback loop, the frequency response of the other elements of the loop need to be considered so that the total loop response is as desired. In the laser frequency stabilisation loop the frequency response of the reference cavity and the frequency actuators need to be included in the loop design.

The response of the reference cavity, and the PZT and EOM frequency actuators is shown in Figure 11.12.

Reference Cavity

The reference cavity is an optical resonator and the response of the cavity to changes in frequency of the carrier is determined by the free spectral range (FSR) and the linewidth of the cavity resonance. Where the FSR of the cavity is the frequency spacing of the fundamental resonant mode. The FSR is determined by the length (L) of the cavity (FSR = $\frac{c}{2L}$). The linewidth, or full-width half-maximum (FWHM), is the ratio of the cavity FSR to finesse ($\mathcal{F}$).

For frequencies below the FSR of the cavity the frequency response is characteristic of a first-order low-pass filter. Up to the half-width half-maximum, also known as the cavity corner frequency (f_{cc}) the response is flat and falls off as $1/f$ above the corner frequency (see Section 1.5.2 for a full description of optical cavities).

The reference cavity used for the stabilisation of the 1550 nm laser is 21.4 cm long and has a finesse of 260 so the cavity corner frequency is 1.35 MHz, meaning the response of the cavity as a frequency discriminator is flat for the full frequency band of the stabilisation loop. However, the cavity corner frequency is close enough to the desired unity gain frequency that the pole needs to be considered.

Laser PZT Actuator

The PZT actuator is a piezo-electric length actuator that changes the resonant condition of the laser cavity. Piezoelectric actuators are resonant devices with a flat voltage-to-length response up to the resonant frequency of the device, above which the response peaks sharply and then drops off as $1/f^2$.

The resonant frequency of the PZT actuator internal to the 1550 nm laser is ~130 kHz, so for the frequency stabilisation servo being designed the PZT frequency actuator can be considered as having a flat response with a fixed actuation strength (measured to be 9.22 MHz/V).

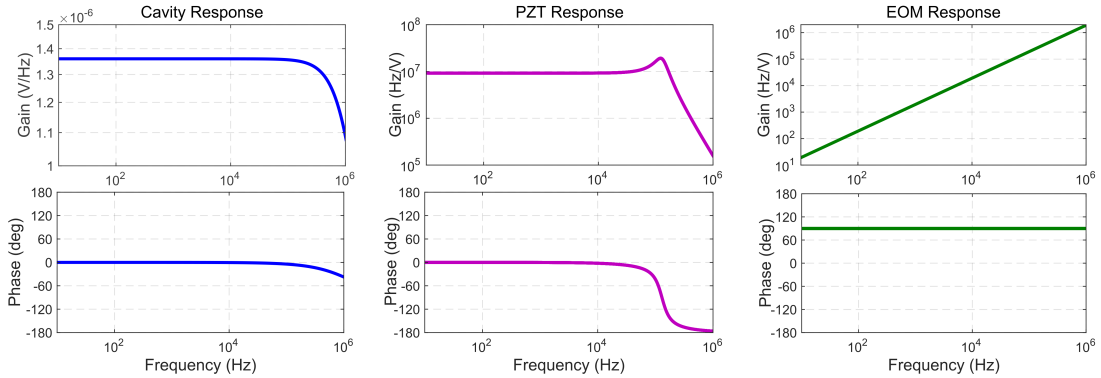


Figure 11.12: Response of cavity (a) as a frequency reference and the PZT (b) and EOM (c) frequency actuators.

Laser EOM Actuator

The EOM is an optical phase modulator, with the effect on the frequency being the integral of the phase modulation. As such the EOM has a sloped frequency response, the magnitude of the induced optical frequency change increases with the actuation frequency. This sloped actuation strength factor has to be taken into consideration when designing a stabilisation system using EOM feedback.

11.2.4 Analogue Servo Filter Design

The performance of a laser stabilisation servo is considered in terms of the amount of noise suppression produced over a given bandwidth and the stability of the loop. For the 1550 nm frequency stabilisation servo an analogue servo was designed with PZT and EOM feedback path shaping to achieve a crossover between the two feedback paths around 10-30 kHz and a unity gain point between 500 kHz -1 MHz.

To achieve the specified 80 dB of noise reduction at 100 Hz, an electronic servo gain of 50 dB is needed.

For a locking servo, the most important performance factor is the stability of the servo loop which can be evaluated by the gain and phase margin of the open loop response. A loop will become unstable if the open loop gain crosses unity gain at any frequency where the open loop phase is less than $-180°$. We can therefore describe quantitatively the stability of the loop by two parameters the phase margin and the gain margin. The phase margin of a open loop response can be calculated as the difference between the phase at the unity gain frequency and $-180°$ or can be expressed as the amount of phase lag that can be added to the loop to make the loop unstable. Similarly, the gain margin is the minimum gain change that would make the loop unstable.

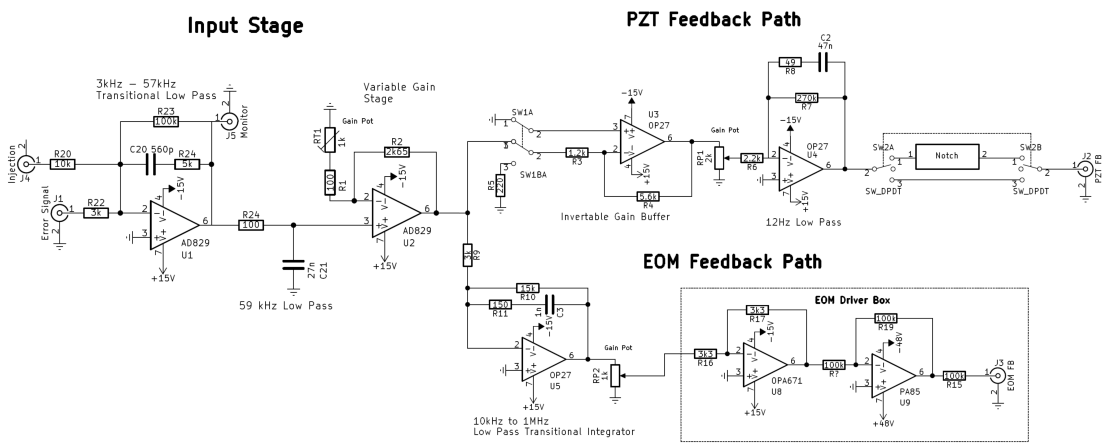


Figure 11.13: Frequency stabilisation servo circuit diagram for the 1550 nm laser system. The input stage of the circuit includes a transitional integrator that in combination with a passive RC low pass filter forms a 3 kHz low pass filter (this design is explained in Section 11.2.4). The PZT feedback path includes a 12 Hz low pass filter and switchable inverting stage. The EOM feedback path includes a transitional integrator and a high voltage driver stage included in a separate box.

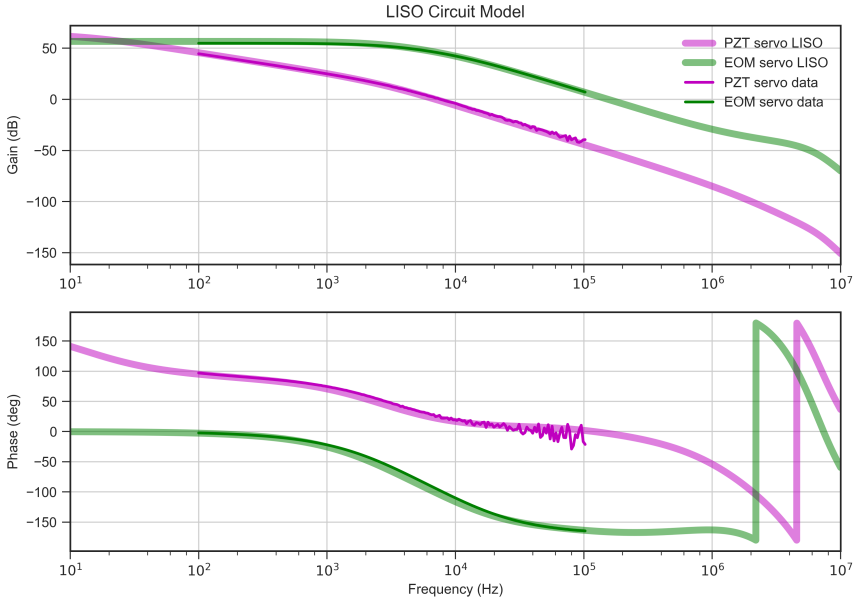


Figure 11.14: Transfer function of the analogue servo for laser stabilisation. The modelled and measured PZT (purple) and EOM (green) feedback paths are shown in the Bode plot.

The minimum phase and gain margin that is considered acceptable will vary from case to case, but in general a phase margin greater than 25° and a gain margin no less the 3 dB are considered as necessary for stable operation. A phase margin of 90° is maximally stable.

The analogue servo circuit diagram is shown in Figure 11.13, with Figure 11.14 showing the modelled and measured transfer functions of the filter electronics.

The servo can be described in three sections; the input stage (or common path) and two parallel feedback paths. Each stage includes analogue filters to shape the feedback signal as required for stable locking.

Input Stage

The input stage of the servo amplifies the error signal and applies the first stage of filtering. The error signal input is low-pass filtered by the common path input stage that includes an active transitional integrator op-amp stage and a passive RC low-pass filter that together form a low-pass filter with pole at 3 kHz.

The common input stage includes a potentiometer that can be used to adjust the overall gain of the stabilisation servo. This stage also includes an injection and monitor point, so the state of the stabilisation loop can be monitored and characterised by injecting signals into the error point of the loop.

Two Analog Devices AD829 [143] op-amps are used in the input stage as they have the high gain bandwidth needed to accept the high frequency error signal. (In Figure 11.15 the open loop gain of three op-amps considered for use in the input stage are compared to the electronic gain of the input stage).

The Analog Devices OP27 [152] is a general purpose op-amp that has good low noise performance, the OP37 [153] is a faster alternative to the OP27 with a similar internal design. But both options lack the speed for this application at the input stage where a high gain bandwidth product is needed. Of the op-amps considered, only the AD829 has enough open loop gain up to high frequency for this application. Without sufficient open loop gain, the op-amp will be unable to fully cancel the signal voltage at the input leading to a differential input voltage across the inverting and non-inverting inputs which can cause harmonic distortion.

The low-pass filter is formed in this two-stage implementation, because the AD829 op-amp used for the active filter is not unity gain stable. To ensure stability, the op-amp is implemented in a transitional integrator format so the gain of the AD829 stage does not pass though unity. The addition of the passive RC filter forms the low-pass filter.

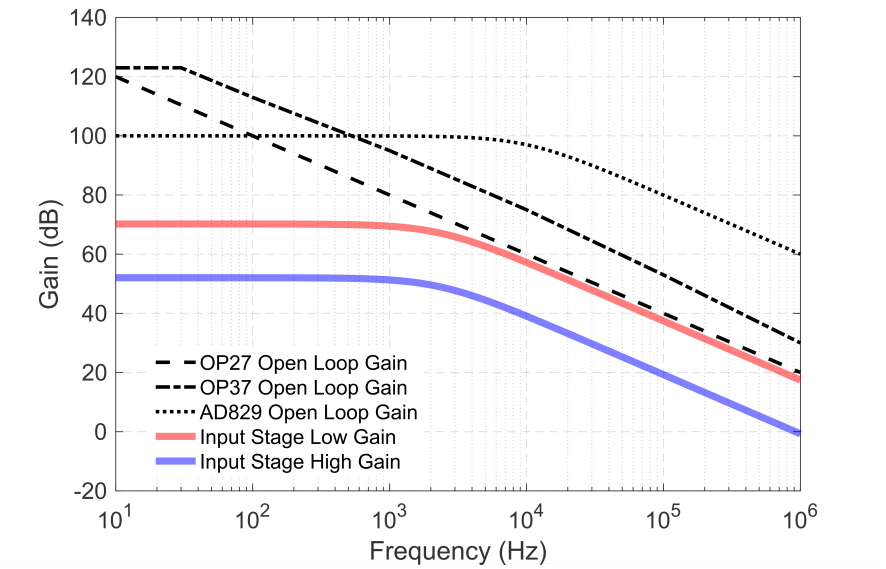


Figure 11.15: Gain of the input stage of the analogue servo for high and low gain settings. The open loop gain of three op-amp choices (OP27, OP37 and AD829) is also plotted based on the data-sheet values [143, 152, 153]. To ensure enough open loop gain the AD829 is the only one of these options suitable for the input stage of the analogue servo.

PZT Feedback

The PZT actuator response is flat up to 130 kHz (see Figure 11.12), so to force the PZT feedback path to crossover with the EOM feedback, which has a slope of $1/f$, an additional low-pass filter is needed in the PZT path. The PZT low-pass filter is implemented with a pole at 12 Hz.

The PZT feedback signal is applied to the laser PZT via a driver box that limits the signal voltage to the safe operating range of the device (see Section B.2).

EOM Feedback

The EOM is a phase actuator with a voltage-to-frequency response proportional to frequency (f). The sensing point of the PDH loop has a flat frequency-to-voltage response up to the cavity corner frequency at 1.35 MHz (see Figure 11.12)

So if the loop gain is to cross unity below the cavity corner frequency with a slope of $1/f$, a servo filter with a slope of $1/f^2$ is needed in the EOM path. In addition to the common path low-pass filter, a transitional integrator filter is implemented in the EOM feedback path with a pole at 10 kHz. The zero of the transitional integrator is set to be at 1 MHz to compensate for the reference cavity pole.

For the EOM feedback an additional high voltage driver stage is needed to apply the frequency correction needed to stabilise this 1550 nm diode laser. This increases the available range of voltage from ±15 V to ±45 V. This driver is included in the circuit diagram in Figure 11.13.

11.3 Results

An in-loop noise measurement represents the degree to which the free running noise of the laser is forced to follow the noise of the frequency reference. In principle the in-loop noise can be suppressed below the noise of the reference at which stage it no longer represents the achieved frequency noise of the stabilised laser.

To characterise the laser frequency stabilisation both in-loop noise measurements and measurement of the loop transfer function are presented in this section along with a consideration of limitations to the performance of the stabilisation system.

11.3.1 Limitations of the Stabilisation Servo

A number of factors affected the implementation of the frequency stabilisation scheme with the original design goals of having a unity gain point at 1 MHz and a crossover between PZT and EOM feedback at 30 kHz.

Laser PZT Instability

It was found that the mechanical interaction between the PZT frequency actuator and the laser module had a destabilising effect on the servo loop. Several mechanical back-action resonances were observed in the loop response transfer functions between 20 kHz and 100 kHz these are understood to be internal mechanical modes of the laser module (see Figures 11.16 and 11.17).

If the PZT actuated mirror inside the laser cavity is considered as a mass on a spring, then the back-action resonances can be understood as acoustic oscillations propagating within the structure of the laser module and reflecting from multiple material interfaces. The fundamental resonance of the PZT is at ~130 kHz but these additional features limit the bandwidth of the laser PZT as a linear frequency actuator.

The original design was to have the PZT and EOM feedback paths crossover at 30 kHz and the frequency shaping of the analogue servo was designed accordingly.

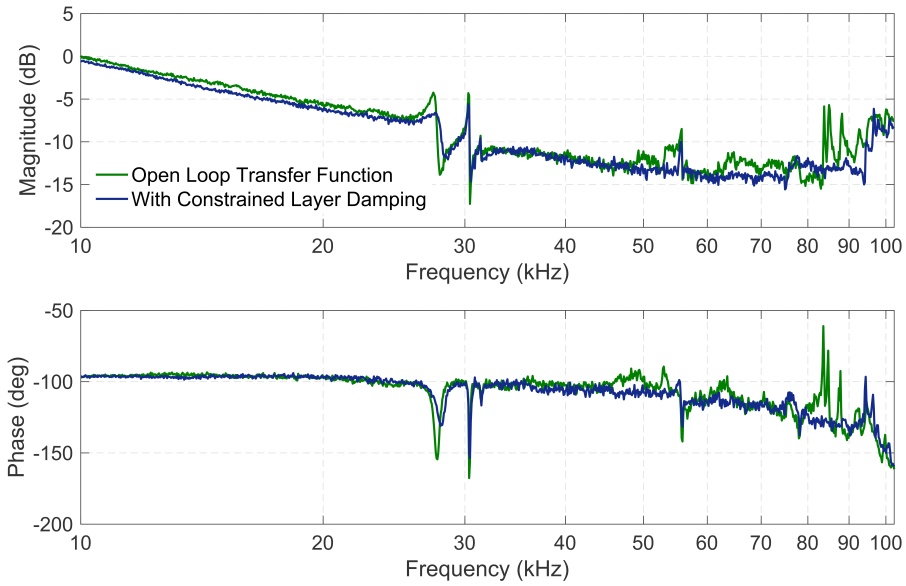


Figure 11.16: Open loop transfer function measurement of the laser locking loop. The green trace shows the resonance feature in the open loop response. The effect of mechanically isolating the laser module is shown in the blue trace where the magnitude of the peaks between 25 kHz and 32 kHz is slightly reduced and above 50 kHz there are fewer peaks in the loop response.

However, if the crossover is above 20 kHz, the PZT resonances cause the loop to become unstable with strong oscillation observed around 30 kHz.

Lowering the crossover to below 15 kHz prevents the mechanical oscillations from ringing up as the EOM actuator becomes the dominant feedback actuator in the 30 kHz region.

Attempts were made to reduce the impact of these resonances with mechanical damping. The laser module is mounted on a aluminium plate (pictured in Figure 11.4 (b)) which was separated from the laser bench by a layer of damping tape between the aluminium plate and a copper board. This provided a level of constrained layer damping [154] that improved the stability of the laser PZT but did not completely remove the resonances (see result in Figure 11.16).

By adding mechanical damping in the form of clamps that apply pressure to the aluminium mounting board, the magnitude of the resonance peaks could be changed. This is shown in Figure 11.17 where the peak at 27 kHz is damped but the peak at 30 kHz is increased in magnitude.

With combination of external mechanical damping and lowering the crossover the effect of these features in the loop response could be improved but not completely mitigated.

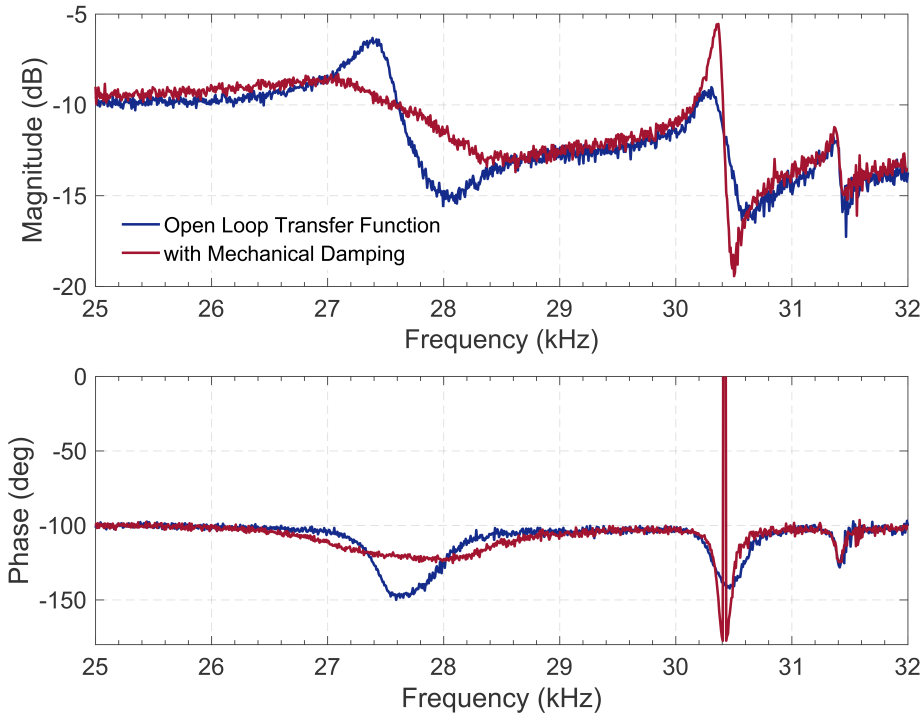


Figure 11.17: Open loop transfer function measurement of the laser locking loop between 25 kHz and 32 kHz showing mechanical resonances of the laser. In the blue trace the laser module was mounted on a aluminium plate, on top of a copper board with a layer of damping tape between the two surfaces; in the red trace the board was firmly clamped.

EOM Saturation

Lowering the crossover between PZT and EOM feedback forces the EOM actuator to correct a greater amount of laser frequency noise.

The EOM response is proportional to frequency, so at 10 kHz the actuator response is ten times smaller than the response at 100 kHz meaning a larger feedback voltage is required for actuation at lower frequencies.

In the original design of the analogue servo the EOM driver was limited to a range of ±15 V. With the lower crossover the EOM driver would run out of range when stabilising the laser frequency, resulting in a saturation in the servo loop.

A high voltage driver was needed to feedback signals to the EOM actuator. A dedicated driver box with an Apex PA85 [142] op-amp was used with an output voltage range of ±45 V. Higher feedback voltages were not used to avoid overheating the EOM by driving RF signals larger than the specificed magnitude into the device.

Slew Rate Limit of Op-amps

The Analog Devices OP27 [152] is a low noise op-amp with good all-round performance for use at audio frequencies. In the initial design of the frequency stabilisation circuit for the 1550 laser system, the OP27 amplifier was used in several of the filter stages.

However, when used to lock the laser to the reference cavity the loop was found to become unstable when the gain of the circuit was increased. The first OP27 op-amp in the PZT path (U3 in the circuit diagram in Figure 11.13) was found to be oscillating due to the slew rate limit of the op-amp. This is in part due to the comparably high noise profile of this 1550 nm laser resulting in large in-loop feedback signals.

An alternative op-amp, the Texas Instruments TL071 [155], with higher slew rate was used as a replacement. The key parameters of these two op-amps are shown in Table 11.1. Although it has a faster slew rate, the TL071 has a lower gain bandwidth product and higher input noise when compared to the OP27.

The impact on the in-loop noise due to this op-amp choice is presented in Figure 11.18.

In the original design that used the OP27, as the gain was increased the noise was reduced across the frequency band apart from around the crossover frequency just below 20 kHz, where the noise increased. As the gain is raised to half of the full range, a strong in-loop oscillation rings up at 10 kHz and odd numbered harmonics, indicating a non-linearity in the loop. The low frequency noise is also

	OP27	TL071
Slew Rate	$2.8\,V/\mu s$	$16\,V/\mu s$
Gain Bandwidth Product	8 MHz	4 MHz
Input Noise	$3.3\,nV/\sqrt{Hz}$	$15\,nV/\sqrt{Hz}$

Table 11.1: Comparison of important datasheet parameters of the OP27 [152] and TL071 [155] op-amps considered for the PZT feedback stage of the analogue frequency stabilisation circuit.

raised.

With the TL071 in the circuit, as the gain is increased the noise across the band reduces with only a slight bulge in the noise around 18 kHz. As the gain reaches the maximum of the adjustment range, a similar oscillation rings but of a significantly smaller peak magnitude and with less effect on noise across the band.

With the TL071 in the stabilisation servo, the target noise level could be achieved with careful tuning of the gain controllers and the crossover frequency.

