



Leavey, Sean Stephan (2017) *Enhancing the sensitivity of future laser-interferometric gravitational wave detectors*. PhD thesis.

<http://theses.gla.ac.uk/7902/>

Available under License Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0):

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

Glasgow Theses Service

<http://theses.gla.ac.uk/>

theses@gla.ac.uk

ENHANCING THE SENSITIVITY OF FUTURE
LASER-INTERFEROMETRIC GRAVITATIONAL
WAVE DETECTORS

SEAN STEPHAN LEAVEY

MSci

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

SCHOOL OF PHYSICS AND ASTRONOMY
COLLEGE OF SCIENCE AND ENGINEERING
UNIVERSITY OF GLASGOW

SEPTEMBER 2016

The first direct detection of gravitational waves last year was the beginning of a new field of astronomy. While we have already learned a great deal from the signals sensed by the LIGO interferometers in their first observation run, research is already underway to improve upon the sensitivity of the state of the art detectors. Novel mirror designs, new interferometer topologies and larger, more advanced detectors are all being considered as future improvements, and these topics form the focus of this thesis.

A reduction in the thermal noise arising from the mirrors within gravitational wave detectors will enhance sensitivity near their most sensitive frequencies, and this can potentially be achieved through the use of *waveguide* mirrors employing gratings. It has been shown that the thermal noise is reduced in waveguide mirrors compared to standard dielectric mirrors whilst retaining the required reflectivity, but an open question regarding their suitability remains due to the potential for increased technical noise coupling created by the substructure. We place an upper limit on this coupling with a suspended cavity experiment, showing that this approach to the design of grating mirrors has promise.

While the use of higher classical laser input initially increases interferometer sensitivity, eventually the Michelson interferometer topology employed in existing detectors reaches the *standard quantum limit* preventing further enhancement. Efforts are being made to test the suitability of so-called *quantum non-demolition* (QND) technologies able to surpass this limit, one of which involves the use of a new interferometer topology altogether. An experiment to demonstrate a reduction in quantum radiation pressure noise in a QND-compatible Sagnac speed meter topology is underway in Glasgow, and we introduce novel techniques to control this suspended, audio-band interferometer to inform the technical design of future detectors wishing to measure beyond the standard quantum limit. In particular, the problem of controlling the interferometer at low frequencies is discussed. Due to the nature of the speed meter topology, the response of the interferometer vanishes towards zero frequency, while the interferometer's noise does not. This creates a control problem at low frequencies where test mass perturbations arising from, for example, seismic and electronic noise, can lead to loss of interferometer sensitivity over the course of minutes to hours. We present a solution involving the blending of signals from different readout ports of the interferometer, facilitating measurements with almost arbitrary integration times.

The longer, larger Einstein Telescope facility planned as part of the next generation of detectors will push the Michelson interferometer topology to the limit. The low frequency interferometer will utilise optomechanical interactions to enhance its sensitivity at low frequencies, and the control problems associated with this technique have not been investigated in detail. Following the approach taken in the current generation of detectors we show that the interferometer can be controlled without adversely affecting its sensitivity to gravitational waves, paving the way for a future technical design.

Revision history

Version	Date	Description
1.0.0	2016-09-14	Submitted for examination
1.1.0	2017-01-27	Publication date

An electronic copy of the latest version of this work can be found at
<https://attackllama.com/thesis/>.

Content from this work may be used under the terms of the Creative Commons *Attribution-NonCommercial-ShareAlike 4.0 International* licence.



Table of contents

Abstract	iii
Revision history	iv
Licence	v
Table of contents	vii
List of tables	xiii
List of figures	xv
Acknowledgements	xxi
Dedication	xxv
1 Gravitational waves	1
1.1 Event GW150914	1
1.2 Scientific outcomes from the first joint observation run	2
1.3 General relativity and gravitational waves	3
1.4 Development of the gravitational wave detector	5
1.4.1 Resonant bars	5
1.4.2 The gravitational wave interferometer	6
1.5 Current and future interferometric detectors	7
1.6 Thesis structure	8
2 Sensitivity and noise in gravitational wave interferometers	11
2.1 Interferometer foundations	11
2.1.1 Measurement of interferometer length fluctuations	12
2.1.2 Optimal operating point	13
2.1.3 Readout	14