Loop Phase Delay

The phase margin at the unity gain frequency is the sum of phase delays in the servo loop. The servo electronics have a finite bandwidth that can be modelled.

The reference cavity acts as a single pole low-pass filter with a pole at 1.35 MHz which adds approximately 45° of phase lag to the loop at this frequency.

The bandwidth of the the stabilisation loop is also limited by the finite signal travel time. The loop is comprised of an optical path length of 9.28 m between the laser and the photodiode in reflection of the reference cavity, 1.46 m in air on the laser bench, 2.12 m inside the vacuum system and 5.7 m inside optical fibres. In addition, the signal cables between the photodiode and the frequency actuator total approximately 6 m where the signal travels a $\frac{2}{3}c$.

The total length of the control loop leads to a loop phase delay of 60 ns, equivalent to 21.6° of phase lag at 1 MHz.

The measured open loop response of the laser stabilisation servo was found to be consistent with a total loop delay of 140 ns (see Figure 11.20).

11.3.2 Frequency Noise Measurements

For the ideal application of a noise stabilisation servo the open loop gain represents the level of noise suppression that is being applied to the laser across the bandwidth of the servo. From cavity locking crossover to the unity gain frequency the inherent noise of the laser is stabilised to the reference value.

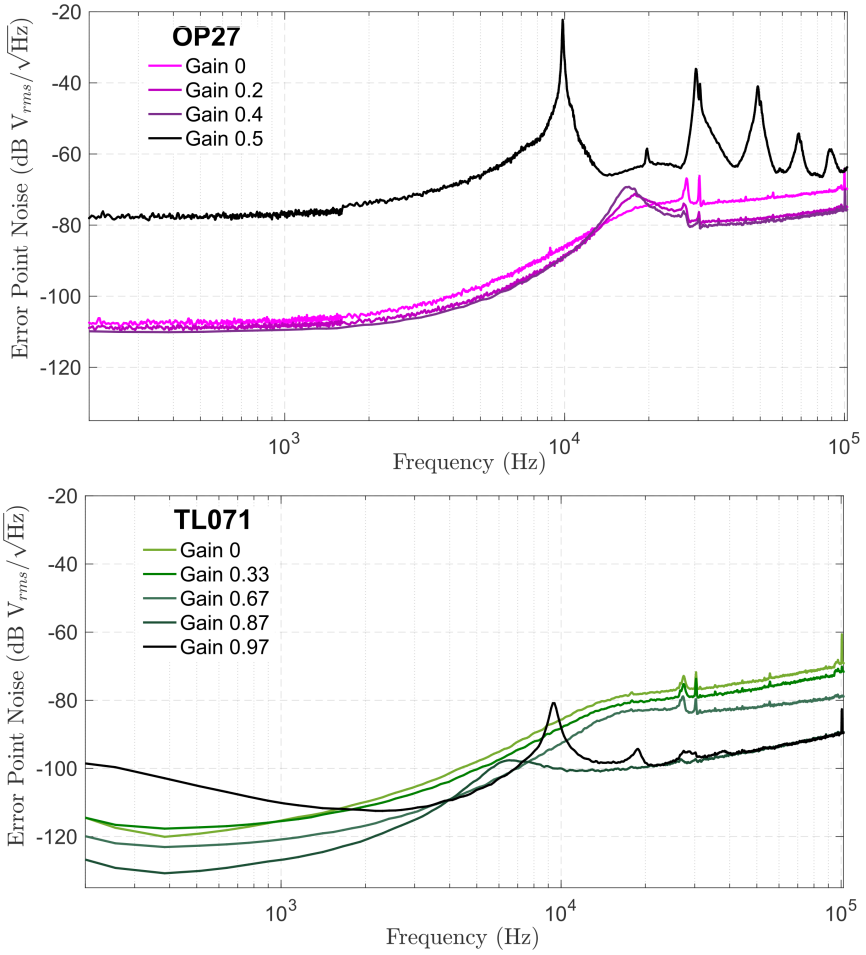


Figure 11.18: Error point noise in-lock with different gain setting of the analogue servo with OP27 (top) and TL071 (bottom) op-amps in the PZT feedback path. On each plot the error point noise for different settings of the overall gain adjustment from 0 to 1 are shown. The top plot shows the oscillating behavior of the OP27 in this circuit, whereas the oscillation is less evident with the implementation of the TL071 in the lower plot and a lower noise level is achieved.

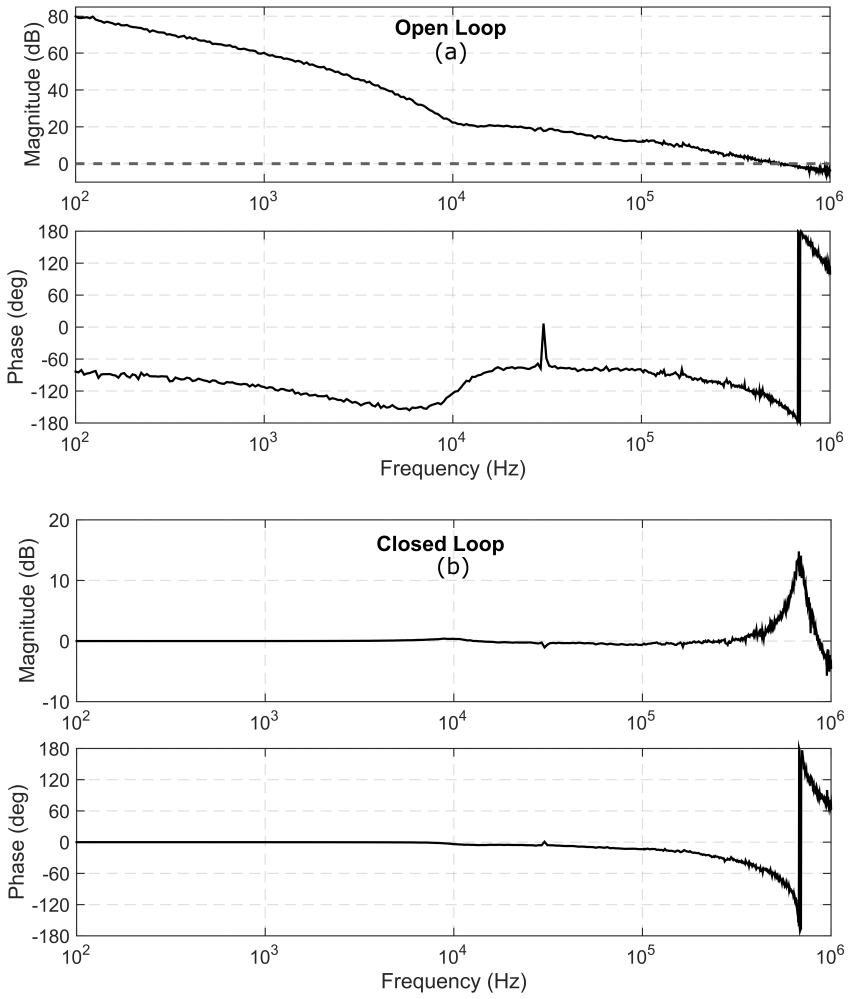


Figure 11.19: In-loop transfer function measurements of the 1550 nm laser locked to the reference cavity. The open (a) and closed (b) loop responses of the laser stabilisation loop are shown. The unity gain point is 526 kHz with a phase margin of 38.8°.

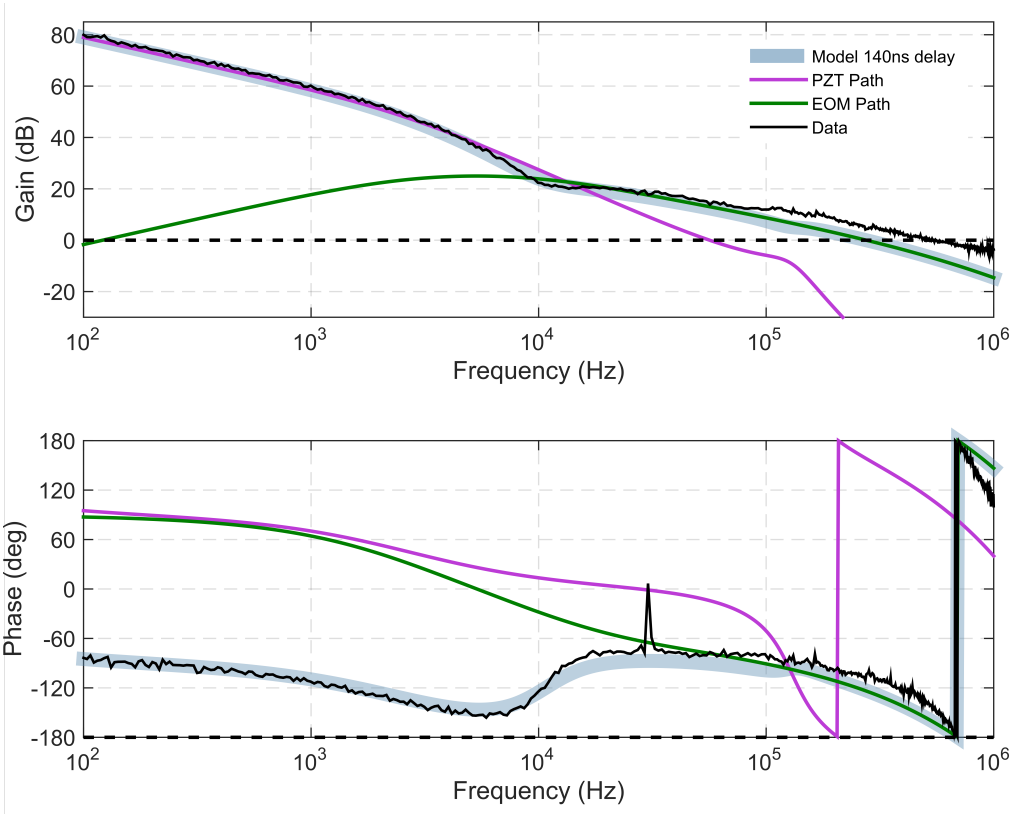


Figure 11.20: Modelling the laser stabilisation servo from the LISO simulation of the servo circuit and the known response functions. The purple trace shows the modelled PZT path response and the EOM path in green. The combined loop response is shown in the blue trace and can be made to fit the measured data by including 140 ns of time delay.

The signal noise at the error point of the locking loop is representative of the frequency noise of the laser. The error signal is the free running noise of the laser suppressed by the open loop gain of the loop below the unity gain point.

In real applications the in-loop measurement of the stabilised laser noise can be mis-leading as a measure of the stabilised noise level. Issues with the servo electronics can add noise to the laser that appears to be stabilised in-loop. To ensure the servo is working as expected, dark noise measurements at different points in the loop can help to diagnose issues with the servo and eliminate the possibility of adding noise to the laser in the servo loop.

The closed loop response of the laser stabilisation loop was measured and from this the open loop response was calculated. Figure 11.19 shows the open and closed loop response of the laser stabilisation loop. In Figure 11.20 the open loop response is matched to a model of the loop response.

An example of an in-loop noise measurement is shown in Figure 11.21. In this case the laser was locked to the reference cavity with the analogue servo set-up to give a unity gain frequency around 500 kHz and a crossover between PZT and EOM feedback around 15 kHz. The measured noise is within a factor of two of the original target noise level of $0.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$.

11.4 Summary

This chapter described a 1550 nm laser system that was installed and tested in the Glasgow 10 m prototype interferometer facility. This laser system was used as the light source for the polarisation Saganc speedmeter experiments described in Section II. This laser system will also be used as the light source for future experiments in the Glasgow prototype facility.

Much of this chapter focussed on the design and implementation of an analogue frequency stabilisation servo for the laser system. In Chapter 12 a digital servo is considered for this laser system.

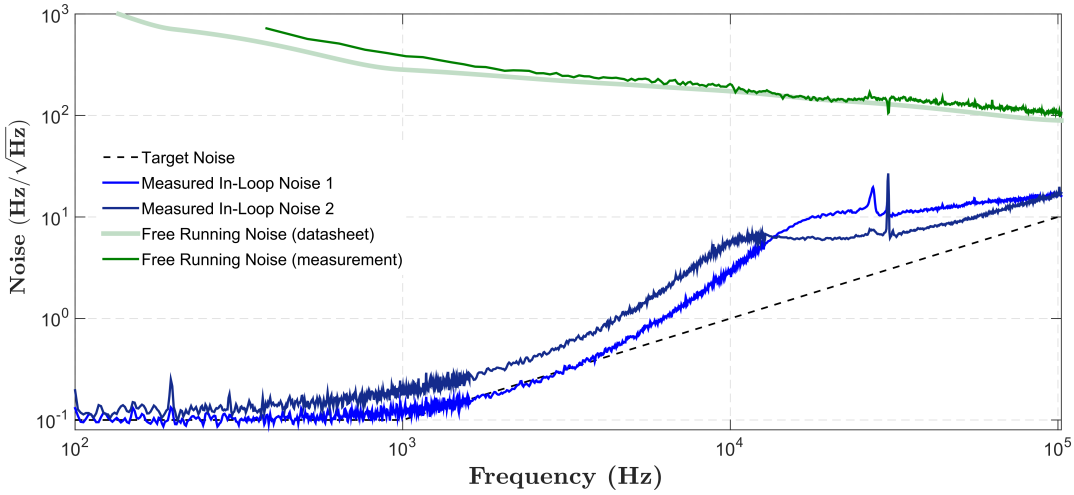


Figure 11.21: In-loop measurement of stabilised laser frequency noise for two different crossover set-points (blue) with respect to the free running laser frequency noise (green). The noise level was measured at the error point of the stabilisation loop locking the laser to the reference cavity. The two blue traces were measured with different gain settings of the electronic servo, for the lighter trace the crossover was higher, lowering the noise below 10 kHz but more strongly ringing up the instability at 27 kHz. The free running noise was inferred from a previous measurement of the in-loop noise and the loop response transfer function and found to be in agreement with the datasheet for the laser.

Chapter 12

Fast Digital Filtering for Laser Frequency Stabilisation

In Chapter 11 a frequency stabilisation system for a 1550 nm laser was presented. This scheme used an analogue servo to filter the locking signal. In this chapter a digital filter system is investigated to replace the electronic servo stage. This alternative laser stabilisation scheme maintained most of the hardware described in Section 11.2.

Digital filtering is used in the LIGO gravitational wave detectors to lock the interferometer using a digital control and data acquisition system, CDS [99]. The same hardware is used on a smaller scale in the Glasgow prototype interferometer, and was utilised to lock the 10 m suspended cavity in the polarisation Sagnac experiment described in Chapter 4.

The same digital control system is not used for laser stabilisation however, as it lacks the speed to adequately stabilise the laser noise across the required bandwidth. The sample rate of CDS is 16 kHz, while the control bandwidth is limited to approximately 10 kHz by the anti-aliasing and anti-imaging filters at the input and outputs (see Section 12.1.3).

For faster digital control a dSPACE [156] processor was used with an available sample rate of 250 kHz. The design target for this stabilisation scheme was to have a unity gain point of 10 kHz with high gain up to 1 kHz while maintaining loop stability. This requires a steep shaping to the frequency response between 1 kHz and the unity gain point (UGP).

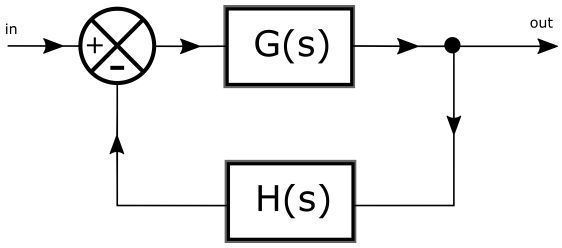


Figure 12.1: Block diagram of a basic control system with negative feedback.

12.1 High Speed Digital Filtering

A block diagram of a basic control system is shown in Figure 12.1. In this process negative feedback is used to apply a corrective error signal to the input of the plant ($G(s)$), with the signal generated by filtering the output of the plant with filter block ($H(s)$) [157].

12.1.1 Discretisation

A transfer function, $H(s)$, is a representation of a continuous-time linear time invariant (LTI) system and can be expressed as the ratio of two frequency dependent polynomials of order n

$$H(s) = \frac{P^n_{num}(s)}{P^n_{den}(s)}. \tag{12.1}$$

A digital filter is by nature a discrete-time LTI system, so the transfer function has to be converted to the discretised z-plane

$$H(z) = \frac{P^n_{num}(z)}{P^n_{den}(z)}. \tag{12.2}$$

The discretisation of the analogue filter to a digital form is an important part of the calculation of the filter parameters, as it approximates the continuous transfer function of a filter as a discrete filter for a given sample rate.

Various method of discretisation exist, with the simplest methods based on the rectangular block digitisation, where a filter is approximated by a series of rectangular sections of height equal to the amplitude at each sampled point.

Tustin's method is based on the trapezoid rule of digitisation and provides a better approximation of an arbitrary filter and ensures that the transform of the filter into the z-plane results in a stable and accurate representation of the desired filter [158].

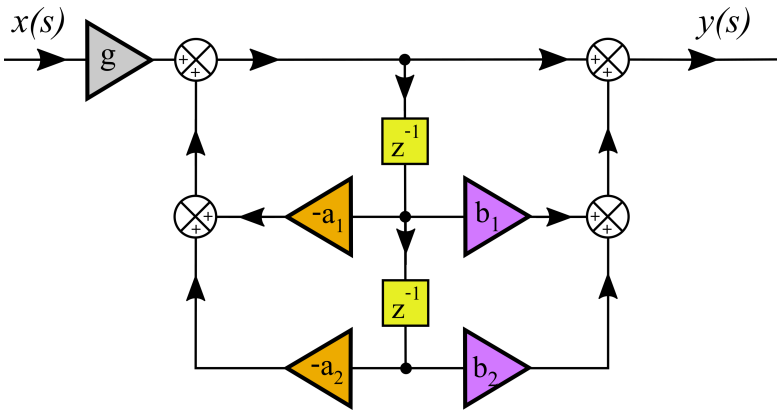


Figure 12.2: Block diagram of an implementation of a digital filter as a second-order section in the 'direct-form-II' layout of filter parameters. The triangular blocks represent the multiplication of the signal with the filter coefficients, and the square yellow blocks represent unit time delays.

12.1.2 Implementation of digital filters

A transfer function that contains two poles and two zeros can be represented by second-order polynomials with the numerator and denominator in the following form

$$H(z) = g\frac{1 + b_1 z^{-1} + b_2 z^{-2}}{1 + a_1 z^{-1} + a_2 z^{-2}}, \tag{12.3}$$

with more complicated functions represented by higher-order polynomials or by a series combination of 'second-order sections' such as this [159].

The transfer function described by Equation 12.3 can be implemented in the block diagram in Figure 12.2, where the filter parameters a_n and b_n multiply the input signal $x(z)$ with time delays z^{-1} to produce the output signal $y(z)$. This block implementation of a transfer function is known as 'direct-form-II' [160].

12.1.3 Aliasing

When a continuous signal is converted to a discreet time signal the signal is sampled at a finite sample rate ($f_s = 1/T_s$). Components of the signal at frequencies greater than half the sample rate will not be faithfully reconstructed [157].

Figure 12.3 illustrates the aliasing effect where sampling a signal with a frequency greater than the sample rate results in a false signal of lower frequency being re-constructed, this is known as signal *aliasing*.

When a digital signal is converted back into an analogue signal a similar process called *imaging* can occur where signals at frequencies (f_{sig}) below the sample frequency are reconstructed as signals either side of the sample frequency, $f_s \pm f_{sig}$ and multiples of the sample rate.

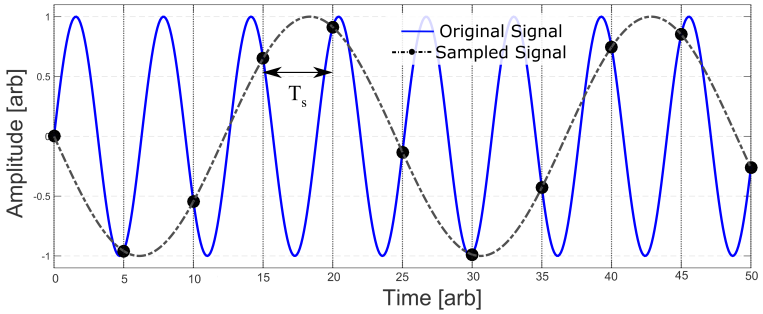


Figure 12.3: Representation of the aliasing problem with digital signal sampling.

To prevent aliasing and imaging from limiting the performance of the digital stabilisation system bi-quadratic low-pass filters are used before and after the digital servo as so called anti-aliasing and anti-imaging filters. This is covered in Section B.5.

12.1.4 dSPACE Digital Processor

For the frequency stabilisation system a dSPACE MicroLabBox digital controller [156] was used, with a 16-bit digital processor and a sample rate of 1 MHz.

The dSPACE MicroLabBox includes analogue-to-digital (ADC) and digital-to-analogue converters (DAC) at the input and outputs, either side of a field programmable gate array (FPGA) used as a multiplexer and de-multiplexer for the digital processor.

In total it takes four clock cycles for a signal to go from input to output, meaning the model running of the processor is limited to a maximum sample rate of 250 kHz. If such speeds can be achieved then it should be easily possible to run a filter model that shapes the error signal of a frequency stabilisation servo to provide stable locking of the laser with a unity gain point in the range of 10-40 kHz.

A Simulink® [161] model is compiled onto the digital processor to implement the digital laser control scheme. Details of this model and the user interface can be found in Appendix D.

The full stabilisation model could not be compiled onto the processor at the maximum sample rate due to the complexity of the model. Instead the model was compiled with a sample rate of 166 kHz.

12.2 Laser Frequency Stabilisation with a Digital Servo

The experimental layout of the frequency stabilisation servo with digital filtering uses similar hardware to the system described in detail in Section 11.2, with the main difference being that this version of the scheme uses digital filtering of the error signal and only feedback to the laser PZT actuator (as described in Figure 12.4).

The unity gain point of a digital servo loop is limited by the speed of the digital processor. However, across the available bandwidth the amount of noise suppression can be enhanced with complex shaping of the servo filters.

In order to reduce the free running laser frequency noise of the laser to a level that is limited by the thermal noise of the reference cavity an open loop gain of 80 dB at 100 Hz is required. The noise suppression bandwidth of the loop was initially set to be 1 kHz to provide a decent separation between the high gain level and the unity gain point.

12.2.1 Digital Filter Design for the Laser Stabilisation System

For designing the stabilisation system, it is important to also consider the bandwidth across which there is to be significant noise suppression. From here on, the corner frequency (f_c) is referred to as the highest frequency in the open loop response with high gain. This defines a turning point in the shape of the open loop gain between the high gain response for noise suppression below f_c and the transition towards the stable unity gain crossing. By increasing the steepness of the gain slope, the separation of corner frequency and unity gain frequency can be reduced without impacting the stability of the loop.

Low-pass Filter

The simplest filter design to meet this requirement is a single pole low-pass filter. This ensures a $1/f$ roll-off above the filter corner frequency. The transfer function of such a filter is shown as the black trace in Figure 12.5.