2.1.4	Control	16
2.2	Measurement noise in interferometers	16
2.2.1	Noise arising from loss and uncertainty	17
2.2.2	Thermal noise	19
2.2.3	Quantum noise and the Standard Quantum Limit	22
2.2.4	Other fundamental noise	25
2.2.5	Technical noise	26
2.3	Sensitivity of the Michelson interferometer	28
2.3.1	Fabry-Perot arm cavities	29
2.3.2	Power recycling	30
2.3.3	Signal recycling	30
2.3.4	Dual recycling with Fabry-Perot arm cavities	31
2.4	Surpassing the Standard Quantum Limit	31
2.4.1	Squeezing	32
2.4.2	Variational readout	32
2.4.3	Light-mirror interactions	33
2.4.4	Modification of the interferometer design	33
2.5	The future of ground-based gravitational wave interferometry	34
3	Measurement of transverse-to-longitudinal phase coupling in a waveguide mirror	35
3.1	Thermal noise in advanced detectors	35
3.2	Waveguide and grating mirrors	36
3.2.1	Transverse to longitudinal coupling in grating mirrors	37
3.3	Experiment	39
3.3.1	Effect of waveguide mirror rotation on cavity length	40
3.3.2	Cavity length signals	41
3.3.3	Experiment infrastructure	42
3.3.4	Measuring cavity length changes	43
3.4	Measurements and analysis	48
3.4.1	Voice coil balancing	48

3.4.2	Actuator calibration	49
3.4.3	Measurement of transverse to longitudinal coupling	49
3.4.4	Analysis of the coupling level	52
3.5	Results	55
3.6	Outlook	56
4	The Sagnac speed meter experiment: introduction and technical design	59
4.1	Position and speed meters	59
4.1.1	Sensitivity of a position meter	59
4.1.2	Sensitivity of a speed meter	64
4.2	The Glasgow Sagnac speed meter experiment	70
4.2.1	Balanced homodyne detection	71
4.2.2	Suspensions	73
4.2.3	Actuation	74
4.2.4	Sensing and control	74
4.3	Topics of particular focus	75
5	Concept for the longitudinal control of the Sagnac speed meter experiment	77
5.1	Introduction	77
5.2	Control of the proof-of-concept experiment	78
5.2.1	Degrees of freedom	78
5.2.2	Sideband frequency	80
5.2.3	Control considerations	81
5.2.4	Required controller precision	82
5.2.5	Balanced homodyne detection	84
5.2.6	Op-amp noise	85
5.2.7	Technical noise	86
5.2.8	Sensor and actuator dynamic range considerations	88
5.3	Velocity control	97
5.3.1	Control loop	97
5.3.2	Low frequency noise projection with velocity feedback	97

5.4	Velocity-displacement control	101
5.4.1	Combined filter	101
5.4.2	Control loop	103
5.4.3	Low frequency noise projection with combined feedback	103
5.5	Matched filter	106
5.5.1	Noise budget	109
5.6	Experimental parameters	109
5.7	Summary and Outlook	110
6	Infrastructure for the control of a plate capacitor electrostatic drive with reduced seismic coupling	113
6.1	Electrostatic drives as actuators in suspended interferometer experiments .	113
6.1.1	Comb electrostatic drive design	114
6.1.2	Plate capacitor electrostatic drive design	114
6.2	Electrostatic drives for the Sagnac speed meter experiment	115
6.2.1	Maximum actuation requirements	116
6.3	High voltage amplifier design and implementation	118
6.3.1	Differential sending and receiving	118
6.3.2	Switchable dewatering	119
6.3.3	Choice of high voltage op-amp	123
6.3.4	Amplifier signal path	123
6.3.5	Practical and safety features	124
6.3.6	Transfer functions and noise measurements	126
6.4	Outlook	133
7	Conceptual longitudinal sensing scheme for the low frequency Einstein Telescope detector	137
7.1	The Einstein Telescope facility	137
7.1.1	New facility	139
7.1.2	Xylophone configuration	139
7.2	Control challenges with the Einstein Telescope	141
7.2.1	Signal recycling and resonant sideband extraction	141

7.2.2	Resonant sideband extraction in ET-LF	142
7.3	Longitudinal control of a dual-recycled Fabry-Perot Michelson interferometer	143
7.3.1	The dc readout technique	143
7.3.2	Degrees of freedom	144
7.3.3	Decoupled control signals	147
7.4	Longitudinal sensing scheme for ET-LF	150
7.4.1	Scope and method	150
7.4.2	Control sideband frequencies	151
7.4.3	Schnupp asymmetry	154
7.4.4	Dark fringe offset	155
7.4.5	Power in each light field	157
7.4.6	Readout ports	158
7.4.7	Control signals	160
7.4.8	Sensitivity of the scheme	163
7.5	Outlook and future work	165
7.5.1	Optimising the sensing matrix	165
7.5.2	Switching between tuned to detuned operation	166
7.5.3	Sensing and control of seismic and gravity gradient noise	166
7.6	Summary	167
8	Conclusions and future work	169
A	Interferometric fundamentals	173
A.1	Interferometer fields	173
A.2	Simple cavities	174
A.2.1	Cavity figures of merit	177
A.3	Signal sidebands	179
B	Control in the frequency domain	183
B.1	Signal to noise ratio	183
B.2	Frequency representation of signals and noise	184
B.2.1	Fourier transform	184