Transitional Integrator Filter

An improvement on the simple low-pass scheme is to use a transitional integrator filter. This is a single pole, single zero filter that has the same $1/f$ slope above the

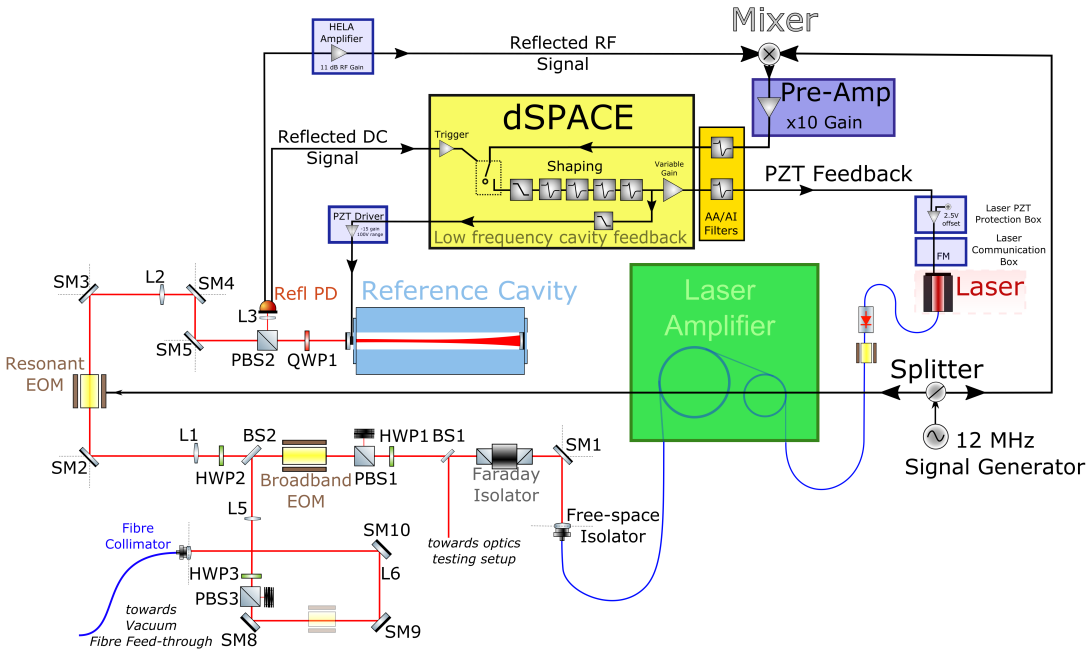


Figure 12.4: Top down view of the full table-top 1550nm frequency stabilisation servo, including table-top optics, laser, electronics and digital servo. The laser system is described in detail in Section 11.1. In this diagram the reference cavity (see Section 11.2.1) is shown on the optical bench; in the final version of the stabilisation servo the reference cavity was installed inside the vacuum system but the optical layout is the same. A PBS and QWP combination is used to pick off the light reflected from the reference cavity that is detected on a photodiode. The DC photodiode signal is used as a trigger to turn on the feedback to the laser once the cavity is close to resonance. The reflected RF signal is demodulated with the same modulation signal used to add the frequency modulation sidebands to the laser light. This signal is amplified before and after the mixer, and the mixer output is low-pass filtered. The output of the digital servo is fed back onto the laser PZT. The servo also features a secondary feedback loop that is used to lock the length of the reference cavity to follow the slow drifts of the laser frequency (see Section 11.2.2).

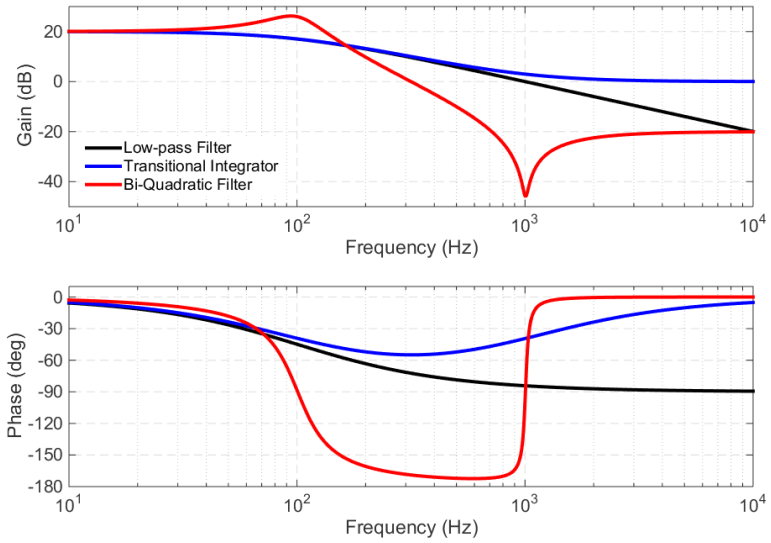


Figure 12.5: Transfer functions of some basic filter types; low-pass (black) formed from a single pole, a transitional integrator (blue) of one pole and one zero, and a bi-quadratic filter (red) of a pair of poles and zeros to form a resonant peak and notch combination. In this example filter the peak has a Q-factor of 2 and the notch has a Q-factor of 20, the higher the Q-factor the steeper the change in gain and phase around the resonance.

corner frequency but a flat gain at high frequencies (blue trace in Figure 12.5). The addition of the zero means the phase of the filter increases towards 0° at high frequencies which means the phase margin of the loop will be improved at the expense of gain margin.

Bi-quadratic Filters

A bi-quadratic filter is a complex two pole, two zero filter. In a low-pass application, this filter features a complex conjugate pair of poles at a frequency lower than the frequency of a complex conjugate pair of zeros. This results in a filter with flat gain above the zero frequency and below the pole frequency with a $1/f^2$ slope between the two (red trace in Figure 12.5).

The difference in the gain level above and below the sloped region increases with the separation between the pole and zero pairs. The phase of the filter is zero when the gain is flat and drops down towards $-180°$ in the sloped region between poles and zeros. The depth of the phase drop depends on the separation of the poles and zeros and the steepness of the phase drop depends on the Q-factors of the pole and zero pairs.

Combined Filter Design

The approach adopted here was to use a transitional integrator filter that would ensure a stable crossover at a chosen unity gain frequency and add bi-quadratic low-pass filters on top of this basic filter to increase the servo gain without changing the unity gain point.

A careful balance needs to be struck between high gain and high corner frequency, and between phase margin and gain margin.

The gain at low frequency can be increased by increasing the separation of the poles and zeros but this increases the separation of corner frequency from the UGP, meaning less noise suppression up to high frequency.

Similarly, the phase margin at the UGP can be improved by increasing the Q-factor of the notch. This increases the rate at which the phase increases back towards 0° about the notch frequency. But by increasing the depth of the notch the gain margin is reduced. If the gain of the open loop response drops below 1 at a frequency where the phase is less then −180° then the loop will not be stable.

The design of the filters for the laser stabilisation scheme and the modelled effect on the open loop gain of the laser stabilisation system is shown in Figure 12.6. To better balance the competing requirements an automated filter design and optimisation scheme was developed. This is discussed in Section 12.4.

12.3 Laser Stabilisation Results

The digital stabilisation filter was used to lock the frequency of the 1550 nm laser to the rigid reference cavity. The open and closed loop responses of the laser stabilisation system are shown in Figure 12.7. In this plot the result is shown for a digital servo with only the basic locking filter applied to shape the locking signal. In this case a unity gain of 9 kHz was measured with a phase margin of 50°.

Also shown in this plot is the result measured with the basic locking filter and two bi-quadratic filter to increase the gain below 1 kHz. The additional filters changed the overall gain slightly such that the unity gain point was measured at 6.5 kHz with a phase margin of 33°. Adding the two stages of bi-quadratic filters increased the gain of the loop at 100 Hz from 40 dB to 60 dB.

The in-loop noise of the stabilised laser is shown in Figure 12.8. This plot shows the noise with the laser stabilised using a basic filter (blue) and the result of adding bi-quadratic filters (red) to the digital servo. The improvement in the noise level could not be measured due to technical issues thought to be related to

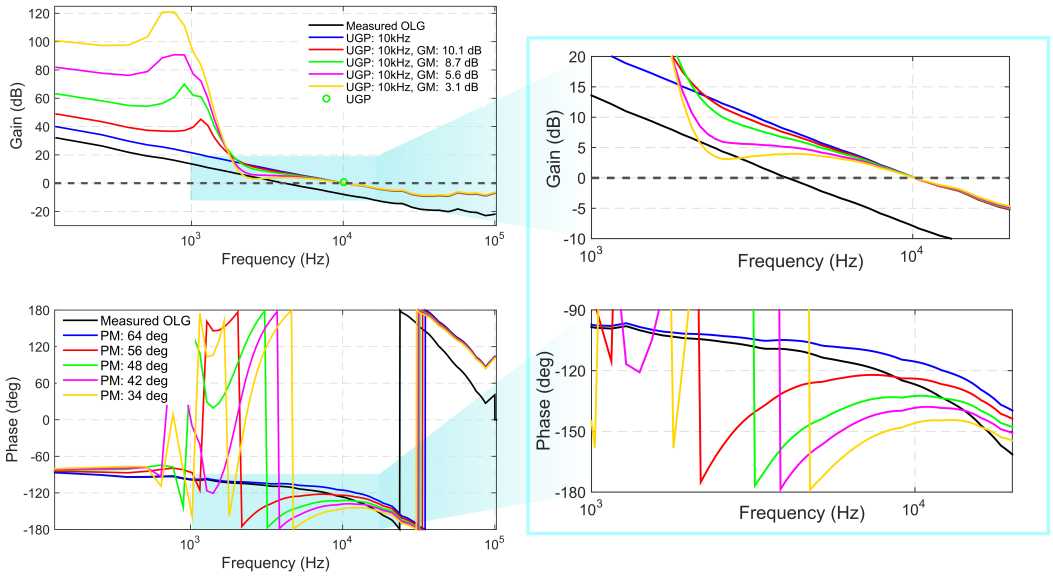


Figure 12.6: Digital filter design for frequency stabilisation servo showing the simulated open loop gain of the stabilisation loop in the main plot and a zoomed in version around unity gain in the insert on the right. The open loop gain was measured with simple low-pass filter (black) to lock the laser. This result was used to simulate the effect of implementing the designed filters with a unity gain of 10 kHz. With a transitional integrator and three bi-quadratic filters the design gain of 80 dB up to 1 kHz is achieved. An additional filter was included to further boost the gain to 100 dB.

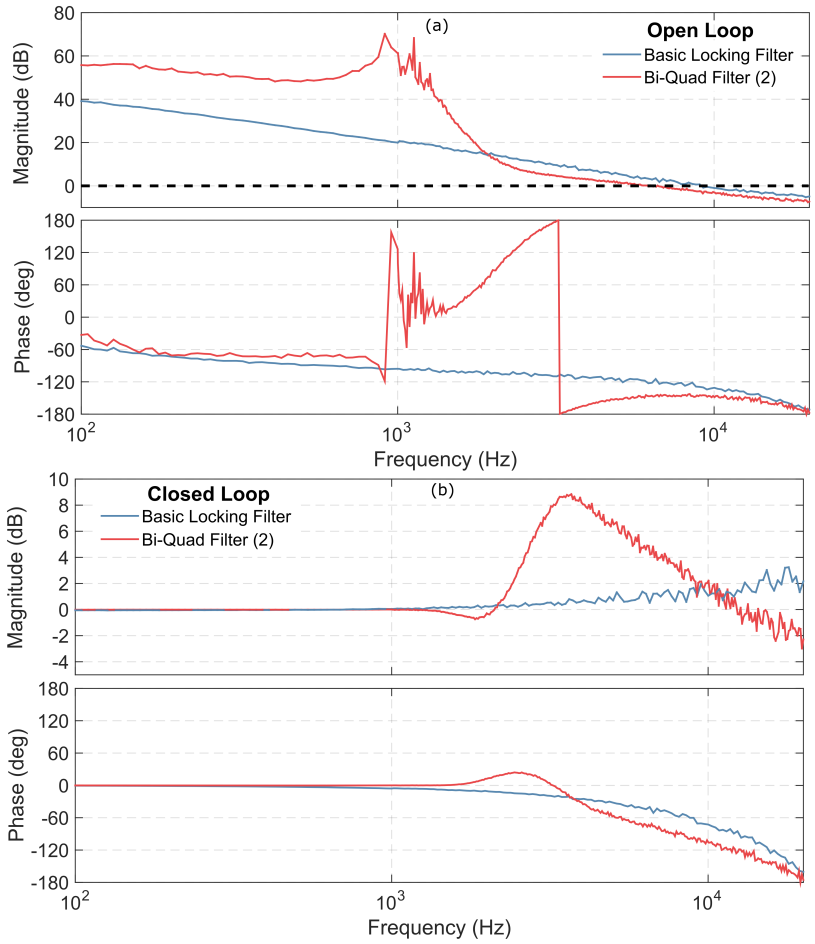


Figure 12.7: Open and closed loop transfer functions of laser locked to reference cavity with basic digital filter (blue) and a two stages of bi-quadratic filters to increase the gain below 1 kHz (green).

the input noise of the digital processor that limited the in-loop noise measurement to $10\frac{\text{Hz}}{\sqrt{\text{Hz}}}$.

12.4 Digital Filter Optimisation

An algorithm was developed to automate the filter design process and use simple statistical methods to optimise the filter design and fine tune the filter parameters.

To develop a successful algorithm two essential elements need to be established; a way of generating the filters to be tested, and a metric by which filters can be compared and evaluated.

Many statistical methods are available to optimise a servo design. Here it was decided to adopt an approach that mixes the stochastic nature of a Monte-Carlo simulation with the more deterministic process of a particle swarm optimisation.

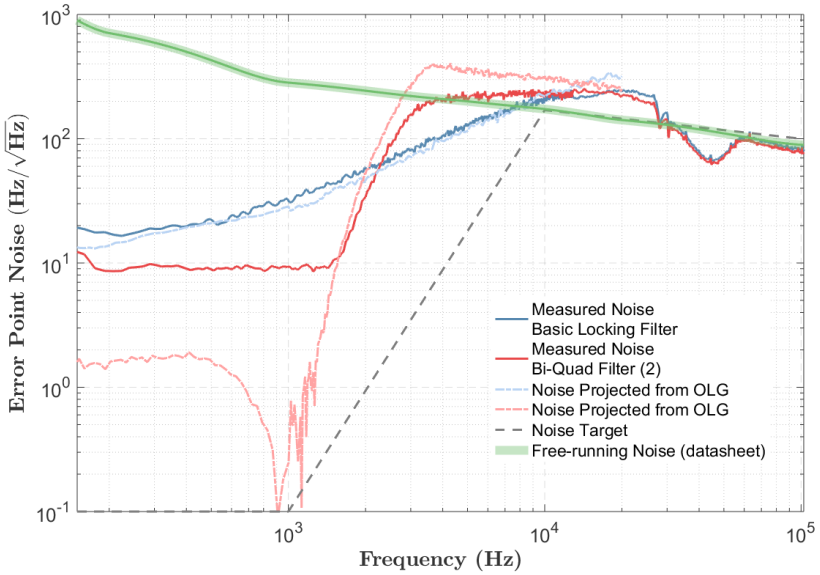


Figure 12.8: Measured laser noise with laser locked to reference cavity with basic digital filter (blue) and with two bi-quadratic filters (red) compared to the free-running noise of the laser (green). The solid traces show the measured in-loop noise, while the fainter dashed traces show the expected noise level based on the measured loop response.

This would allow the fine tuning of the filter parameters for complex filter design from a minimal starting definition of an optimal stable filter. The Monte-Carlo method provides a fast way of sampling a large parameter space while not getting locked into a local sub-optimal solution. Then the particle swarm optimisation allows the optimisation of a stochastic solution with greater local optimisation.

12.4.1 Optimisation Algorithm

A method of optimising the design of digital filters was developed to see if it was possible to balance the requirements of noise suppression and stability.

Stochastic filter generation

In the first step of the algorithm an array is created of randomised filter parameters where each row described a complete filter option. Each random filter option is then analysed and a numerical cost function is ascribed to each filter. The method for assigning a numerical cost function is based on the gain of the filter and the modelled stability of the resulting loop (this is described later in more detail). This can be used to compare different filters and choose the optimal solution.

For the filter design algorithm, it is desired to optimise the parameters of the

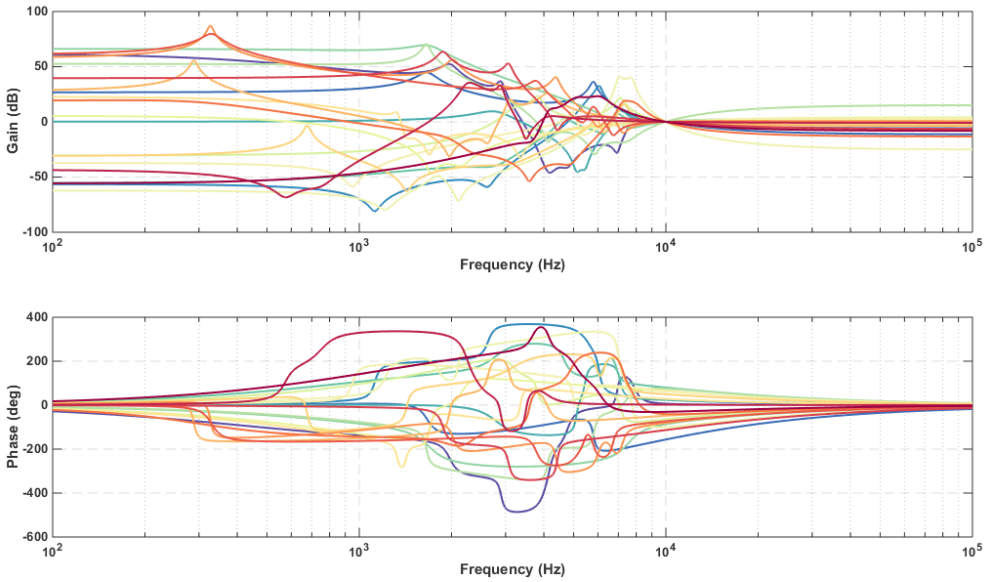


Figure 12.9: Bode plots of randomised filters. Here 20 example filters are shown with a specified unity gain point of 10 kHz. The frequency of the poles and zeros was limited to be between 100 Hz and the unity gain frequency. The Q-factors of the poles and zeros was randomly generated in the range $0.5 - 30$.

filters poles and zeros. Each pole or zero can be fully described by a frequency and a Q-factor. Alternatively, the poles and zeros can be described in complex form.

For the stochastic filter generation, the frequency and Q-factors of each individual pole and zero is a randomised value in a specified parameter range. For the Q-factors a lower limit of 0.5 was set to ensure the real part of the poles was always positive.

The final parameter needed to specify a servo filter is the overall gain factor. In this approach the overall gain is not a random factor to be optimised but is determined by the randomised combination of poles and zeros generated. For each filter the gain factor is set such that the filter crosses unity gain at some specified frequency. In this way the efficiency of the simulation is improved so that all filters generated have a unity gain crossing at the desired frequency. Other than these stipulations no other requirements are set on the initial filters generated in the simulation.

An example of the generation of randomised filters is shown in Figure 12.9. In this plot only 20 filters are plotted, in a full simulation of the order 10^5 filters would be simulated. Even in this small set of random filters there are a few that could be implemented in a servo and be stable. These therefore represent possible optimal solutions.

Particle swarm filter optimisation

A selection of optimal filter solutions identified from the stocastic filter generation are then taken as the starting point for a particle swarm optimisation.

A particle swarm optimisation works by generating a population of solutions (or particles) randomly distributed in a defined parameter space. Each particle is then analysed with a cost function based on each particles position in the parameter space. The best placed particle is identified and its position noted as an optimal location. In the next iteration of the optimisation each particle 'moves' towards the optimal position. The 'inertia' of each particle is determined by the relative cost function and the separation of the particles previous location and the identified optimal location. There is an element of randomness added to the particles 'velocity' vector. At each iteration the cost function for each particle is re-evaluated again and the best placed particle location is updated. In this way the population of particles converge to the optimal solution.

This method of optimisation is easy to visualise for a two or even three-dimensional problem, however the method is also valid for higher dimensional parameters spaces.

The particle swarm step of the filter optimisation algorithm takes a randomly generated filter that has already been established to represent a stable filter for frequency stabilisation and improve the performance of the filter. A population of particles is generated in a bounded parameter space centred on the pole and zero parameters of the filter. Each particle represents a similar filter with modified parameters, and over the course of many iterations the swarm converges on the pole and zero parameters that correspond to the smallest cost function as calculated by the metric ascribed.

An example of a particle swarm optimisation of a filter option is shown in Figure 12.10 with the inner workings of the alorithm represented in Figure 12.11. Here the simulation used only one pole and one zero complex conjugate pair. In Figures 12.12-12.15 the same is shown for a filter optimisation using four pole and zero pairs.

Figure of merit metric

Assigning a single numerical figure of merit to a filter can be a simple process but thought has to be given to both the desirable noise suppression and the stability of the loop. So in considering a single cost function for the optimisation algorithm one approach is to combine several figures of merit. Detailed here is a simple figure of merit metric consideration.

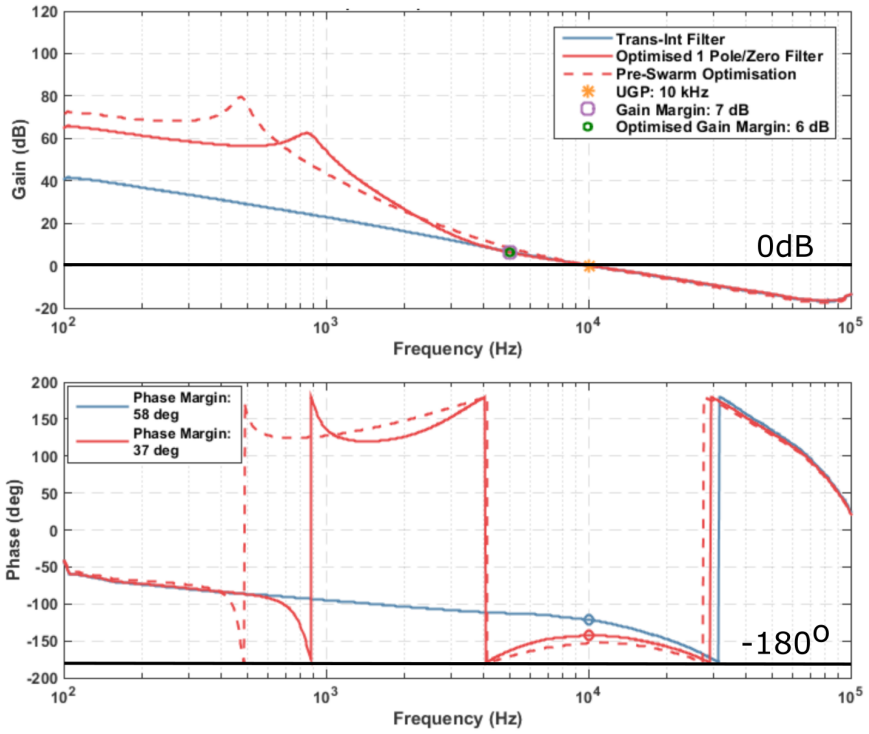


Figure 12.10: Example of an optimised filter generated for one pole and one zero complex conjugate pair. The randomly generated filter option is shown in the red dashed trace, the swarm optimisation improved the phase margin and increased the high gain corner frequency. This was considered a better solution by the alorithm, but with only one pole and one zero to change the optimisation was limited. The target filter was set to have a gain of 60 dB up to a corner of 1 kHz. The bode plot shows the simulated open loop response. The blue trace shows the measured loop response with a single transitional integrator, and the red traces show the simulated response if the designed filter was also implemented. The dashed transfer function represents the filter generated in the stochastic search, and the solid trace shows the final result of the particle swarm optimisation.

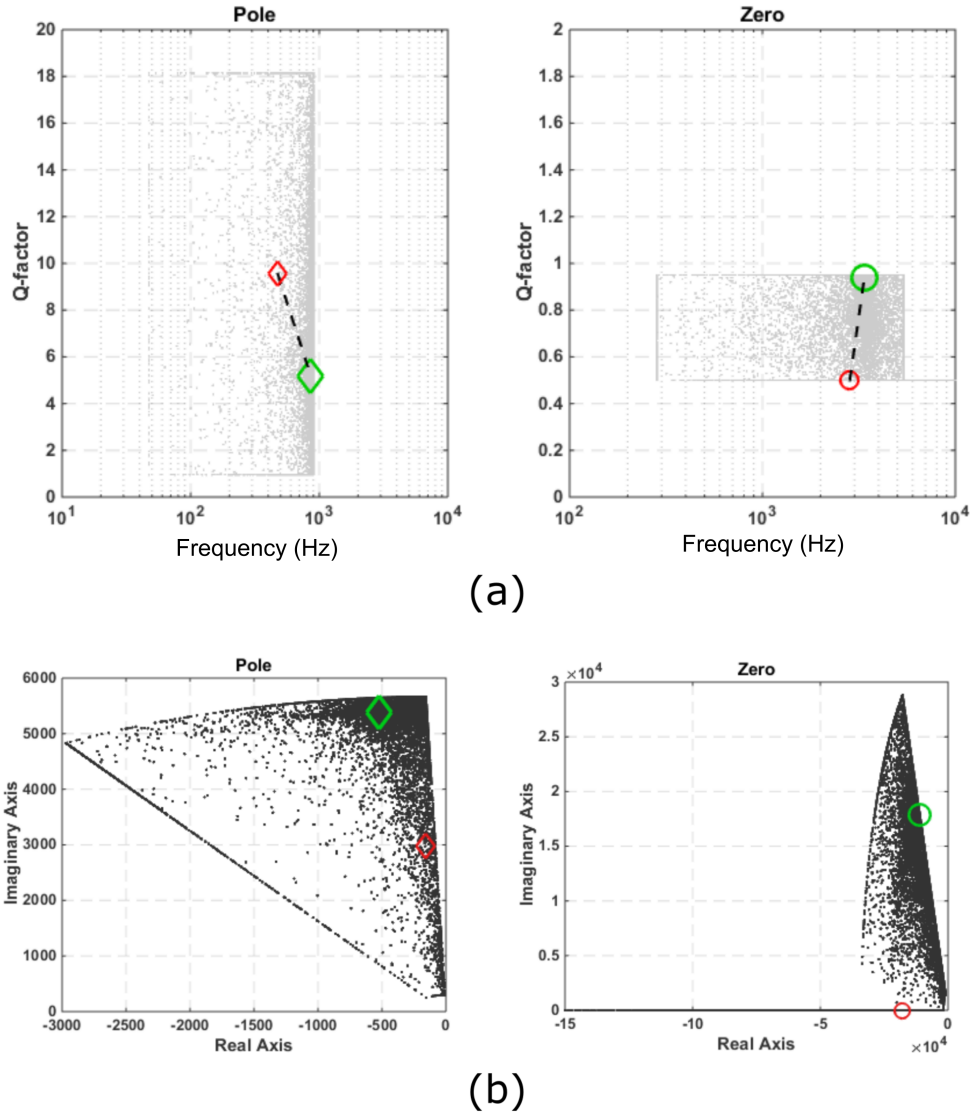


Figure 12.11: Example of an optimised filter generated in Figure 12.10 for one pole and one zero complex conjugate pair. The inner workings of the particle swarm are represented in the pole (diamonds) and zero (circles) maps in (a) and (b). The grey points in these plots show all the filter combinations tested in the search in the frequency/Q-factor parameter space (a), and the same is shown in (b) in the complex representation. The red and green markers show the initial and final pole and zero parameters respectively.

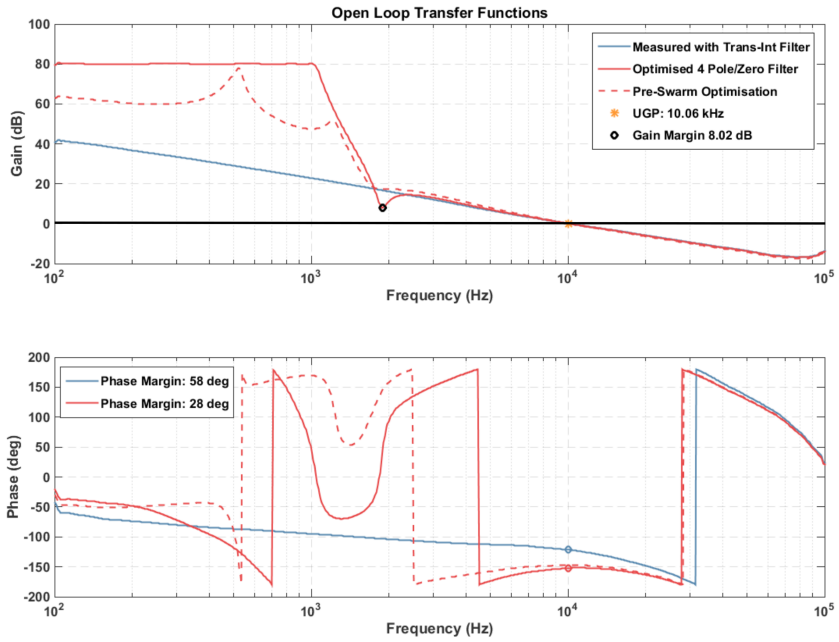


Figure 12.12: Example of an optimised filter generated for four pole and four zero complex conjugate pairs. The target filter was set to have a flat gain of 80 dB up to a corner of 1 kHz. The blue trace shows the measured loop response with a single transitional integrator, and the red traces show the simulated response if the designed filter was also implemented. The dashed transfer function represents the filter generated in the stochastic search, and the solid trace shows the final result of the particle swarm optimisation.

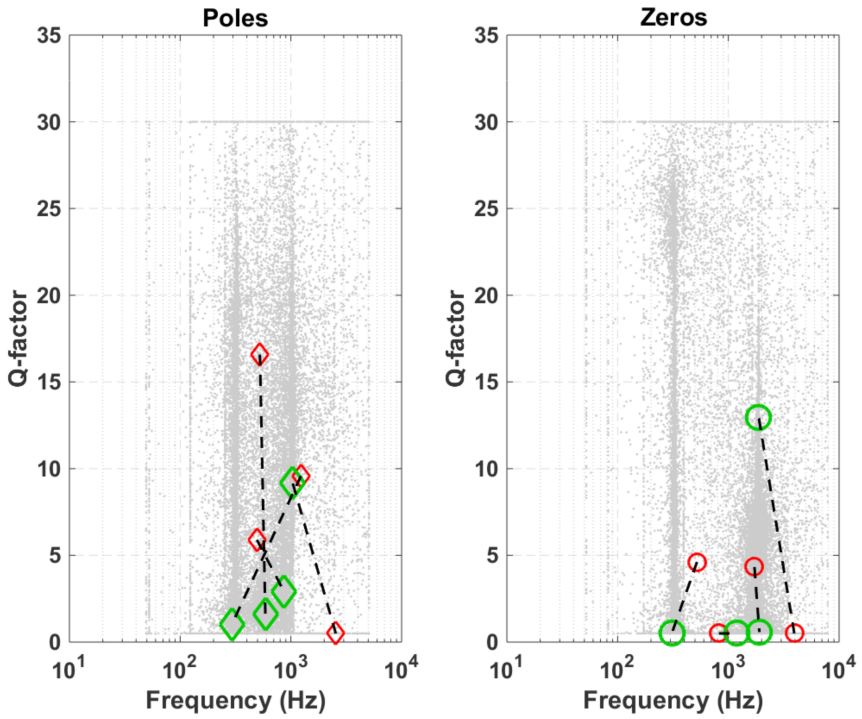


Figure 12.13: The inner workings of the particle swarm are represented in the pole (diamonds) and zero (circles) maps. The grey points show all the filter combinations tested in the search in the frequency/Q-factor parameter space. The red and green markers show the initial and final pole and zero parameters respectively.

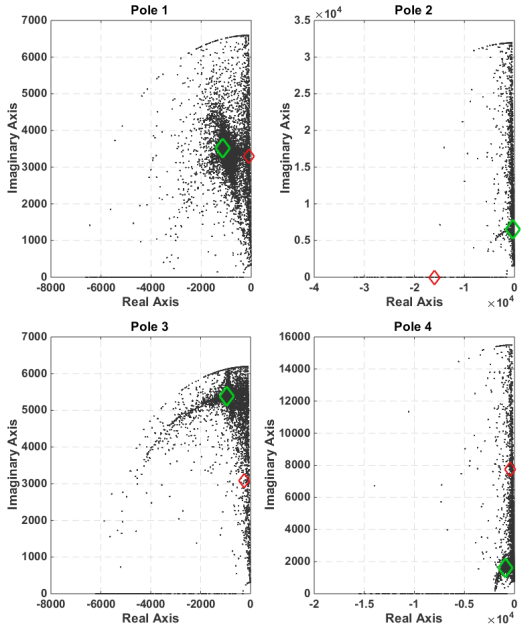


Figure 12.14: The inner workings of the particle swarm are represented in the pole maps. The scatter points show all the filter combinations tested in the search on the real and imaginary axes. The red and green markers show the initial and final pole and zero parameters respectively.

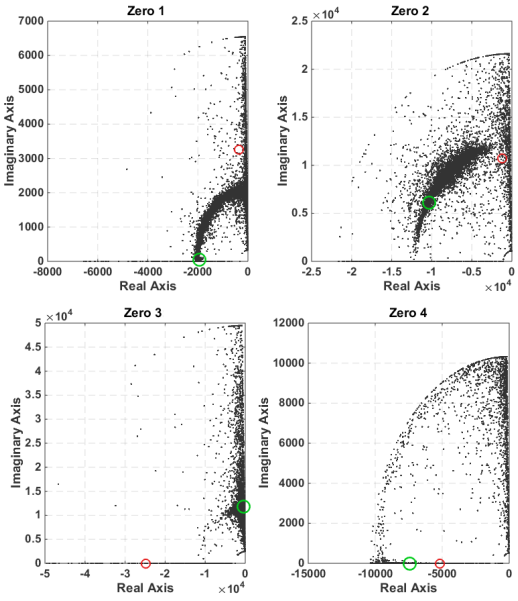


Figure 12.15: The inner workings of the particle swarm are represented in the zero maps. The scatter points show all the filter combinations tested in the search on the real and imaginary axes. The red and green markers show the initial and final pole and zero parameters respectively.

For this metric we consider the phase margin and gain margin as a measure of the stability of the loop, and the unity gain frequency, gain at low frequency and the noise suppression bandwidth as measures of the noise suppression.

The UGP is the frequency at which the open loop gain crosses unity (0 dB). Above the UGP the noise of the laser is not suppressed, below the unity gain point feedback to the laser frequency actuator suppresses the noise of the laser and holds it to the reference value. The gain of the loop determines by how much the free-running noise of the laser is reduced.

The unity gain frequency is often referred to as the bandwidth of the loop but we also want to consider the effective noise suppression bandwidth, that being the bandwidth over which a significant amount of noise suppression is achieved. Just below unity gain the gain is low so the noise of the laser is barely reduced. To add this consideration to the metric we define a corner frequency of the filter (f_c) to be the effective noise suppression bandwidth and we consider the gain at this frequency as one of the figures of merit of the filter.

To consider multiple figures of merit in a single metric each must be weighted according to what outcome is most desirable, a higher noise suppression performance or a greater loop stability.

In addition, for each figure of merit there is a point beyond which increasing that parameter is of no use. For the gain margin there is an upper limit beyond which we no longer consider an increase in this figure of merit to be an improvement to the filter. A loop will become difficult to use if the gain margin is below 3 dB but a gain margin of 10 dB is no more stable then a gain margin of 6 dB.

Similarly for the gain parameters, there is a limit above which higher gain is not useful as the level to which the loop can stabilise the laser is limited by the length noise of the reference cavity which is a finite level determined either by shot noise of the light at the sensing photodiode or thermal noise of the cavity. Higher gain can be applied in-loop but the resultant noise performance of the laser is still limited by the noise of the cavity.

In addition to the figures of merit above, an additional metric is implemented to encourage the algorithm to optimise for flat open loop gain profile across the band of interest. The band of interest being between 0 Hz and the specified corner frequency, f_c. The flatness of the open loop gain is calculated as the sum of the squared difference between the simulated function and the target gain.

The multiple figures of merit are combined in a simple expression as

$$\xi_{combined} = \frac{\xi_{PM} + \xi_{GM} + \xi_{Gain(f_{100})} + \xi_{Gain(f_c)}}{\xi_{flatness}}. \tag{12.4}$$

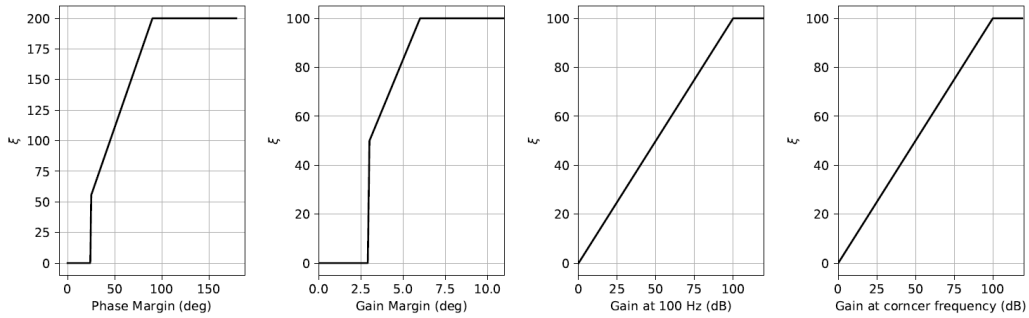


Figure 12.16: Weighted figures of merit (ξ) used in the filter optimisation metric showing the upper and lower limits of the four parameters and the related figures of merit.

Figure 12.16 shows an example of the weighting function of the four figures of merit combined in the numerator of Equation 12.4.

In addition to the scaling, a condition is applied such that is any of the individual figures of merit are equal to zero then the combined figure of merit is also set to zero independent of the other figures of merit. This prevents any filter being considered optimal that is in fact unstable.

Linear regression metric

An alternative method to assign a quality factor to a filter is to define a target transfer function and see how close the simulated filter matches the desired outcome. In this case the cost function is calculated using a linear regression statistical method.

To avoid over-engineering a solution the target transfer function should be as simple as possible or the optimisation algorithm loses some of its utility.

This algorithm can be used in cases where the frequency response of the loop is more complicated than the one considered for the 1550 nm laser system. If there are features in the loop response such as resonance peaks of notch-like structures, these can be compensated for by constructing more complicated filters using more poles and zeros. The algorithm can take a measured open loop gain as the starting point so this information of the response of the loop is taken into consideration.

If the figure of merit metric is used, then only features below the specified corner frequency will be compensated for as the optimisation looks to make the gain flat across this band. For the linear regression metric then any feature in the measured open loop response can be compensated for in the optimisation process. If enough poles and zeros are used and enough stochastic simulations are preformed, then a filter can be found that will produce the desired open loop gain

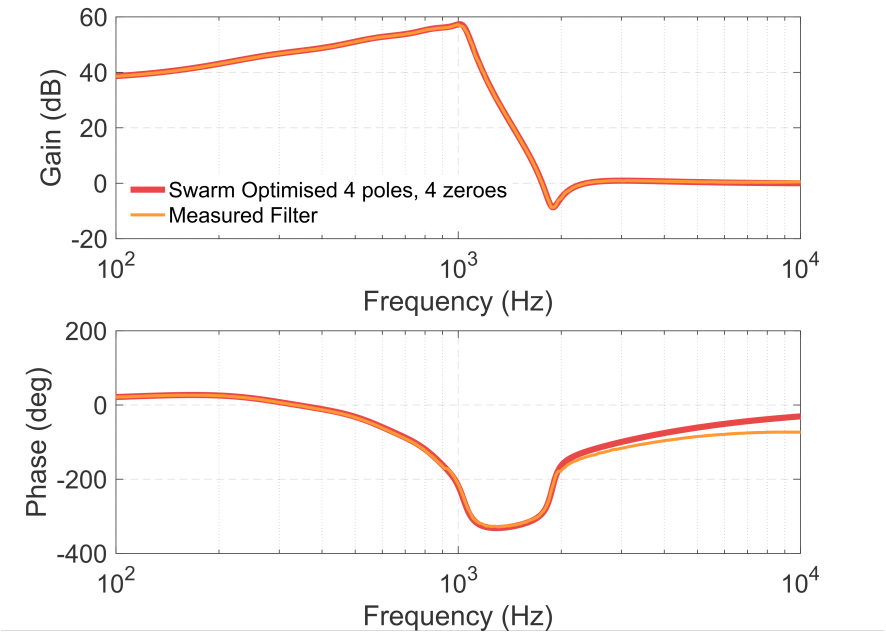


Figure 12.17: Measured transfer function (orange) of optimised filter compared to modelled result (red). The filter was implemented using four second order sections in the dSPACE digital processor and measured from input to output with anti-aliasing and anti-imaging filters which produces the phase discrepancy between measurement and model above 2 kHz.

for any measured loop response.

12.4.2 Results

Figure 12.17 shows the measured transfer function from the input to output of dSPACE with the optimised filter design from Figure 12.12 implemented in the digital regime. This filter, when combined with the loop gain will in principle result in the specified target gain of 80 dB being achieved up to the corner frequency of 1 kHz.

This is just one option produced by the filter optimisation algorithm but shows that arbitrary filters can be conceived of and implemented in this system. The complexity of filters that can be implemented is limited by the computational speed of the digital processor. With the digital processor running at a sample rate of 166 kHz four five second order sections could be used in series.

The key outcomes of this Chapter are summarised in Chapter 13.

Chapter 13

Conclusion

13.1 Summary of LIGO Laser Work

In Chapter 9 it was shown how the addition of a new laser amplifier in the Advanced LIGO laser system doubled the amount of laser power available to the interferometer between the second and third observing runs of the advanced gravitational wave detector network. The work to integrate this laser amplifier in the existing system was presented. For the third observing run (that is in progress at the time of writing) the LIGO detectors were operated with 40 W of laser power injected into the interferometer. Along with other upgrades to the instrument, this contributed to an improved sensitivity and greater detection rate compared to the second observing run.

Two investigations relating to the laser stabilisation system of the LIGO Livingston detector that were conducted as part of this thesis work were presented in Chapter 10.

The frequency stabilisation system was found to be limited at high frequencies by the bandwidth of the electronic servo that was deployed in the LIGO Livingston laser stabilisation system.

The operation of the LIGO Livingston detector was found to be directly impacted by regular dropouts in the actuator signal of the intensity stabilisation system. These dropouts caused the full interferometer to lose lock and were shown to be related to the stability of the control loop employed to stabilise the intensity of the laser. By improving the loop stability these dropouts were mitigated and no longer impacted the running of the detector.

13.2 Summary of 1550 nm Laser Work

A frequency stabilisation system for a 1550 nm laser was developed and while this system was plagued by technical issues, it was possible to reduce the free-running laser noise to the target noise level of $0.1\,\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ between 100 Hz and 1 kHz.

When compared to a 1064 nm NPRO laser that is the current standard for gravitational wave interferometry, this particular 1550 nm diode laser has a significantly higher noise level of a shape that results in a significant amount of noise at high frequencies. This presented a range of challenges and required a significant technical effort to reduce the noise to a suitable level for high-precision interferometry experiments.

Results shown in Chapters 11 and 12 showed that stabilisation of the frequency of this 1550 nm laser was possible but challenging and as such shows that this particular laser is unsuitable for future gravitational wave interferometry experiments.

In Chapter 12 the use of a digital processor for the frequency stabilisation of the 1550 nm laser was investigated. This work was successful in showing the usefulness and limitations of such a scheme. While digital filtering offers great flexibility, the computational limitations of the processor limited the performance of this system.

The mechanical resonances of the 1550 nm laser around the specified unity gain frequency of the digital servo system impacted the practicality of the digital process for this application.

The flexibility of the digital servo was exploited with an investigation of filter optimisation and design automation. This was successful in developing a method of optimising the design of the stabilisation servo that manipulated the poles and zeros of the filter in a way native to digital implementations of signal filters.

Future Work

Future work to stabilise the frequency of the 1550 nm laser system will combine the digital filtering with the high-speed analogue stabilisation system that was successfully developed and demonstrated. This would allow for stable operation with filter tunability.

Additionally, faster digital processors will be investigated as alternatives to the dSPACE controller. High speed FPGA processors are available and would allow for feedback up to higher frequencies that would circumvent the laser PZT instability issues that limited the practicability of the digital servo system deployed here.

Part IV

Appendices

Appendix A

Hardware for Linear Cavity Experiment

A.1 Vacuum System

The linear cavity experiment was housed in the Glasgow 10 m vacuum system. The vacuum system is comprised of nine vacuum systems in a 'W' layout. The polarisation cavity experiment used one arm of the system and the central vertex.

The system includes gate valves between the two 10 m arms and the central vacuum volume. Each section of the system can be isolated and individually pumped down to vacuum.

The prototype uses a system of two types of vacuum pumps, scroll pumps are used to pump the system in from air down to 10^{-1} mbar. To get to lower vacuum pressures a combination of scroll and turbo pumps are used. The system can be pumped down to a level of 10^{-5} mbar.

A.2 Seismic Isolation

The optics inside the vacuum system are isolated from external seismic noise with a simple passive isolation system built up from multiple layers of rubber and metal stacks.

The top layer is a 60×60 cm optical breadboard which is mount in a square ring of aluminium extrusion. This square ring sits on top of a hexagonal ring of thicker extrusion with a layer of soft rubber between the two metal layers. The hexagonal ring sits on top of the three tank feet, with a further layer of sort rubber between the hexagonal ring and the pedestals mounted on the tank feet.