B.2.2	Spectral density	184
B.2.3	Estimation of spectral density	185
B.3	Root mean square amplitude	186
B.4	Control loops	187
B.4.1	Control loop figures of merit	188
B.4.2	Control bandwidth	188
B.4.3	Stable loops	189
C	Interferometer modelling	191
C.1	Software	191
C.1.1	Finesse	191
C.1.2	Optickle	191
C.2	Modelling an interferometer	192
C.2.1	Optics	192
C.2.2	Propagation	193
C.2.3	Fields	194
C.2.4	Drive and field maps	194
C.2.5	Calculation of field amplitudes	195
C.2.6	Probe signals	196
C.2.7	Calculation of transfer functions	196
C.2.8	Probe quantum noise and sensitivity	196
C.3	Calculation of field transfer matrices in Optickle	197
	Acronyms	199
	Bibliography	203

List of tables

3.1	Design parameters of the waveguide mirror for the experiment to measure transverse to longitudinal coupling	40
3.2	Cavity parameters	46
3.3	Spot positions on the end test mass	51
3.4	Distributions assumed for each of the free parameters in the waveguide model	54
5.1	Noise categories within the Sagnac speed meter experiment	97
5.2	Assumed experimental parameters for the Sagnac speed meter control loop modelling	111
6.1	Plate capacitor and optic parameters for the Sagnac speed meter experiment	116
6.2	Truth table for digital switching of dewhitening filters in the high voltage amplifier	122
6.3	Performance specifications for various high voltage operational amplifiers	124
7.1	Driving coefficients for each mirror and each degree of freedom of ET-LF .	147
7.2	Light power in ET-LF in the detuned configuration	158
7.3	Probes that sense the light fields propagating within ET-LF	159
7.4	Gradients of the error signals from each degree of freedom to each probe in ET-LF at dc	163
7.5	Normalised gradients of the suggested sensors to be used for the control of ET-LF's degrees of freedom	163
7.6	Updated parameters for ET-LF in the detuned configuration following the development of the conceptual sensing scheme	165

List of figures

1.1	Measured strain from the LIGO Hanford and Livingston detectors around the time of event GW150914	2
1.2	Plus and cross polarisations of a propagating gravitational wave	4
1.3	Simple Michelson interferometer	7
1.4	Worldwide interferometric gravitational wave detector network	8
2.1	Advanced LIGO noise budget	18
2.2	Some entry points for noise in a Michelson interferometer	19
2.3	Standard quantum limit and quantum noise with various input powers	25
2.4	Fabry-Perot Michelson interferometer	29
2.5	Dual-recycled Fabry-Perot Michelson interferometer	31
3.1	Propagation of light within a waveguide mirror	37
3.2	Brownian thermal noise in an Advanced LIGO style end test mass at room temperature versus that of a waveguide mirror at cryogenic temperature, as a function of transmissivity	38
3.3	Optical path length changes ζ_a and ζ_b due to transverse motion of a Littrow grating	38
3.4	Two ways in which light can be scanned across the surface of a waveguide mirror	39
3.5	Geometrical end test mass longitudinal coupling	41
3.6	Simulations of indicative cavity longitudinal error signals during end test mass rotation for different levels of waveguide mirror coupling	42
3.7	Experimental setup for measurement of waveguide transverse to longitudinal coupling	44