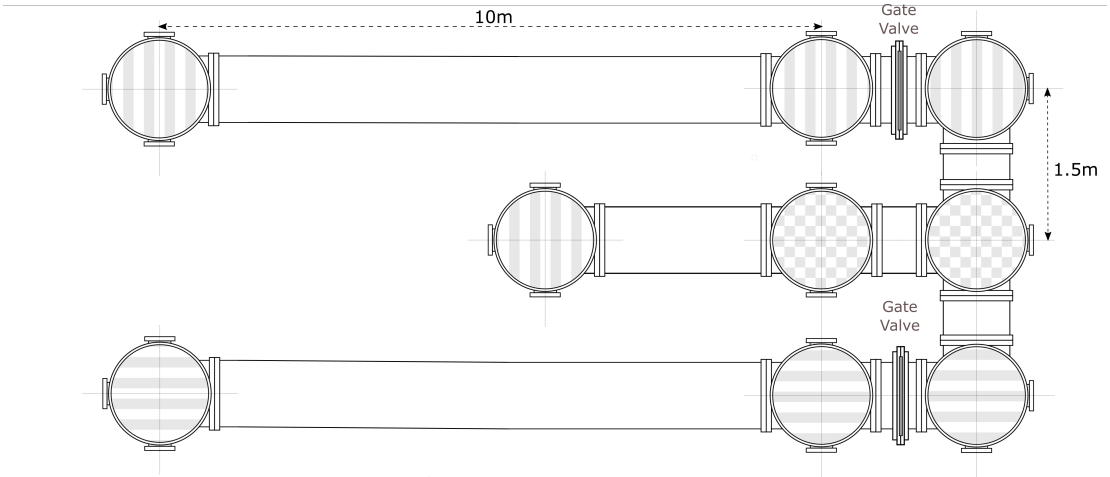


Figure A.1: Birds-eye view of the Glasgow 10 m prototype vacuum system. The polarisation cavity experiment used the tanks indicated with horizontal and square shading.

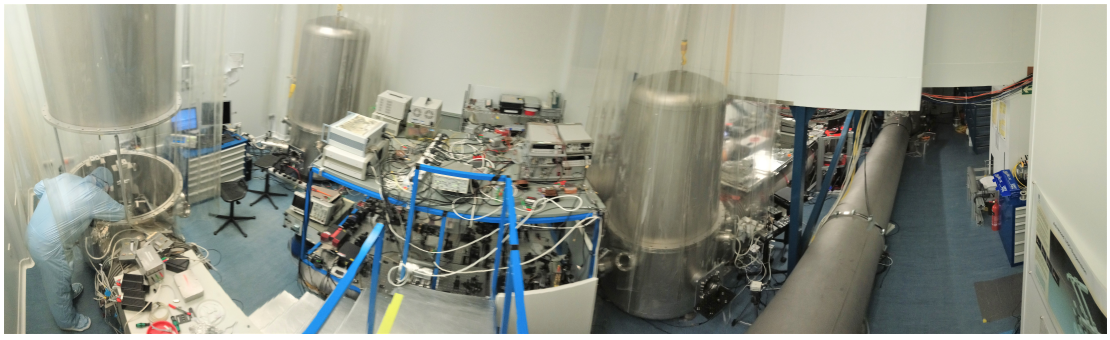


Figure A.2: Panoramic view of the Glasgow 10 m prototype vacuum system. The open tank on the left is the ETM tank of the polarisation Sagnac experiment, and the beam tube seen on the right is the arm of the 10 m cavity.

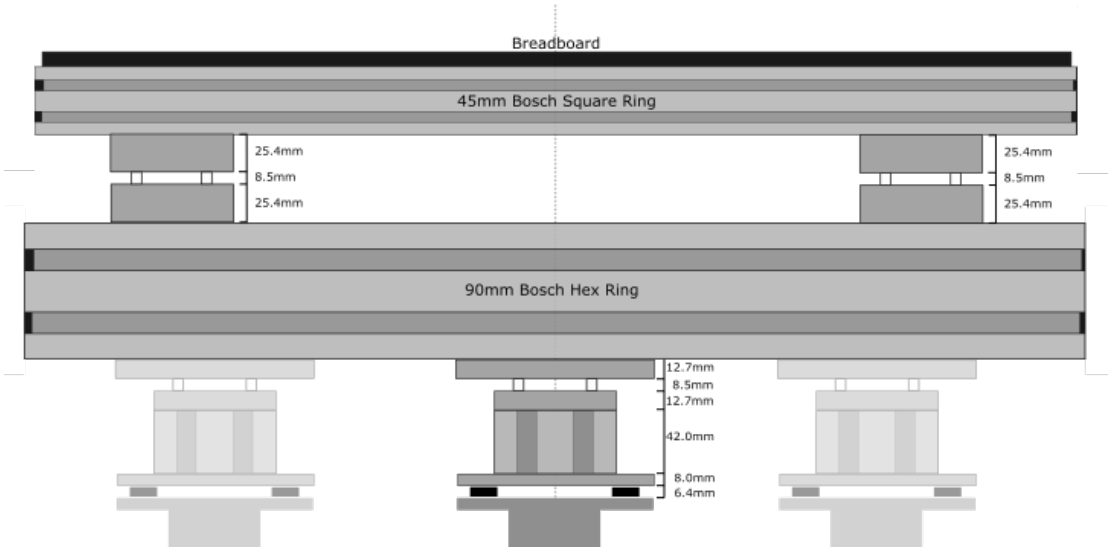


Figure A.3: Side-view of seismic isolation stack for in-vacuum breadboards for the Glasgow 10 m prototype constructed from aluminium extrusions from the Bosch company.

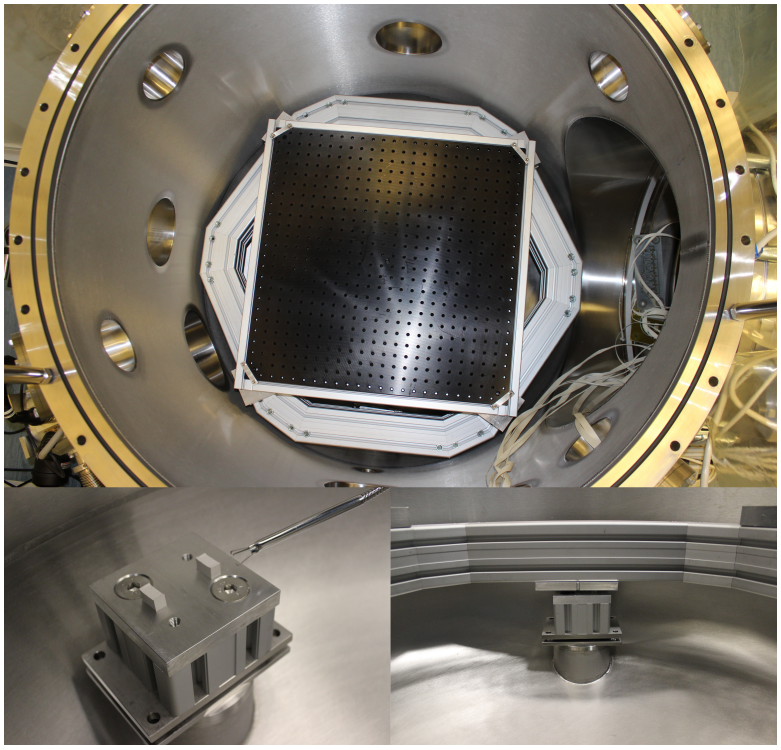


Figure A.4: Picture of the seismic isolation of the optical breadboard. Top image shows birds-eye view of the installed breadboard and isolation system in the ETM tank. Bottom left shows rubber and metal stack on the tank feet. Bottom right shows the hexagonal ring sitting on the bottom stack of rubber and metal.

A.3 Suspensions

The suspensions used in this experimental work were constructed to an existing design, details of which can be found in [98]. This suspension design has a compact footprint of 20x20cm tailored to the Glasgow 10 m prototype system so that multiple suspensions can be installed in each vacuum tank. The suspension comprises an aluminium frame, a steel anchor plate sitting atop the frame, isolated with a layer of rubber, and a double pendulum steel wire suspension hanging from the anchor plate. This suspension design is shown in Figure A.5.

A.3.1 Suspension Modes

The six modes of motion (or degrees of freedom) of a suspended optic are shown in Figure A.6. This represents the three translational and three rotational directions in which the suspended optic can move. The terminology used here is based on the standard description of the principal axes of aircraft motion, where the rotational degrees of freedom are called the pitch, roll and yaw.

A.3.2 Suspension Control

Coil-magnet actuators can be added to upper suspension stage or the lower test mass stage. For the intermediate mass stage coils of 80 turns are wound onto copper cup shaped coil formers that surround the magnets on the suspended masses. These provide an increased damping force and drive as the magnitude of the generated magnetic field is greater. This stage also includes additional passive copper dampers above the intermediate mass to damp additional motion of the mass.

The coil magnets pairs are arranged in a diamond formation on the rear of the mass so that the position of the mass can be modulated in the pitch, yaw or longitudinal degrees of freedom. The cylindrical magnets used on the upper suspension stage had a diameter of 3 mm and a length of 6 mm.

In addition, the coils provide passive damping of the suspended masses displacement due to eddy current damping. If the magnet moves relative to the coils, an eddy current is generated in the copper that results in a magnetic field that push the magnet in the opposite direction of motion cancelling out and damping the displacement of the mass.

Actuation of the upper suspension stage is referred to as local control and is used for alignment of the experimental system.

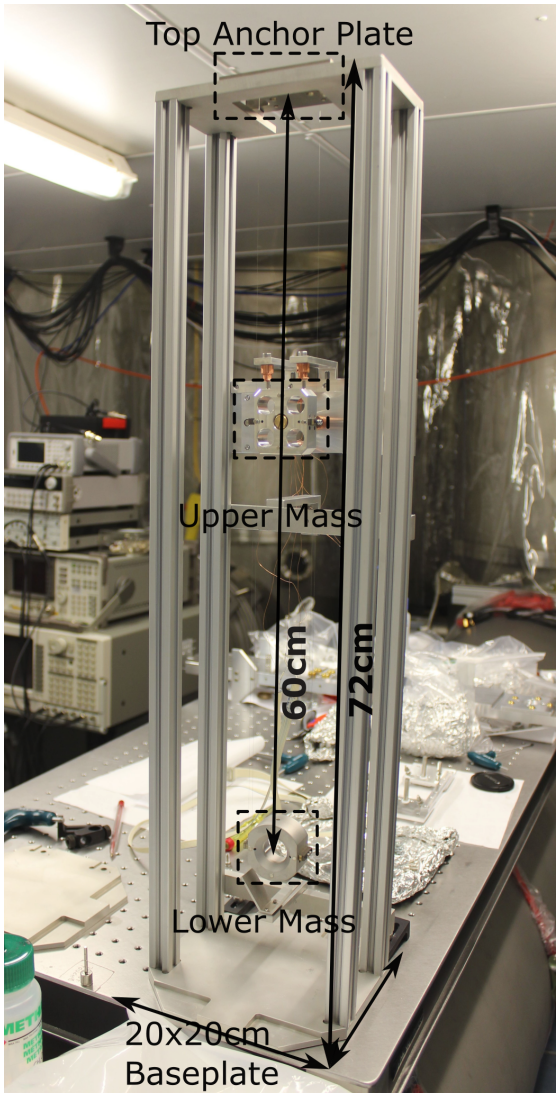


Figure A.5: Picture of the suspension during construction. The suspension frame is made from four aluminium extrusion struts between the baseplate and the top plate. The suspension is hung from a steel anchor plate that sits on rubber to isolate it from the top plate. The double suspension includes the upper and lower masses hanging from steel wire.

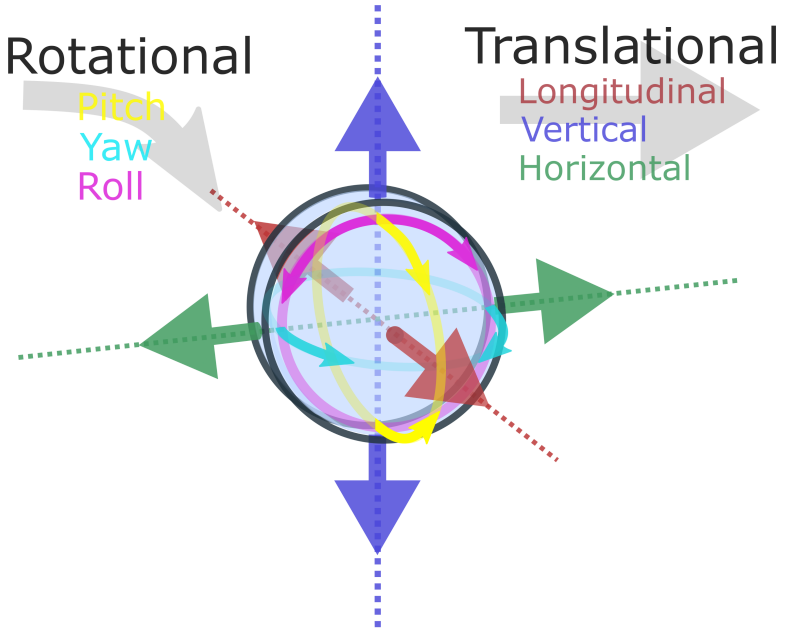


Figure A.6: Modes of motion of a suspended optic. The optic can move in three translational degrees of freedom and three rotational degrees of freedom.

On the lower (test mass) stage the coils of 30 turns are wound onto plastic coil formers and smaller magnets of 3 mm diameter and 1 mm length were used. Less actuation strength is needed to drive the lower stage of the pendulum.

Direct actuation of the test mass position is used for cavity locking and signal injection and is referred to as global control. Only the cavity mirror suspensions were built with global control hardware. The signals driving the coils are generated from the CDS digital control system.

The intermediate mass stage is used for so called local control to set the position of each suspended optic individually and is used for alignment. The local control signals are generated by Arduino controllers that communicate with a local server with a network graphical user interface. From this system the alignment of all suspensions on the system can be controlled.

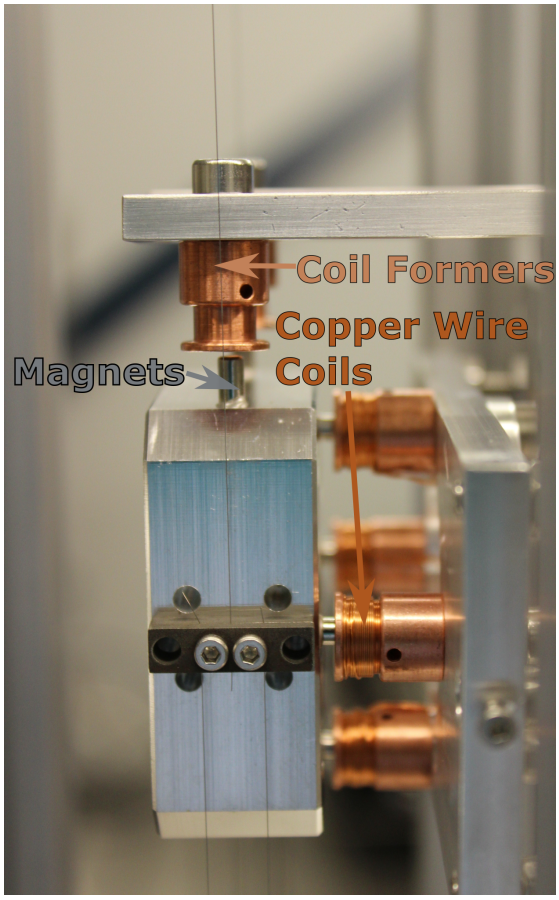


Figure A.7: Picture of the suspension actuation at the upper stage. Current is passed through the copper coils that are wound onto the coil formers that are mounted on the suspension frame. This generates a magnetic field that applies a force to the permanent magnets that are fixed to the suspended mass. The copper coils formers also provide eddy current damping of the suspended mass. At this upper stage six copper former and magnet pairs are utilised to damp all six modes of motion. Coils are included on the four formers behind the mass so that the longitudinal, yaw and pitch modes of motion can be actuated.

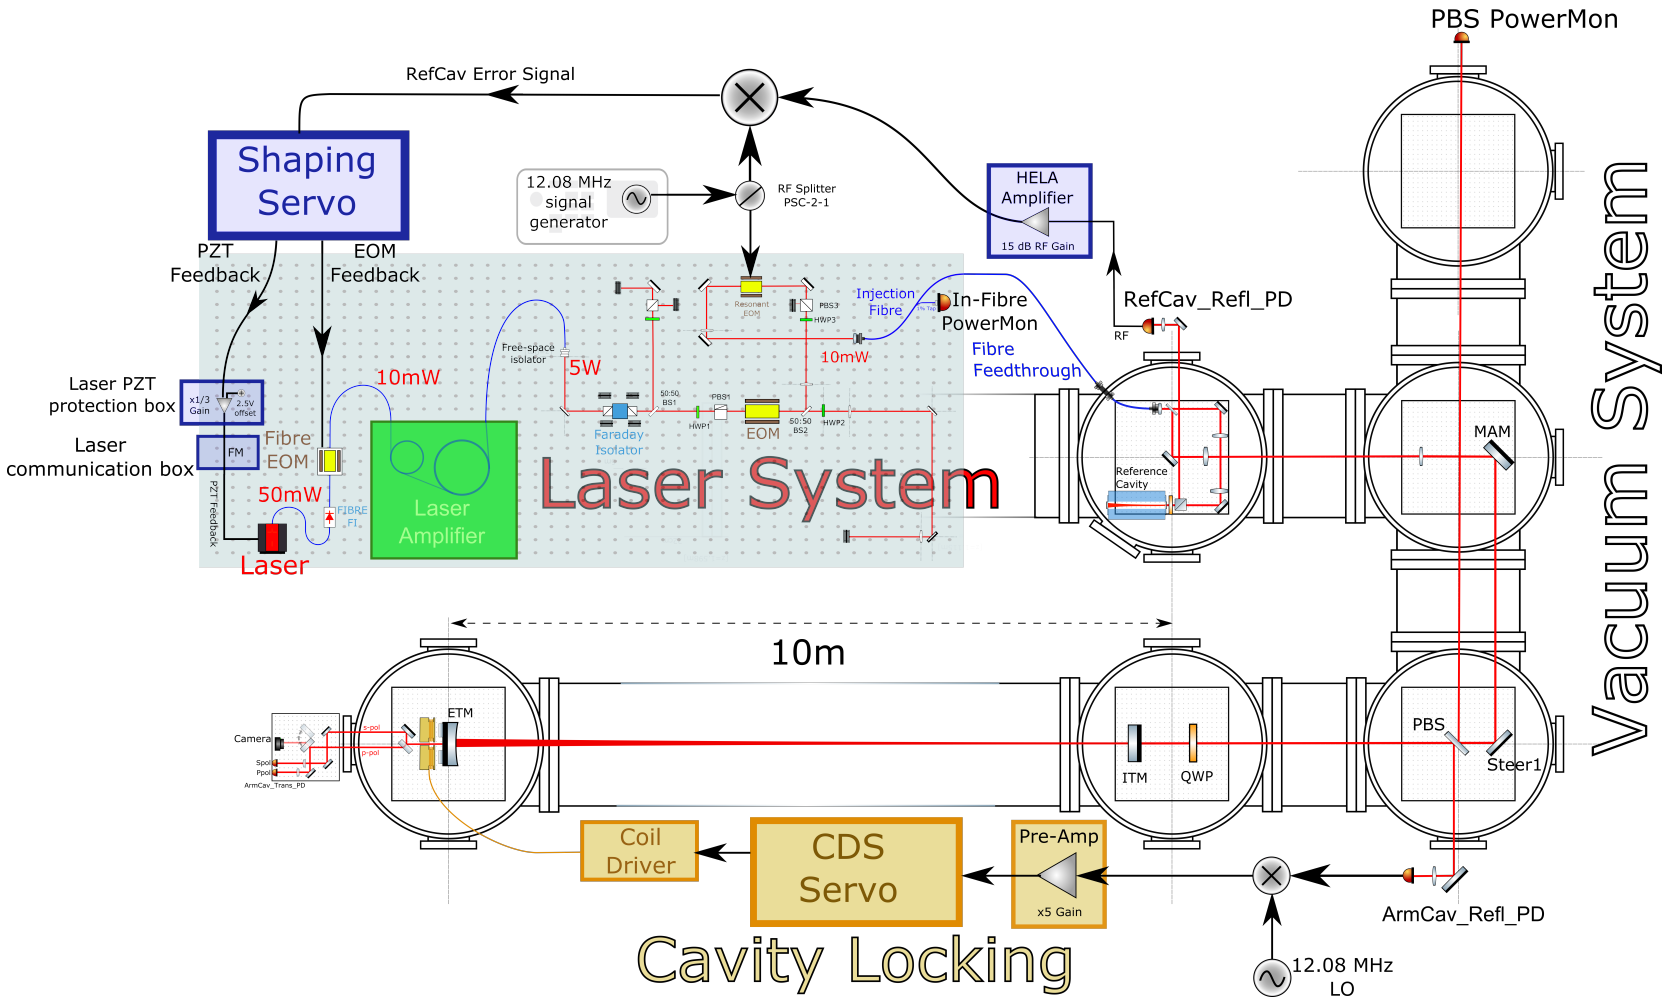


Figure A.8: Full layout of the 1550 nm laser stabilisation and 10 m suspended cavity with polarising optics experimental set-up.

Appendix B

Electronics for the 1550 nm Laser System

B.1 Cavity PZT Driver

A dedicated driver is needed to increase the voltage of the signal driven onto the PZT device that actuates the length of the reference cavity. The circuit diagram of this driver is shown in Figure B.1.

The OP454 driver can output signals of 100 V. A DC-DC voltage converter was used to generate a supply voltage of +88 V for the OP454. To avoid being able to drive large negative voltages onto the PZT that could damage the device, the driver is supplied with -12 V on the negative power supply line.

The circuit includes two stages of low pass filtering, with one pole at 1.6 MHz and 10 Hz. The PZT driver is used as part of the frequency stabilisation servo to feedback signals to the cavity to compensate for slow drifts in the laser frequency, and keep the laser resonant with the reference cavity (see Section 11.2.2 for more details).

The cavity locking feedback is designed to crossover with the laser stabilisation around 1 Hz. The 10 Hz pole in the driver box along with the digital filters implemented before the driver box, ensures that the response of the cavity locking loop falls of rapidly above 10 Hz.

B.2 Laser PZT Driver

The 1550 nm laser has an internal laser frequency actuator. Signals driven onto the FM channel of the laser communication box controls a PZT device that controls the length of the internal lasing cavity of the laser.

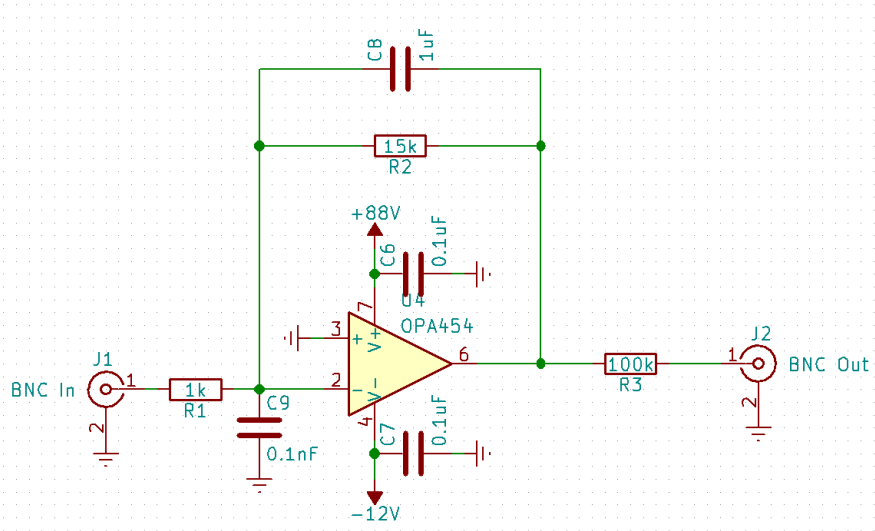


Figure B.1: Circuit diagram of the PZT driver circuit designed to send signals to the reference cavity PZT for length actuation.

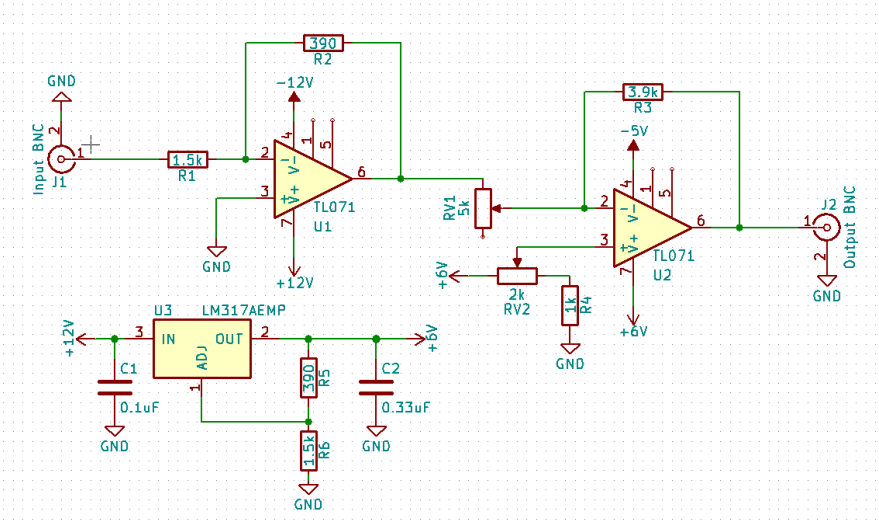


Figure B.2: Circuit diagram of the PZT driver circuit designed to send signals to the internal frequency actuator of the 1550 nm laser.

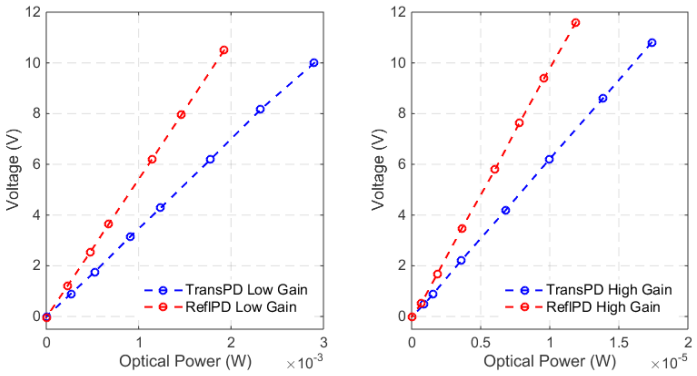


Figure B.3: Calibration of characterisation PDs in transmission (blue) and reflection (red) of test optic for the high and low gain transimpedance settings.

This channel can take voltages in the range of 0-5 V. The output of the stabilisation servo has a range of ±10 V this is reduced to ±2.5 V by the gain the first stage of the circuit (shown in Figure B.2). An offset is added to ensure the signal applied to the laser is in the range of 0-5 V.

B.3 Photodiodes

B.3.1 Polarisation Optics Characterisation Photodiodes

For the characterisation of the polarisation reported on in Chapter 3 a pair of 2 mm InGaAs biased photodiodes were used. These photodiode packages are sold by Thorlabs [162]. The output of these photodiodes were amplified by a transimpedance circuit that had high and low gain settings. The calibration of these photodiodes is shown in Figure B.3.

B.3.2 Reference Cavity Locking Photodiode

The photodiode monitoring the reflection path from the reference cavity is used to produce the reflected DC signal used as the trigger for the auto-locking system. The photodiode also includes an AC coupled output that is demodulated and used to generate the locking error signal for the laser frequency stabilisation system; the circuit diagram of this photodiode is shown in Figure B.4.

B.3.3 10 m Cavity Locking Photodiode

The circuit diagram of the photodiode used to lock the 10 m suspended cavity is shown in Figure B.5. This photodiode is placed in the reflection path of the 10 m

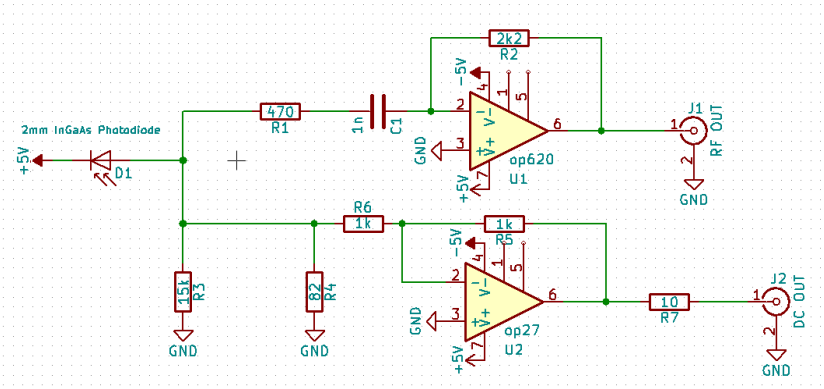


Figure B.4: Circuit diagram of the reference cavity locking photodiode.

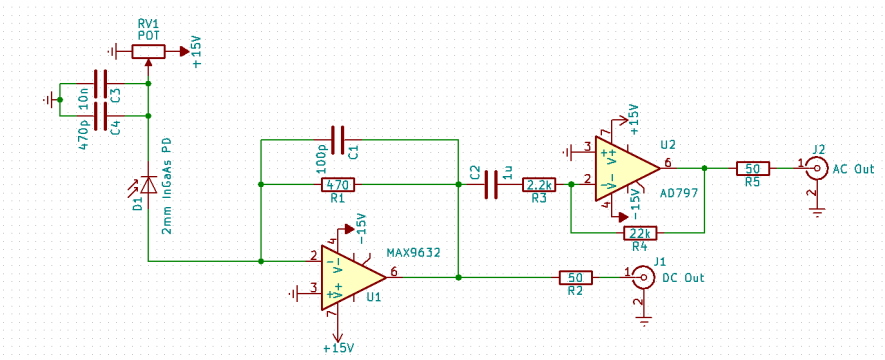


Figure B.5: Circuit diagram of the photodiode in the reflection path of the 10 m suspended cavity. AC coupled output is used to generate the cavity locking error signal.

cavity and includes DC and AC coupled outputs. The AC coupled output signal is demodulated to generate the cavity locking error signal. The DC coupled signal is also recorded to monitor the locked state of the cavity.

The circuit diagram of the cavity locking photodiode is shown in Figure B.5.

B.3.4 10 m Cavity Transmittion Photodiode

Two photodiodes of identical design (see Figure B.6) were used to measure the optical power transmitted by the 10 m cavity in orthogonal polarisations.

The calibration of these photodiodes is shown in Figure B.7.

B.4 Notch Filter

A resonant notch filter was used in the frequency stabilisation scheme after the mixer to remove the RF signals before the analogue servo box.

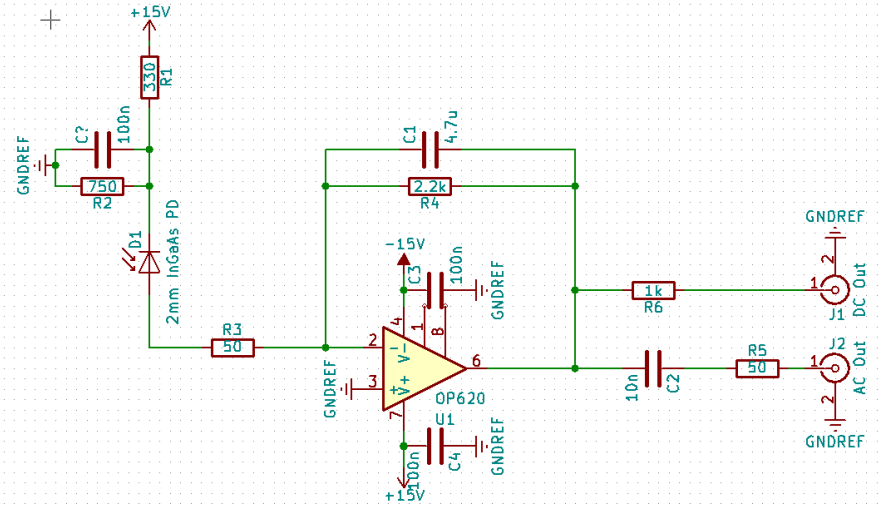


Figure B.6: Circuit diagram of the photodiodes in the transmission path of the 10 m suspended cavity.

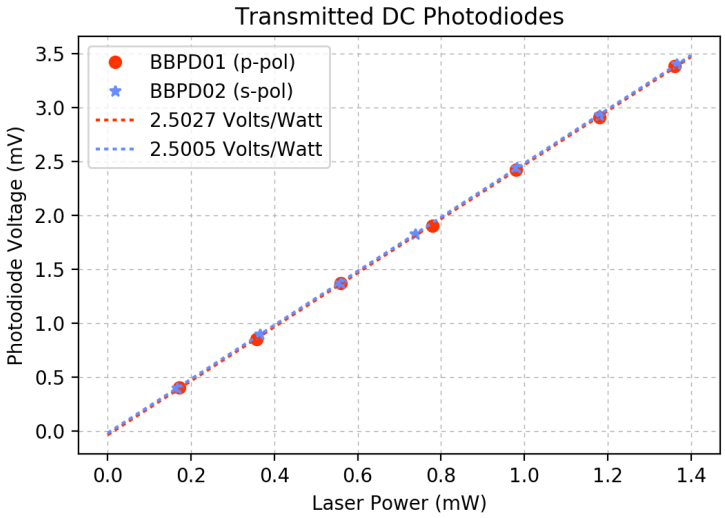


Figure B.7: Calibration of the two photodiodes used in transmission of the 10 m cavity as part of the ellipticity readout scheme.

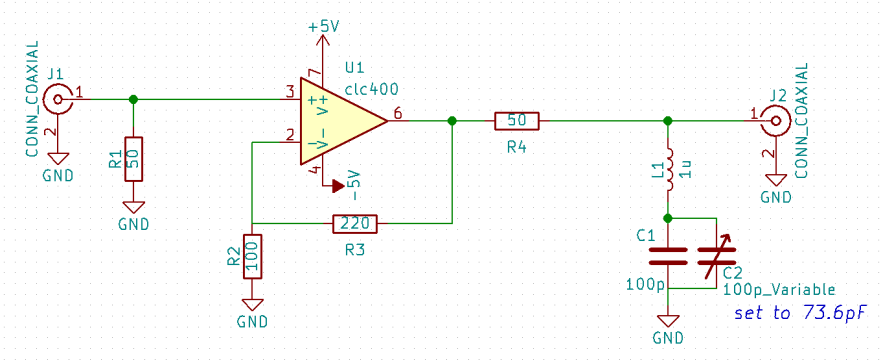


Figure B.8: Circuit diagram of the notch filter design to filter out the 12.08 MHz local oscillator signal from the 1550 nm frequency stabilisation system.

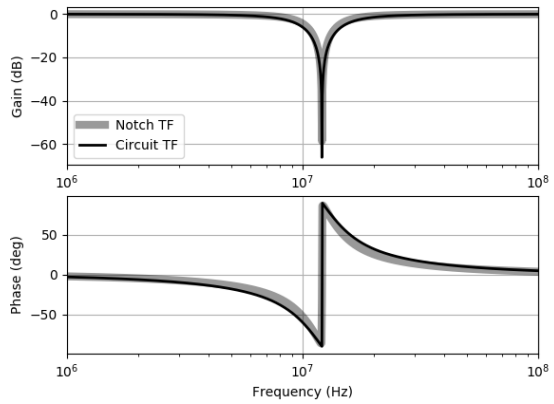


Figure B.9: Simulated response of notch filte designed for a 50 Ω load impedance. The grey trace is the simulated notch filter at 12.08 MHz and the black trace is the modelled response of the circuit design presented in Figure B.8

The circuit diagram of this filter is shown in Figure B.8 and the simulated response in Figure B.9.

B.5 AA/AI Filter

To eliminate aliasing and imaging effects from the dSPACE digital processor, filters were designed to be used at the input and output in the servo path to add strong notches at the sample frequency and to attenuate any signal at frequencies above the sample rate. The design of such a filter is presented in Figure B.10.

A bi-quadratic filter can be formed of a resonant peak and a notch, in a low pass arrangement such as the one used here the notch is at a higher frequency then the resonant peak. Such a filter can be realised with analogue electronics in a circuit such at the one shown in Figure B.10(a) [163].

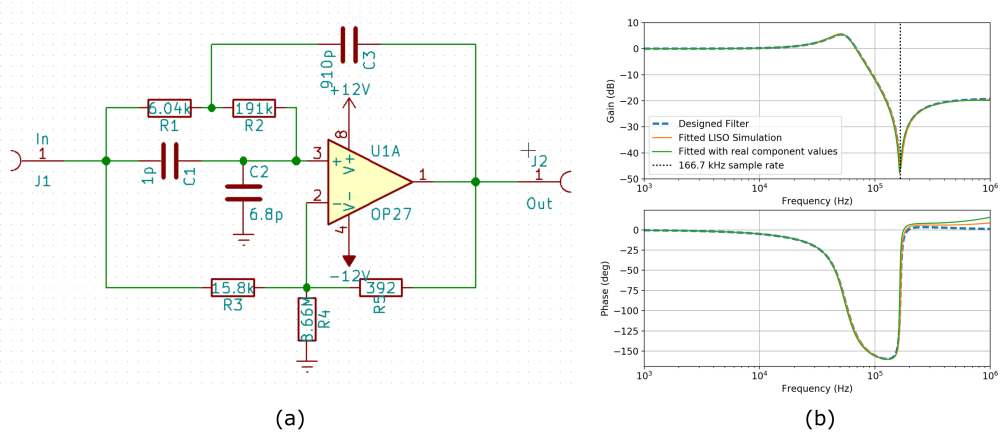


Figure B.10: Design of an anti-imaging/anti-aliasing filter for the dSPACE digital process with a sample rate of 166.7 kHz. The filter is a low pass bi-quadratic filter with a peak at 55.6 kHz and a notch at 166.7 kHz. The circuit diagram of the filter is shown in (a) with all component values shown. The simulated transfer function of this circuit is shown in (b). The LISO simulation package was used to design the circuit. The filter transfer function wanted is shown in the blue trace on (b), a LISO code was then run to optimise the circuit shown in (a) to match the design transfer function by changing the component values. The green curve then shows the transfer function with the best fit values that represent real components with values actually available. These 'fitted real' values are the component values shown in (a).

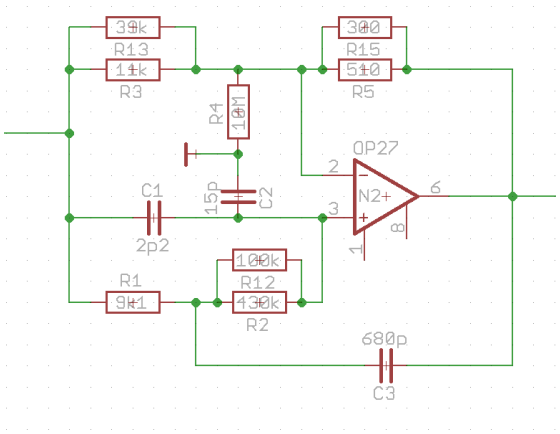


Figure B.11: Circuit diagram of the filter with the used component values. Parallel combinations of resistors were used to produce values of resistance closer to the desired values than the available component values allowed for.

The depth of the notch is set to be -50 dB to provide strong attenuation at the sample frequency of the digital process. The q-factor of the peak is kept low to prevent amplification of any other features of the loop, but cannot be set too low as this would significantly effect the behaviour of the loop at 10 kHz, the chosen unity gain point of the PZT actuated frequency stabilisation loop. The q-factor of the peak is tuned so that the phase lag at 10 kHz is less than 5°.

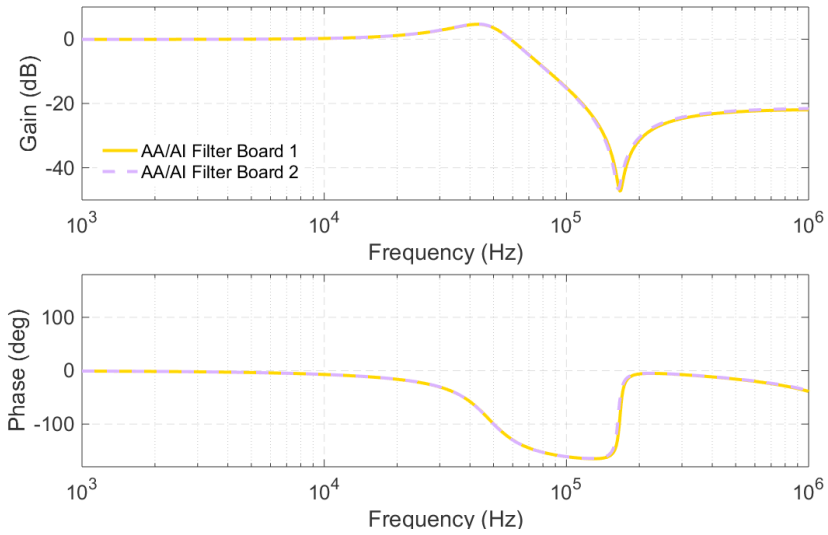


Figure B.12: Measured transfer function of the anti-aliasing/anti-imaging filter designed for the digital feedback system using the circuit design presented in Figure B.11

.

Appendix C

Jones Calculus

C.1 Jones Calculus for Optics and Light

C.1.1 Light Fields

Polarised optical fields can be described by Jones vector. For linearly polarised light the vector is

$$\begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}, \tag{C.1}$$

with θ being the orientation of the polarisation relative to the horizontal axis. An elliptically polarised field is described by the vector

$$\begin{bmatrix} \cos\theta \exp(i\phi_x) \\ \sin\theta \exp(i\phi_y) \end{bmatrix}, \tag{C.2}$$

with $\phi = \phi_x - \phi_y$ being the phase difference between orthogonal polarisations.

The commonly considered polarisation states and the corresponding Jones vectors are contained in Table C.1.

C.1.2 Optics

Polarising optics can either be *polarisers* or *waveplates*. A polariser has a given reflectivity and transmissivity that is dependent on the polarisation of the light interacting with the optic, whereas waveplates change the polarisation of the light by introducing a phase difference between light polarised in orthogonal axes.

Polarisation	Jones Vector
p-polarisation	$\begin{bmatrix}1\\0\end{bmatrix}$
s-polarisation	$\begin{bmatrix}0\\1\end{bmatrix}$
$\pm 45°$ linear polarisation	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\\pm 1\end{bmatrix}$
left-hand circular polarisation	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\i\end{bmatrix}$
right-hand circular polarisation	$\frac{1}{\sqrt{2}}\begin{bmatrix}1\\-i\end{bmatrix}$

Table C.1: Jones Vectors to describe polarisation states of optical fields.

Polarisers

Polarisers such as polarising beam-splitters are described by two Jones matrices for the transmission port and the reflection port for a given input direction. For transmission

$$\mathbf{J}_{\mathrm{PBS}_T} = \begin{bmatrix}\tau_p & 0\\ 0 & \tau_s\end{bmatrix} = \begin{bmatrix}\sqrt{T_p} & 0\\ 0 & \sqrt{T_s}\end{bmatrix}, \tag{C.3}$$

and for reflection

$$\mathbf{J}_{\mathrm{PBS}_R} = \begin{bmatrix}\rho_p & 0\\ 0 & \rho_s\end{bmatrix} = \begin{bmatrix}\sqrt{R_p} & 0\\ 0 & \sqrt{R_s}\end{bmatrix}. \tag{C.4}$$

Waveplates

An waveplate is described by the Jones matrix

$$\mathbf{J}_{\mathrm{WP}} = \begin{bmatrix}e^{i\frac{\eta}{2}} & 0\\ 0 & e^{-i\frac{\eta}{2}}\end{bmatrix} \tag{C.5}$$

where η is the birefringence or the relative phase difference between orthogonally polarised light. In the case of quarter waveplate $\eta = \frac{\pi}{2}$ and for a half waveplate $\eta = \pi$.

For an optic that is not aligned to the optical field, rotation matrices ($\mathbf{Rot}(\theta)$) are used. For example for a waveplate with matrix $\mathbf{J'}_{\mathrm{WP}}$, with fast axis orientated

at an angle θ relative to the horizontal polarisation of the optical field,

$$\mathbf{J}_{\mathrm{WP}} = \mathbf{Rot}(-\theta) \cdot \mathbf{J'}_{\mathrm{WP}} \cdot \mathbf{Rot}(\theta) \tag{C.6}$$

$$= \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \sin\theta \end{bmatrix} \cdot \begin{bmatrix} e^{i\frac{\eta}{2}} & 0 \\ 0 & e^{-i\frac{\eta}{2}} \end{bmatrix} \cdot \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}, \tag{C.7}$$

which can be expressed as a single matrix,

$$\mathbf{J}_{WP} = \begin{bmatrix} e^{i\frac{\eta}{2}}\cos^2\theta + e^{-i\frac{\eta}{2}}\sin^2\theta & (e^{i\frac{\eta}{2}} - e^{-i\frac{\eta}{2}})\cos\theta\sin\theta \\ (e^{i\frac{\eta}{2}} - e^{-i\frac{\eta}{2}})\cos\theta\sin\theta & e^{i\frac{\eta}{2}}\cos^2\theta + e^{-i\frac{\eta}{2}}\sin^2\theta \end{bmatrix}, \tag{C.8}$$

Rotation matricies are also useful when considering multiple optical components each with a different alignment.

C.2 Linear Cavity Model

The interaction of light with an optic is described by the multiplication of the Jones vector describing the light before the optic with the Jones matrix that describes the optic. The result of the multiplication represents the polarisation state of the optical field after the optic. To build up a model of an interferometer, multiple matrices are multiplied in series to calculate the optical field at the output of a specified path.

The amplitude transmission and reflection from a linear cavity are described as

$$E_T = \frac{\tau_1\tau_2 E_I e^{-i\delta}}{1 - \rho_1\rho_2 e^{-i2\delta}} \tag{C.9}$$

and

$$E_R = \rho_1 E_I - \frac{\rho_2\tau_1^2 E_I e^{-i2\delta}}{1 - \rho_1\rho_2 e^{-i2\delta}}, \tag{C.10}$$

where δ is the cavity detuning and E_I is the input field amplitude.

To model the linear cavity with birefringent mirrors each mirror is considered to be a waveplate of birefringence η_n $(n = 1, 2)$ with fast axis orientated at an angle θ_n relative to the horizontal axis.