3.8	Positions of the magnets on the rear surface of the end test mass used in the waveguide experiment	45
3.9	Waveguide cavity frequency stabilisation servo open loop gain	48
3.10	Experiment to measure the effect of misaligned voice coil actuation	51
3.11	Change in force as a function of transverse displacement from voice coil axis	52
3.12	Posterior probability density distribution of the waveguide mirror coupling levels	56
3.13	Posterior probability density distributions of scaling and spot position standard deviation factors applied to the waveguide model	57
3.14	Measurements and simulations of the cavity length signal for spot positions with respect to the end test mass's centre of yaw	58
4.1	Response of a Fabry-Perot Michelson interferometer to differential arm cavity motion	62
4.2	Quantum noise of a Fabry-Perot Michelson interferometer at the output port	63
4.3	Sensitivity of a Fabry-Perot Michelson interferometer at the output port to differential arm cavity motion	64
4.4	Layout of a Fabry-Perot Michelson interferometer with a sloshing cavity .	65
4.5	Layout of a zero-area Sagnac speed meter	67
4.6	Response of a Sagnac speed meter to differential arm cavity motion	68
4.7	Quantum noise of a Sagnac speed meter at the output port	69
4.8	Sensitivity of a Sagnac speed meter at the output port to differential arm cavity motion	69
4.9	Sagnac speed meter experiment layout	71
4.10	Predicted quantum noise limited sensitivity of the Sagnac speed meter experiment	73
5.1	Transfer functions from important mirrors/combinations of mirrors in the Sagnac speed meter experiment to the balanced homodyne detector	79
5.2	Effective $L_{(-)}$ seismic noise contribution from M_7	80
5.3	Simplified layout of the Sagnac speed meter experiment including a basic velocity feedback loop	81
5.4	The frequency response of the Sagnac speed meter experiment's differential arm cavity degree of freedom to the balanced homodyne readout	83

5.5	Electronic schematic for the balanced homodyne readout	84
5.6	Electronic schematic for the measurement of noise from the balanced homodyne readout	85
5.7	Time series of the noise measured from the balanced homodyne readout electronics	86
5.8	Amplitude spectral density of the noise measured from the balanced homodyne readout electronics	87
5.9	The effect of whitening on a signal	89
5.10	Input and output whitening filters in the Sagnac speed meter experiment .	90
5.11	Anti-aliasing and anti-imaging filter transfer functions	91
5.12	End test mass suspension actuator control loop block diagram	92
5.13	Simulated end test mass suspension actuator open-loop transfer functions	95
5.14	End test mass suspension actuator control loop block diagram showing injection points	96
5.15	Modelled Sagnac speed meter experiment control loop using velocity feedback	98
5.16	Spectral density noise projection for $L_{(-)}$ using velocity feedback	100
5.17	Root-mean-square noise projection for $L_{(-)}$ using velocity feedback	100
5.18	Simplified layout of the Sagnac speed meter experiment including both displacement and velocity feedback paths	101
5.19	The frequency response of the differential arm cavity degree of freedom to the Pound-Drever-Hall readout	102
5.20	Modelled Sagnac speed meter experiment control loop using both displacement and velocity feedback	104
5.21	Spectral density noise projection for $L_{(-)}$ using both displacement and velocity feedback	105
5.22	Root-mean-square noise projection for $L_{(-)}$ using both displacement and velocity feedback	105
5.23	Simulated controller open loop gain	106
5.24	Matched filters to combine the balanced homodyne and Pound-Drever-Hall signals for different values of M_0 reflectivity	108
5.25	Control noise budget for the Sagnac speed meter experiment	110

6.1	Simulations of the actuation force produced by the proposed electrostatic drive design	117
6.2	Maximum end test mass displacement the electrostatic drive can create . .	117
6.3	Active inverting dewhitening circuit	120
6.4	Simulated dewhitening filter frequency response	121
6.5	Digital signalling between the control and data acquisition system and a channel within the high voltage amplifier	122
6.6	High voltage amplifier signal schematic	125
6.7	High voltage amplifier soft-start schematic	127
6.8	Minimum breakdown voltage between the two plates of the electrostatic drive for different separations	128
6.9	High voltage amplifier interlock schematic	129
6.10	Frequency response of the high voltage amplifier's channels with dewhitening enabled	130
6.11	Transfer functions of the high voltage amplifier input to monitor output with the dewhiteners on and off	131
6.12	High voltage amplifier cross-channel coherence	132
6.13	High voltage amplifier output noise	133
6.14	High voltage amplifier output projected end test mass displacement noise .	134
6.15	High voltage amplifier board, top view	135
6.16	High voltage amplifier board, bottom view	136
7.1	Sensitivity curves for the Einstein Telescope	140
7.2	Cavity lengths in a dual-recycled Fabry-Perot Michelson interferometer . .	144
7.3	Schnupp asymmetry and differential arm cavity offset in a dual-recycled Fabry-Perot Michelson interferometer	148
7.4	Reflected power from ET-LF as a function of power recycling mirror transmissivity	152
7.5	Power of the control sidebands in the cavities of ET-LF in detuned configuration as a function of Schnupp asymmetry	155
7.6	Power of the control sidebands in the signal recycling cavity of ET-LF in detuned configuration as a function of signal recycling cavity length . . .	156