The reflectivity and transmissivities can be replaced with Jones matrices, $\mathbf{R}_n$ and $\mathbf{T}_n$,

$$\mathbf{R}_n = \rho_n \, \mathbf{J}_n \tag{C.11}$$

and

$$\mathbf{T}_n = \tau_n \mathbf{J}_n, \tag{C.12}$$

where $\mathbf{J}_n$ is the Jones matrix of a birefringent element as given above, which can be used in the modified expression of the cavity reflection and transmission field amplitudes

$$\mathbf{J}_{\mathrm{Cav}_T} = \frac{\mathbf{T}_1\mathbf{T}_2\vec{E_I}e^{-i\delta}}{\mathbf{I} - \mathbf{R}_1\mathbf{R}_2\, e^{-i2\delta}} \tag{C.13}$$

and

$$\mathbf{J}_{\mathrm{Cav}_R} = \mathbf{R}_1\vec{E_I} - \frac{\mathbf{R}_2\mathbf{T}_1^2\,\vec{E_I}\, e^{-i2\delta}}{\mathbf{I} - \mathbf{R}_1\mathbf{R}_2\, e^{-i2\delta}}. \tag{C.14}$$

C.3 Speedmeter Model

In the case of a single arm of the polarisation Sagnac speedmeter (or the linear cavity experiment) the path from input to desired output is described by the following combination of Jones matrices

$$E_{out} = \mathbf{J}_{\mathrm{PBS}_R} \cdot \mathbf{J}_{\mathrm{QWP}} \cdot \mathbf{J}_{\mathrm{Cav}_R} \cdot \mathbf{J}_{\mathrm{QWP}} \cdot \mathbf{J}_{\mathrm{PBS}_T} \cdot E_{in}, \tag{C.15}$$

whereas the path of leakage light is described by

$$E_{out} = \mathbf{J}_{\mathrm{PBS}_T} \cdot \mathbf{J}_{\mathrm{QWP}} \cdot \mathbf{J}_{\mathrm{Cav}_R} \cdot \mathbf{J}_{\mathrm{QWP}} \cdot \mathbf{J}_{\mathrm{PBS}_T} \cdot E_{in}. \tag{C.16}$$

Similar combinations of matrices can be constructed to describe any possible axial path though the interferometer. In this way a model is built that can calculate the amplitude of fields coupled to the output of the interferometer along different paths of interest.

C.4 Derivation of Limits to Speedmeter

C.4.1 PBS

Using the Jones Matrix for an ideal quarter waveplate and no cavity birefringence, the limit on the amount of Michelson signal at the output due to finite PBS extinction ratio can be calculated as,

$$\mathbf{Lim}_{\mathrm{PBS}} = \frac{1}{2\sqrt{2}} \begin{bmatrix} \sqrt{1-\frac{1}{\mathrm{PER}_T}} & 0 \\ 0 & \sqrt{\frac{1}{\mathrm{PER}_T}} \end{bmatrix} \cdot \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \cdot \begin{bmatrix} \sqrt{1-\frac{1}{\mathrm{PER}_T}} & 0 \\ 0 & \sqrt{\frac{1}{\mathrm{PER}_T}} \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{C.17}$$

This can be multiplied out to give the amplitude vector due to PBS extinction

$$\mathbf{Lim}_{\mathrm{PBS}} = \frac{i}{\sqrt{2}} \begin{bmatrix} 0 \\ \sqrt{1 - \frac{1}{\mathrm{PER}_T}} \sqrt{\frac{1}{\mathrm{PER}_T}} \end{bmatrix}. \tag{C.18}$$

In a Sagnac speedmeter there are two circulating field so there will be two such leakage fields at the output of the interferometer, the total optical power at the output of Michelson signals due only to the PBS extinction ratio is

$$P_{\mathrm{PBS}} = \frac{\mathrm{PER}_T - 1}{\mathrm{PER}_T^2}. \tag{C.19}$$

C.4.2 QWP

$$\mathbf{Lim}_{\mathrm{QWP}} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} \cos\frac{\eta}{2} & i\sin\frac{\eta}{2} \\ i\sin\frac{\eta}{2} & \cos\frac{\eta}{2} \end{bmatrix} \cdot \begin{bmatrix} \cos\frac{\eta}{2} & i\sin\frac{\eta}{2} \\ i\sin\frac{\eta}{2} & \cos\frac{\eta}{2} \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{C.20}$$

This reduced down to the Jones vector of the output field,

$$\mathbf{Lim}_{\mathrm{QWP}} = \frac{1}{\sqrt{2}} \left[\cos^2\tfrac{\eta}{2} - \sin^2\tfrac{\eta}{2}\right]. \tag{C.21}$$

Using the double angle identity and again considering two field with equivalent paths though the symmetric speedmeter, the power at the output is simply given by

$$P_{\mathrm{QWP}} = \cos\eta. \tag{C.22}$$

Appendix D

Digital Servo

D.1 Simulink Model

The block diagram of the Simulink® model used in the digital laser stabilisation system described in Chapter 12 is shown in Figure D.1.

The model includes five second-order section blocks that are used to implement digital filters to shape the feedback signal that is used to actuate the laser PZT. An additional second-order section filter block is used to further shape the feedback signal that is sent to the PZT that actuates the length of the reference cavity.

The model takes two inputs, the demodulated error signal (Channel 1) that is used for the feedback, and the reflected DC photodiode signal (Channel 2) that is used for the locking trigger. In the triggered state, the locking signal is only delivered if the cavity is on resonance. The resonance condition is determined from the reflected DC signal.

The model also includes a signal generator that can be used to apply a ramp signal to the PZT that scans the length of the reference cavity.

A graphical user interface was developed for the digital frequency stabilisation servo developed for the 1550 nm laser system, using the dSPACE ControlDesk software. A screenshot of the working interface is shown in Figure D.2, this includes filter selection, auto-locker and signal generator control.

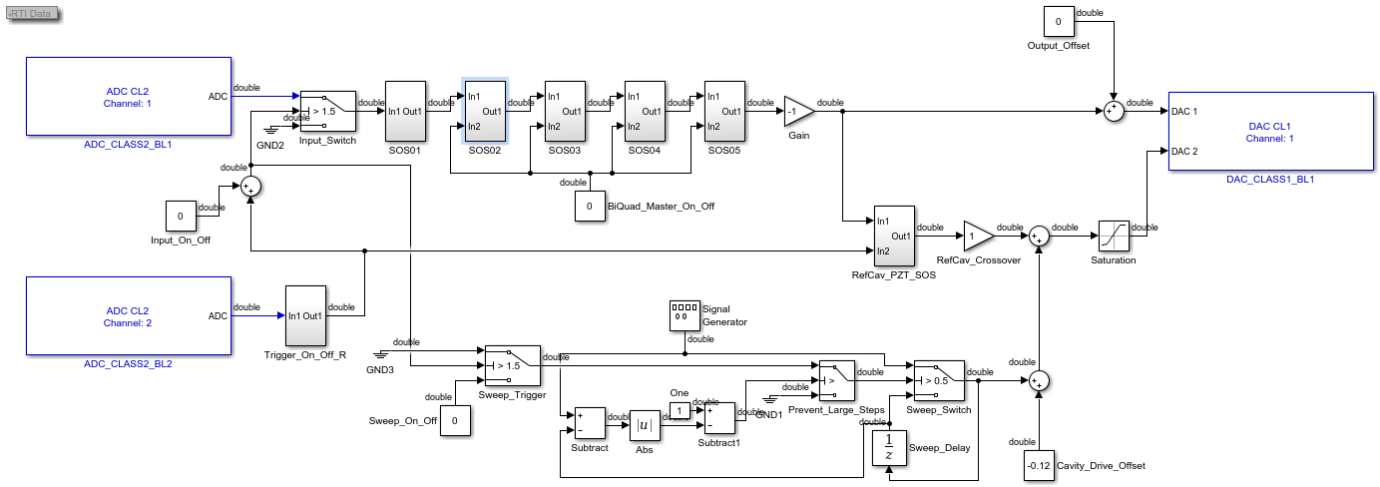


Figure D.1: Block diagram of Simulink® model for frequency stabilisation courtesy of B. Barr including five second-order sections for shaping of the feedback signal. The model also includes a signal generator and triggering for auto-locking.

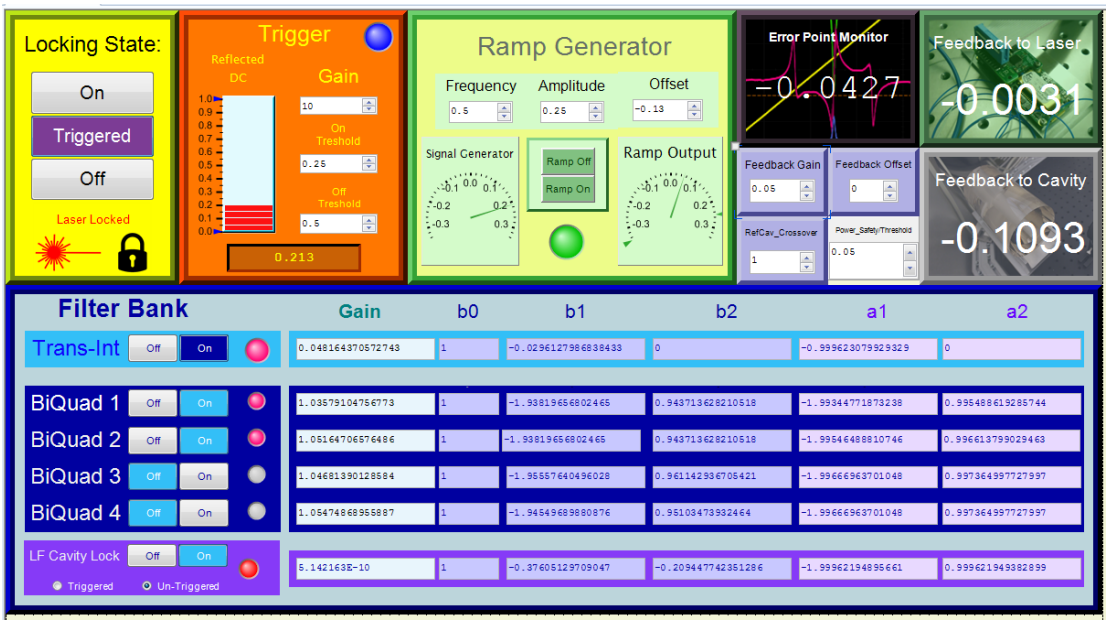


Figure D.2: User interface for the digital laser stabilisation servo. The locking state can be set with the buttons in the yellow box in the top left. When triggered, the feedback is only engaged when the reflected power from the cavity drops below a threshold, indicting the cavity is close to a resonance. The trigger level is displayed in the bar plot in the orange box where the trigger thresholds can also be set. The parameter of the ramp signal driven onto the cavity length actuator can be set in the green box. A readout of the signal being applied to the two actuators (cavity length and laser PZT) is shown in the top right of the interface along with the error signal at the input. The bottom half of the user interface is the filter bank where the parameters of the second-order sections can be changed and individual filters can be turned on or off. A automated script could also be run to load in the filter parameters from a saved file. In this image the laser is locked to the reference cavity with the basic filter and two bi-quadratic filters in the triggered state.

Bibliography

[1] M. Pitkin, S. Reid, S. Rowan, and J. Hough. Gravitational Wave Detection by Interferometry (Ground and Space). *Living Review in Relativity*, 14(5), 2011.

[2] B.P Abbot et al (LIGO Scientific Collaboration and Virgo Collaboration). Observation of Gravitational Waves from a Binary Black Hole Merger. *Phy. Rev. Lett.*, 116:061102, 2016.

[3] D. Martynov et al. (LIGO Scientific Collaboration and Virgo Collaboration). Sensitivity of the Advanced LIGO Detectors at the Beginning of Gravitational Wave Astronomy. *Phys. Rev. D*, 93:112004, 2016.

[4] B.P Abbot et al. (LIGO Scientific Collaboration and Virgo Collaboration). Properties of the Binary Black Hole Merger GW150914. *Phys. Rev. Lett.*, 116:241102, Jun 2016.

[5] B.P Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration). GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9:031040, Sep 2019.

[6] B.P Abbot et al. (LIGO Scientific Collaboration and Virgo Collaboration). GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.

[7] B.P Abbot et al., (LIGO Scientific Collaboration, Virgo Collaboration, Fermi GBM, INTEGRAL, IceCube Collaboration, AstroSat Cadmium ZincTelluride Imager Team, IPN Collaboration, The Insight-HXMT Collaboration, ANTARES Collaboration, The Swift Collaboration, AGILE Team, The 1M2H Team, The Dark Energy Camera GW-EM Collaboration and the DES Collaboration, The DLT40 Collaboration, GRAWITA: GRAvitational Wave Inaf TeAm, The Fermi Large Area Telescope Collaboration,

ATCA: Australia Telescope CompactArray, ASKAP: Australian SKA Pathfinder, Las Cumbres Observatory Group, OzGrav, DWF(Deeper, Wider, Faster Program), AST3 and CAASTRO Collaborations, The VINROUGE Collaboration, MASTER Collaboration, J-GEM, GROWTH, JAGWAR, Caltech-NRAO, TTU-NRAO and NuSTAR Collaborations, Pan-STARRS, TheMAXITeam, TZACConsortium, KU Collaboration, NordicOptical Telescope, ePESSTO, GROND, Texas Tech University, SALT Group, TOROS: Transient Robotic Observatory of the SouthCollaboration, The BOOTES Collaboration, MWA: Murchison Widefield Array, The CALET Collaboration, IKI-GW Follow-upCollaboration, H.E.S.S. Collaboration, LOFAR Collaboration, LWA: Long Wavelength Array, HAWC Collaboration, The Pierre Auger Collaboration, ALMA Collaboration, Euro VLBI Team, Pi of the Sky Collaboration, The Chandra Team at McGill University, DFN:Desert Fireball Network, ATLAS, High Time Resolution Universe Survey, RIMAS and RATIR, and SKA South Africa/MeerKAT). Multi-messenger Observations of a Binary Neutron Star Merger. *The Astrophysical Journal*, 848(2):L12, Oct 2017.

[8] B.P Abbot et al. (LIGO Scientific Collaboration and Virgo Collaboration). Binary Black Hole Mergers in the First Advanced LIGO Observing Run. *Phys. Rev. X*, 6:041015, Oct 2016.

[9] B.P Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration). Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo. *arXiv*, 2018.

[10] J. Aasi et al. (The LIGO Scientific Collaboration). Advanced LIGO. *Class. Quantum Grav.*, 32:074001, 2015.

[11] B. P. Abbott et al. (LIGO Scientific Collaboration). LIGO: the Laser Interferometer Gravitational-wave Observatory. *Rep. Prog. Phys.*, 72:076901, 2009.

[12] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.

[13] A. Cumming, A. Bell, L. Barsotti, M. Barton, G. Cagnoli, D. Cook, L. Cunningam, M. Evans, G. Hammond, G. Harry, A. Heptenstall, J. Hough, R. Jones, R. Kumar, R. Mittleman, N. Robertson, S. Rowan, B. Shapiro, K. Strain, K. Tokmakov, C. Torrie, and A. van Veggel. Design and Development of the Advanced LIGO Monolithic Fused Silica Suspension. *Class. Quantum Grav.*, 29:035003, 2012.

[14] K. Izumi and D. Sigg. Advanced LIGO: Length Sensing and Control in a Dual Recycled Interferometric Gravitational Wave Antenna. *Classical and Quantum Gravity*, 34(1):015001, Dec 2016.

[15] B. Willke, K. Danzmann, C. Fallnich, M. Frede, M. Heurs, P. King, D. Kracht, P. Kwee, SavageR., F. Seifert, and R. Wilhelm. Stabilized High Power Laser for Advanced Gravitational Wave Detectors. *Journal of Physics: Conference Series*, 32:270–275, Mar 2006.

[16] C. S. Unnikrishnan. IndIGO and LIGO-India: Scope and plans for gravitational wave research and precision metrology in India. *Int. J. Mod. Phys.*, D22:1341010, 2013.

[17] S. Fairhurst. Improved Source Localization with LIGO-India. *Journal of Physics: Conference Series*, 484:012007, Mar 2014.

[18] R. Adhikari, N. Smith, A. Brooks, L. Barsotti, B. Shapiro, B. Lantz, D. McClelland, E. Gustafson, D. Martynov, V. Mitrofanov, D. Coyne, K. Arai, and C. Torrie, C. Wipf. LIGO Voyager Upgrade: Design Concept. *LIGO DCC*, T1400226, 2014.

[19] T. Accadia et al. (Virgo Collaboration). Virgo: A Laser Interferometer to Detect Gravitational Waves. *J. Inst.*, 7:P03012, 2012.

[20] F. Acernese et al. (Virgo Collaboration). Advanced Virgo: A Second-Generation Interferometric Gravitational Wave Detector. *Class.Quantum Grav.*, 32(2):024001, 2015.

[21] B.P Abbot et al. (LIGO Scientific Collaboration and Virgo Collaboration). GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 119:141101, Oct 2017.

[22] B. Willke et al. The GEO 600 Gravitational Wave Detector. *Class. Quantum Grav.*, 19:1377–1387, 2002.

[23] H. Grote, A. Freise, M. Malec, G. Heinzel, B. Willke, H. Lück, K.A. Strain, J. Hough, and K. Danzmann. Dual recycling for GEO 600. *Classical and Quantum Gravity*, 21(5):S473–S480, Feb 2004.

[24] J. Abadie et al. A Gravitational Wave Observatory Operating Beyond the Quantum Shot-Noise Limit: Squeezed light in application. *Nature Phys.*, 7:962–965, 2011.

[25] H. Grote, K. Danzmann, K. Dooley, R. Schnabel, J. Slutsky, and H. Vahlbruch. First Long-Term Application of Squeezed States of Light in a Gravitational-Wave Observatory. *Phys. Rev. Lett.*, 110:181101, 2013.

[26] T. Uchiyama, K. Furuta, M. Ohashi, S. Miyoki, O. Miyakawa, and Y. Saito. Excavation of an Underground Site for a km-scale Laser Interferometric Gravitational-Wave Detector. *Classical and Quantum Gravity*, 31(22):224005, Nov 2014.

[27] T. Akutsu et.al (KAGRA Scientific Congress). Construction of KAGRA: an underground gravitational-wave observatory. *Progress of Theoretical and Experimental Physics*, 2018(1), Jan 2018.

[28] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, H. Yamamoto, and (The KAGRA Collaboration). Interferometer Design of the KAGRA Gravitational Wave Detector. *Phys. Rev. D*, 88:043007, 2013.

[29] M. Punturo et al. The Einstein Telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, Sep 2010.

[30] S. Hild et al. Sensitivity studies for third-generation gravitational wave observatories. *Class. Quantum Grav.*, 28:094013, 2011.

[31] S. Hild, S. Chelkowski, A. Freise, J. Franc, N. Morgado, R. Flaminio, and R. DeSalvo. A Xylophone Configuration for a Third Generation Gravitational Wave Detector. *Class. Quantum Grav.*, 27:015003, 2010.

[32] F. Matichard, B. Lantz, R. Mittleman, K. Mason, J. Kissel, B. Abbott, S. Biscans, J. McIver, et al. Seismic Isolation of Advanced LIGO: Review of strategy, instrumentation and performance. *Class. Quant. Grav.*, 32(18), 2015.

[33] C. Torrie. *Development of Suspensions for the GEO 600 Gravitational Wave Detector*. PhD Thesis, University of Glasgow, 2001.

[34] GWINC-The Gravitational Wave Interferometer Noise Calculator. `https://awiki.ligo-wa.caltech.edu/wiki/GWINC`.

[35] Y. Levin. Internal Thermal Noise in the LIGO Test Masses: A direct approach. *Phys. Rev. D*, 57(659), 1998.

[36] M. Husman, J. Hough, and N. Robertson. Thermal Noise in a Pendulum Suspended by Multiple Wires. *Classical and Quantum Gravity*, 21:1371, 2004.

[37] G.M. Harry et al. Thermal Noise in Interferometric Gravitational Wave Detectors Due to Dielectric Optical Coatings. *Class. Quant. Grav.*, 19:897–918, 2002.

[38] G. Gonzalez. Suspensions Thermal Noise in the LIGO Gravitational Wave Detector. *Classical and Quantum Gravity*, 17:4409–4435, 2000.

[39] A. Gillespie and F. Raab. Thermal Noise in the Test Mass Suspensions of a Laser Interferometer Gravitational-Wave Detector Prototype. *Physics Letters A*, 178:357–363, 1993.

[40] G. Hammond, A. Cumming, J. Hough, R. Kumar, K. Tokmakow, S. Reid, and S. Rowan. Reducing the Suspension Thermal Noise in Advanced Gravitational Wave Detectors. *Classical and Quantum Gravity*, 29, 2012.

[41] G. Harry, T. Bodiya, and R. DeSalvo. Optical Coatings and Thermal Noise in Precision Measurement. *Cambridge University Press*, 2012.

[42] M. Granata, E. Saracco, N. Morgado, A. Cajgfinger, G. Cagnoli, J. Degallaix, V. Dolique, D. Forest, J. Franc, C. Michel, L. Pinard, and R. Flaminio. Mechanical Loss in State-of-the-art Amorphous Optical Coatings. *Phys. Rev. D*, 93:012007, 2016.

[43] R. Flaminio, J. Franc, C. Michel, N. Morgado, L. Pinard, and B. Sassolas. A Study of Coating Mechanical and Optical Losses in View of Reducing Mirror Thermal Noise in Gravitational Wave Detectors. *Class. Quantum Grav.*, 27:084030, 2010.

[44] S.A. Hughes and K.S. Thorne. Seismic Gravity-Gradient Noise in Interferometric Gravitational-Wave Detectors. *Phys. Rev. D*, 58:122002, Nov 1998.

[45] M.G. Beker, J.F.J. van den Brand, E. Hennes, and D.S. Rabeling. Newtonian Noise and Ambient Ground Motion for Gravitational Wave Detectors. *Journal of Physics: Conference Series*, 363:012004, Jun 2012.

[46] A. Freise and K. Strain. Interferometer Techniques for Gravitational-Wave Detection. *Living Review in Relativity*, 13(1), 2010.

[47] E. Collett. *Field Guid to Polarization*. SPIE Press, 2005.

[48] W. Shurcliff. *Polarized Light Production and Use*. Harvard University Press, 1962.

[49] R. Paschotta. *Encyclopedia of Laser Physics and Technology*. Wiley-VCH, 12th edition, 2008.

[50] G. Hammond, S. Hild, and M. Pitkin. Advanced Technologies for Future Groung-Based, Laser-Interferometric Gravitational Wave Detectors. *Journal of Modern Optics*, 61(S10-S45), 2014.

[51] M. Bassan, editor. *Advanced Interferometers and the Search for Gravitational Waves*, volume 404. Springer, 2014.

[52] G. Heinzel, K. Strain, J. Mizuno, K. Skeldon, B. Willke, W. Winkle, R. Schilling, and K. Danzmann. An Experimental Demonstration of Dual Recycling on a Suspended Interferometer. *Phys. Rev. Lett.*, 81:5493–5496, 1998.

[53] K. Strain and B. Meers. Experimental Demonstration of Dual Recycling for Interferometric Gravitational-Wave Detectors. *Phys. Rev. Lett.*, 66(11):1391–1394, 1991.

[54] FINESSE - Frequency domain Interferometer Simulation Software. `https://www.gwoptics.org/finesse/`.

[55] C. Caves. Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer. *Phys. Rev. Lett.*, 45:75–79, 1980.

[56] C. Caves. Quantum-Mechanical Noise in an Interferometer. *Phys. Rev. D*, 23(8):1693–1708, 1981.

[57] S. Danilishin and F. Khalili. Quantum Measurement Theory in Gravitational Wave Detectors. *Living Reviews in Relativity*, 15(5), 2012.

[58] S. Hild. A Basic Introduction to Quantum Noise and Quantum-Non-Demolition Techniques In: Bassan M. (eds) Advanced Interferometers and the Search for Gravitational Waves. *Astrophysics and Space Science Library. Springer, Cham*, 404(11), 2014.

[59] V. Braginsky, F. Khalili, and K. Thorne. *Quantum Measurement*. Cambridge University Press, 1995.

[60] K.A Strain, K. Danzmann, J. Mizuno, P.G. Nelson, A. Rüdiger, R. Schilling, and W. Winkler. Thermal Lensing in Recycling Interferometric Gravitational Wave Detectors. *Physics Letters A*, 194(1):124 – 132, 1994.

[61] W. Ling-Aan, X. Min, and H. Kimble. Squeezed States of Light from an Optical Parametric Oscillator. *Journal of the Optical Society of America*, 4(10), 1987.

[62] H. Vahlbruch. *Squeezed Light for Gravitational Wave Astronomy.* PhD Thesis, Gottfried Wilhelm Leibniz Universitat Hannover, 2008.

[63] E. Oelker, L. Barsotti, S. Dwyer, D. Sigg, and N. Mavalvala. Squeezed Light for Advanced Gravitational Wave Detectors and Beyond. *Opt. Express*, 22(17):21106–21121, Aug 2014.

[64] J.P Zendri et al. AdVirgo: Frequency Independent Squeezing Status Report. *LIGO DCC*, G1900600, 2019.

[65] K. Thorne, R. Drever, C. Caves, M. Zimmermann, and V. Sandberg. Quantum Nondemolition Measurements of Harmonic Oscillators. *Phy. Rev. Lett.*, 40:667, 1978.

[66] F. Khalili. Quantum Speedmeter and Laser Interferometric Gravitational-Wave Antennae. *arXiv:gr-qc/0211088*, 2002.

[67] V. Braginsky and F. Khalili. Gravitational Wave Antenna with QND Speed Meter. *Phy. Lett. A*, 147:251, 1990.

[68] V. Braginsky, M. Gorodetsky, F. Khalili, and K. Thorne. Dual-Resonator Speed Meter for a Free Test Mass. *Phys. Rev. D*, 61:4002, 2000.

[69] P. Purdue. Analysis of a Quantum Nondemolition Speed-Meter Interferomter. *Phys. Rev. D*, 66:022001, 2002.

[70] S. Danilishin. Quantum Noise in the Next Generation of GW Detectors and how to suppress it. *Invited talk at the GRASS 2018 Workshop in Padova*, LIGO DCC(G1800304), 2018.

[71] Y. Chen. Sagnac Interferometer as a Speed-Meter-Type, Quantum-Nondemolition Gravitational-Wave Detector. *Phy. Rev. D*, 67:122004, 2003.

[72] D. Pascucci. *On Optics Surface Imperfections and Their Effects on the Sensitivity of Speed Meters.* PhD Thesis, University of Glasgow, 2019.

[73] S. Danilishin. Sensitivity Limitations in Optical Speed Meter Topology of Gravitational-Wave Antennas. *Phy. Rev. D*, 69:102003, 2004.

[74] M. Wang, C. Bond, D. Brown, F. Bruckner, L. Carbone, R. Palmer, and A. Freise. Realistic Polarizing Sagnac Topology with DC Readout for the Einstein Telescope. *Phy. Rev. D*, 87:096008, 2013.

[75] P. Purdue and Y. Chen. Practical Speed Meter Designs for Quantum Nondemolition Gravitational-Wave Interferometers. *Phy. Rev. D*, 66:122004, 2002.

[76] A. Wade, K. McKenzie, Y. Chen, D. Shaddock, J. Chow, and D. McClelland. Polarization Speed Meter for Gravitational-Wave Detection. *Phys. Rev. D*, 86:062001, 2012.

[77] S. Danilishin, E. Knyazev, N. Voronchev, F. Khalili, C. Gräf, S. Steinlechner, J. Hennig, and S. Hild. A New Type of Quantum Speed Meter Interferometer Measuring Speed to Search for Intermediate Mass Black Holes. *Light Science and Applications*, 7(11), 2018.

[78] S. Huttner, S. Danilishin, B. Barr, A. Bell, C. Gräf, J. Hennig, S. Hild, E. Houston, S. Leavey, D. Pascucci, B. Sorazu, A. Spencer, S. Steinlechner, J. Wright, T. Zhang, and K. Strain. Candidates for a Possible Third-Generation Gravitational Wave Detector: Comparison of ring-Sagnac and sloshing-Sagnac speedmeter interferometers. *Class. Quantum Grav.*, 34(2), 2016.

[79] F.Y. Khalili. Reducing the Mirrors Coating Noise in Laser Gravitational-Wave Antennae by Means of Double Mirrors. *Physics Letters A*, 334(1):67 – 72, 2005.

[80] A.G. Gurkovsky, D. Heinert, S. Hild, R. Nawrodt, K. Somiya, S. Vyatchanin, and H. Wittel. Reducing Thermal Noise in Future Gravitational Wave Detectors by Employing Khalili Etalons. *Physics Letters A*, 375(46):4147–4157, Nov 2011.

[81] C. Gräf, B. Barr, A. Bell, F. Campbell, A. Cumming, S. Danilishin, N. Gordon, G. Hammond, J. Hennig, E. Houston, S. Huttner, R. Jones, S. Leavey, H. Luck, J. Macarthur, M. Marwick, S. Rigby, R. Schilling, B. Sorazu, A. Spencer, S. Steinlechner, K. Strain, and S. Hild. Design of a Speed Meter Interferometer Proof of Principle Experiment. *Classical and Quantum Gravity*, 31(21), 2014.

[82] S. Danilishin, C. Gräf, S. Leavey, J. Hennig, E. Houston, D. Pascucci, S. Steinlechner, J. Wright, and S. Hild. Quantum Noise of Non-Ideal Sagnac Speed Meter Interferometer with Asymmetries. *New Journal of Physics*, 17, 2015.

[83] T. Zhang, S. Danilishin, S. Steinlechner, B. Barr, A. Bell, P. Dupej, C. Gräf, J. Hennig, E. Houston, S. Huttner, S. Leavey, D. Pascucci, B. Sorazu, A. Spencer, J. Wright, K. Strain, and S. Hild. Effects of Static and Dynamic Higher-Order Optical Modes in Balanced Homodyne Readout for Future Gravitational Waves Detectors. *Phys. Rev. D*, 95:062001, 2017.

[84] H.J. Eichler, J. Eichler, and O. Lux. *Lasers: Basics, Advances and Applications*. Springer, 2018.

[85] Layertech. `https://www.layertec.de/en/`.

[86] Union Optic. `https://www.u-optic.com/`.

[87] F. Brandi, F. Della Valle, A.M. De Riva, P. Micossi, F. Perrone, C. Rizzo, G. Ruoso, and G. Zavattini. Measurement of the Phase Anisotropy of Very High Reflectivity Interferential Mirrors. *Applied Physics B*, 65:351–355, Sept 1997.

[88] F. Della Valle, A. Ejlli, U. Gastaldi, G. Messineo, E. Milotti, R. Pengo, G. Ruoso, and G. Zavattini. The PVLAS Experiment: Measuring vacuum magnetic birefringence and dichroism with a birefringent Fabry-Perot cavity. *The European Physical Journal C*, 76(1):24, Jan 2016.

[89] R. Battesti, B. Pinto Da Souza, S. Batut, et al. The BMV Experiment: A novel apparatus to study the propagation of light in a transverse magnetic field. *Eur. Phys. J. D*, 46(2):323–333, 2008.

[90] F. Della Valle, E. Milotti, A. Ejlli, G. Messineo, L. Piemontese, G. Zavattini, U. Gastaldi, R. Pengo, and G. Ruoso. First Results From the New PVLAS Apparatus: A new limit on vacuum magnetic birefringence. *Phys. Rev. D*, 90(9):092003, 2014.

[91] F. Bielsa, A. Dupays, M. Fouche, R. Battesti, C. Robilliard, and C. Rizzo. Birefringence of Interferential Mirrors at Normal Incidence (Experimental and Computational study). *Appl. Phys. B*, 97(2):457–463, 2009.

[92] N. Wei-Tou. Cryogenic Test and Interferonletric Test of QED and Search for Axions and Anomalous Spin-Dependent Interactions. *Chinese Journal of Physics*, 34(3S):962–967, 1996.

[93] S. Moriwaki, H. Sakaida, T. Yuzawa, and N. Mio. Measurement of the Residual Birefringence of Interferential Mirrors Using Fabry-Perot Cavity. *Appl. Phys. B*, 65:347–350, 1997.

[94] C. Wood, S. C. Bennett, J. L. Roberts, D. Cho, and C. E. Wieman. Birefringence, Mirrors, and Parity Violation. *Optics & Photonics News*, 7:54–55, Oct 1996.

[95] J.L. Hall, J. Ye, and M. Long-Sheng. Measurement of Mirror Birefringence at the sub-ppm Level: Proposed application to a test of QED. *Phys. Rev. A*, 62:013815, Jun 2000.

[96] D. Jacob, M. Vallet, F. Bretenaker, A. Le Floch, and M. Oger. Supermirror Phase Anisotropy Measurement. *Opt. Lett.*, 20(7):671–673, Apr 1995.

[97] K.D. Skeldon, D.A. Clubley, G.P. Newton, S. Thieux, M. Von Gradowski, and B.W. Barr. Measurements of an Ultra-Low Loss Polarizer for λ =1064 nm Using a High Finesse Optical Cavity. *Journal of Modern Optics*, 48:695–702, Mar 2001.

[98] N. Gordon. *Characterisation and Control of Coupled Optical Springs for Future Gravitational Wave Detectors*. PhD Thesis, University of Glasgow, 2015.

[99] R. Bork, R. Abbott, D. Barker, and J. Heefner. An Overview of the LIGO Control and Data Acquisition Systems. *Proc. ICALEPCS*, pages 19–23, Jan 2001.

[100] R. Drever, J. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward. Laser Phase and Frequency Stabilization Using an Optical Resonator. *Appl. Phys. B*, 31:97–105, 1983.

[101] A.E. Siegman. *Lasers*. University Science Books, 1986.

[102] O. Svelto. *Principles of Lasers*. Springer, 1998.

[103] B. Willke, K. Danzmann, M. Frede, P. King, D. Kracht, P. Kwee, O. Puncken, R.L Savage, B. Schulz, F. Seifert, C. Veltkamp, S. Wagner, P. Weßels,

and L. Winkelmann. Stabilized Lasers for Advanced Gravitational Wave Detectors. *Classical and Quantum Gravity*, 25(11):114040, May 2008.

[104] W. Li-Wei. *High-Power Laser System for Advanced Virgo Gravitational Wave Detector: Coherently combined master oscillator fiber power amplifiers*. PhD Thesis, Université Nice Sophia Antipolis, Dec 2015.

[105] I. Zawischa, M. Brendel, K. Danzmann, C. Fallnich, M. Heurs, S. Nagano, V. Quetschke, H. Welling, and B. Willke. The GEO 600 Laser System. *Classical and Quantum Gravity*, 19:1775–1781, March 2002.

[106] T.J. Kane and R.L. Byer. Monolithic, Unidirectional Single-Mode Nd:YAG Ring Laser. *Opt. Lett.*, 10(2):65–67, Feb 1985.

[107] P. Kwee and B. Willke. Automatic Laser Beam Characterization of Monolithic Nd:YAG Nonplanar Ring Lasers. *Appl. Opt.*, 47(32):6022–6032, Nov 2008.

[108] M. Fukuda, T. Hirono, T. Kurosaki, and F. Kano. 1/f noise behavior in semiconductor laser degradation. *IEEE Photonics Technology Letters*, 5(10):1165–1167, Oct 1993.

[109] C. Henry. Phase noise in semiconductor lasers. *Journal of Lightwave Technology*, 4(3):298–311, 1986.

[110] K. Kikuchi. Effect of 1/f-type fm noise on semiconductor-laser linewidth residual in high-power limit. *IEEE Journal of Quantum Electronics*, 25(4):684–688, 1989.

[111] Kenji Numata, Jordan Camp, Michael A. Krainak, and Lew Stolpner. Performance of planar-waveguide external cavity laser for precision measurements. *Opt. Express*, 18(22):22781–22788, Oct 2010.

[112] D. J. Richardson, J. Nilsson, and W. A. Clarkson. High Power Fiber Lasers: Current status and future perspectives. *J. Opt. Soc. Am. B*, 27(11):B63–B92, Nov 2010.

[113] A. Malouf, O. Henderson-Sapir, M. Gorjan, and D. Ottaway. Numerical Modeling of 3.5 micron Dual-Wavelength Pumped Erbium Doped Mid-Infrared Fiber Lasers. *IEEE Journal of Quantum Electronics*, PP, Jun 2016.

[114] L. Winkelmann, O. Puncken, R. Kluzik, C. Veltkamp, P. Kwee, J. Poeld, C. Bogan, B. Willke, M. Frede, J. Neumann, P. Weßels, and D. Kracht.

Injection-Locked Single-Frequency Laser with an Output Power of 220 W. *Applied Physics B*, 102(3):529–538, Mar 2011.

[115] M. Frede, R. Wilhelm, R. Gau, M. Brendel, I. Zawischa, C. Fallnich, F. Seifert, and B. Willke. High-Power Single-Frequency Nd:YAG Laser for Gravitational Wave Detection. *Classical and Quantum Gravity*, 21(5):S895–S901, Feb 2004.

[116] M. Steinke, H. Tunnermann, V. Kuhn, T. Theeg, M. Karow, O. de Varona, P. Jahn, P. Booker, J. Neumann, P. Weßels, and D. Kracht. Single-Frequency Fiber Amplifiers for Next-Generation Gravitational Wave Detectors. *IEEE Journal of Selected Topics in Quantum Electronics*, PP:1–1, Oct 2017.

[117] N. Mavalvala, D.E. McClelland, G. Mueller, D. H. Reitze, R. Schnabel, and B. Willke. Lasers and Optics: Looking towards third generation gravitational wave detectors. *General Relativity and Gravitation*, 43(2):569–592, Feb 2011.

[118] V. Quetschke et al. Electro Optic Modulators and Modulation for Enhanced LIGO and beyond. In *Coherent Optical Technologies and Applications*, page CMC1. Optical Society of America, 2008.

[119] V.P Mitrofanov and K.V Tokmakov. Effect of Heating on Dissipation of Mechanical Energy in Fused Silica Fibers. *Physics Letters A*, 308:212–219, 2003.

[120] J. Degallaix, J. Komma, D. Forest, G. Hofmann, M. Granata, D. Heinert, C. Schwarz, R. Nawrodt, L. Pinard, C. Michel, R. Flaminio, and G. Cagnoli. Measurement of the Optical Absorption of Bulk Silicon at Cryogenic Temperature and the Implication for the Einstein Telescope. *Classical and Quantum Gravity*, 31(18):185010, Aug 2014.

[121] E. Hirose, D. Bajuk, G. Billingsley, T. Kajita, B. Kestner, N. Mio, M. Ohashi, B. Reichman, H. Yamamoto, and L. Zhang. Sapphire Mirror for the KAGRA Gravitational Wave Detector. *Physical Review D*, 89, Feb 2014.

[122] R. Kumar, D. Chen, A. Hagiwara, T. Kajita, T. Miyamoto, T. Suzuki, Y. Sakakibara, H. Tanaka, K Yamamoto, and T. Tomaru. Status of the Cryogenic Payload System for the KAGRA Detector. *Journal of Physics: Conference Series*, 716:012017, May 2016.

[123] P. Kwee. *Laser Characterization and Stabilization for Precision Interferometry*. PhD Thesis, Gottfried Wilhelm Leibniz Universitat Hannover, 2010.

[124] J. B. Camp, H. Yamamoto, S. E. Whitcomb, and D. E. McClelland. Analysis of Light Noise Sources in a Recycled Michelson Interferometer with Fabry-Perot Arms. *J. Opt. Soc. Am. A*, 17(1):120–128, Jan 2000.

[125] K. Somiya, Y. Chen, S. Kawamura, and N. Mio. Frequency Noise and Intensity Noise of Next-Generation Gravitational-Wave Detectors with RF/DC Readout Schemes. *Phys. Rev. D*, 73:122005, Jun 2006.

[126] B.W Barr, K.A. Strain, and C.J Killow. Laser Amplitude Stabilization for Advanced Interferometric Gravitational Wave Detectors. *Classical and Quantum Gravity*, 22:4279, Oct 2005.

[127] E. Black. An Introduction to Pound-Drever-Hall Laser Frequency Stabilization. *Am J Phys.*, 69(1):79–87, 2000.

[128] M. Frede, B. Schulz, R. Wilhelm, P. Kwee, F. Seifert, B. Willke, and D. Kracht. Fundamental Mode, Single-Frequency Laser Amplifier for Gravitational Wave Detectors. *Opt. Express*, 15(2):459–465, Jan 2007.

[129] M. Frede, R. Wilhelm, D. Kracht, and C. Fallnich. Nd:YAG Ring Laser with 213 W Linearly Polarized Fundamental Mode Output Power. *Opt. Express*, 13(19):7516–7519, Sep 2005.

[130] B. Willke, P. King, R. Savage, and P. Fritschel. Pre-Stabilized Laser Design Requirements. *LIGO DCC*, T050036, 2005.

[131] P. Kwee, C. Bogan, K. Danzmann, M. Frede, H. Kim, P. King, J. Pold, O. Puncken, R. Savage, F. Seifert, P. Weßels, L. Winkelmann, and B. Willke. Stabilized High-Power Laser System for the Gravitational Wave Detector Advanced LIGO. *Opt. Express*, 20(10):10617–10634, 2012.

[132] H. Kim, P. King, C. Kramer, P. Kwee, J. Pold, R. Savage, P. Weßels, and B. Willke. PSL Final Design. *LIGO DCC*, T0900649, 2009.

[133] O. Puncken. Lectures on the Advanced LIGO PSL. *LIGO DCC*, G1301316, 2013.

[134] B. Willke, N. Uehara, E. K. Gustafson, R. L. Byer, P. J. King, S. U. Seel, and R. L. Savage. Spatial and Temporal Filtering of a 10-W Nd:YAG Laser

with a Fabry-Perot Ring-Cavity Premode Cleaner. *Opt. Lett.*, 23(21):1704–1706, Nov 1998.

[135] J. C. Driggers et al. Improving Astrophysical Parameter Estimation via Offline Noise Subtraction for Advanced LIGO. *Phys. Rev. D*, 99(4):042001, 2019.

[136] neoLase. `http://https://neolase.com/en/`.

[137] F. Thies, N. Bode, P. Oppermann, M. Frede, B. Schulz, and B. Willke. Nd:YVO4 High-Power Master Oscillator Power Amplifier Laser System for Second-Generation Gravitational Wave Detectors. *Opt. Lett.*, 44(3):719–722, Feb 2019.

[138] J. Gleason, R. Martin, and T. Hardwick. As Built PSL/IO Table Layout for L1. *LIGO DCC*, D1300347-v2, 2019.

[139] D. Sigg, P. Schwinberg, and J. Myers. TTFSS: Circuit diagram. *LIGO DCC*, D040105-x0, 2004.

[140] LISO: Program for Linear Simulation and Optimization of analog electronic circuits. `http://www2.mpq.mpg.de/~ros/geo600_docu/soft/liso/manual.pdf`.

[141] G. Heinzel. *Advanced Optical Techniques for Laser-Interferometric Gravitational-Wave Detectors*. PhD Thesis, Universitat Hannover, 1999.

[142] Apex PA85 High Voltage Power Operational Amplifiers. `https://www.apexanalog.com/resources/products/pa85u.pdf`.

[143] Analog Devices AD829 High Speed Low Noise Video Operational Amplifier. `https://www.analog.com/media/en/technical-documentation/data-sheets/AD829.pdf`.

[144] D. Robertson, E. Morrison, J. Hough, S. Killbourn, B. Meers, G. Newton, N. Robertson, K. Strain, and H. Ward. The Glasgow 10m prototype laser interferometric gravitational wave detector. *Review of Sceintific Instruments*, 66:4447, 1995.

[145] N.A Gordon, B.W Barr, A. Bell, C. Gräf, S. Hild, S.H Huttner, S.S. Leavey, J. Macarthur, B. Sorazu, J. Wright, and K.A. Strain. Experimental Demonstration of Coupled Optical Springs. *Classical and Quantum Gravity*, 34(3):035020, Jan 2017.

[146] J. MacArthur. *Towards Surpassing the Standard Quantum Limit Using Optical Springs*. PhD Thesis, University of Glasgow, 2014.

[147] Pure Photonics. `https://www.pure-photonics.com/`.

[148] NuFern. `https://www.nufern.com/`.

[149] Thorlabs - P3-1550PM-FC -1550 nm PANDA PM FC/APC Patch Cable. `https://www.thorlabs.com/thorproduct.cfm?partnumber=P3-1550PM-FC-1`.

[150] Extremely Low Thermal Expansion Glass Ceramic ZERODUR. `https://www.schott.com/advanced_optics/english/products/optical-materials/zerodur-extremely-low-expansion-glass-ceramic/zerodur/index.html`.

[151] Mini-circuits HELA-10 RF Amplifier. `https://ww2.minicircuits.com/pdfs/HELA-10+.pdf`.

[152] Analog Devices OP27 Low Noise, Precision Operational Amplifier. `https://www.analog.com/media/en/technical-documentation/data-sheets/OP27.pdf`.

[153] Analog Devices OP37 Low Noise, Precision, High Speed Operational Amplifier. `https://www.analog.com/media/en/technical-documentation/data-sheets/OP37.pdf`.

[154] Edward M. Kerwin. Damping of Flexural Waves by a Constrained Viscoelastic Layer. *The Journal of the Acoustical Society of America*, 31(7):952–962, 1959.

[155] Texas Instruments TL071 Low Noise JFET-Input Operational Amplifier. `http://www.ti.com/lit/ds/symlink/tl071.pdf`.

[156] dSPACE. `https://www.dspace.com/en/ltd/home.cfm`.

[157] J. Bechhoefer. Feedback for Physicists: A tutorial essay on control. *Review of Modern Physics*, 77, 2005.

[158] G.F. Franklin, J.D. Powell, and M. Workman. *Digital Control of Dynamic Systems*. Prentice Hall, 1997.

[159] M. Evans. Digital Filter Noise. *LIGO DCC*, G0900928, 2009.

[160] J.O. Smith. *Introduction to Digital Filters*. Julius Smith, 2007.

[161] Simulink: Simulation and Model-Based Design. `https://uk.mathworks.com/products/simulink.html`.

[162] Thorlabs - DET20C2 - InGaAs Detector. `https://www.thorlabs.com/thorproduct.cfm?partnumber=DET20C2#ad-image-0`.

[163] M. Herpy and J.C. Berka. *Active RC Filter Design*. Elsevier, 1986.