7.7	Carrier power at the output port of ET-LF in detuned configuration as a function of differential arm cavity offset	157
7.8	Some available readout ports for sensing and control in ET-LF in detuned configuration	159
7.9	Sweeps through the zero-crossings of the chosen error signals in ET-LF . . .	162
7.10	ET-LF quantum noise limited sensitivity using the conceptual control scheme	164
A.1	Fabry-Perot cavity with its mirror input and output coefficients	174
A.2	Reflected field transfer function for a simple Fabry-Perot cavity	176
A.3	Resonant enhancement of input light in a cavity	178
A.4	Sideband structure for amplitude and phase modulation	181
B.1	The effect of windowing on a power spectral density estimate	186
B.2	A basic control loop	187
B.3	Limiting a servo to enhance gain in a certain band	189

Acknowledgements

No thesis represents solely the work of its author. The presented results are the culmination of the effort of many individuals, so this is my attempt to give thanks to the brilliant people I've had the pleasure of working with. Credits specific to the chapters are listed at the end.

To my supervisor Stefan (Hild), for your constant positivity and enthusiasm for all things, and for having faith that this programmer could achieve something useful in the lab. Your broad research interests seem to have rubbed off on me, if this eclectic thesis is anything to go by.

To Ken, for your approachable wisdom and uncanny knack for spotting problems before they arise, and for having patience to teach this boy who failed Circuits and Systems all about electronics.

To Bryan, for all of the help on the front line and for accepting the constant pestering from yet another student (almost like you're used to it). Maybe one day I'll get you those pints.

To Stefan (Danilishin), for the discourse and the cycle trips, and the heroic feat of translating the algebra of interferometers into the verbiage I so desired without resorting to finger puppets.

To the other marvellous postdocs I've had the pleasure of working with: Christian, Sebastian, Conor and Tobin. Thanks for the endless patience and help; I've learned so much from all of you.

Office buddies of 427a past and present, thanks for the absurd chat and the welcome distractions. The same sentiments go to my fellow denizens of the common room coffee club; the jitteringly strong coffee provided the lucid moments of clarity I sometimes needed.

To the rest of the IGR: it has been a pleasure to work in the wonderful environment you all help to create. Keep up the relaxed and open atmosphere, it's what makes the IGR so great!

To the Science and Technology Facilities Council for funding my PhD¹ and secondment to Hanover, and the Institute of Physics and Leibniz Universität Hannover for generous travel grants.

Finally to Mum, Andy, Jill and Jessica: thanks for the constant love and support. I promise I'll stop having fun and get a real job now.

¹Grant number ST/k502005/01.

Chapter 1 The data for the plot of GW150914 were provided by the LIGO Scientific Collaboration through its open science center.

Chapter 3 The experimental work was conducted in close cooperation with Bryan Barr. Stefan Hild and Ken Strain assisted with conceptual and technical development of the experiment. Neil Gordon, John Macarthur, Angus Bell, Borja Sorazu and Sabina Huttner assisted with construction and measurements. Russell Jones and Steven Craig provided help with the design and manufacturing of the reaction masses. Stephanie Kroker produced the waveguide mirror. Stefan Hild, Chris Messenger and Matt Pitkin assisted with the analysis of the results.

Chapter 4 The conceptual design of the experiment was developed by Stefan Hild, Christian Gräf, Sebastian Steinlechner and Ken Strain. The optical layout for the experiment was developed by Jan Hennig and Roland Schilling. Additional input to the experiment came from Jan Hennig, Alasdair Houston, Daniela Pascucci and Teng Zhang.

Chapter 5 The control design was developed in cooperation with Stefan Danilishin, Andreas Gläfke, Stefan Hild, Ken Strain, Bryan Barr, Kentaro Somiya and Harald Lück. The definition of degrees of freedom was jointly conducted in cooperation with Andreas Gläfke, Christian Gräf and Stefan Hild. The quantum noise model was developed with the help of Stefan Danilishin. The state-space ETM suspension model was developed by Ken Strain and Liam Cunningham, and based on the mechanical design developed by Russell Jones. The suspension control scheme was developed in cooperation with Ken Strain. The control software package *SimulinkNb* was developed by Christopher Wipf, who generously provided email support. The necessary modifications to the simulation tool *Optickle* were made with the help of Matt Evans.

Chapter 6 The ANSYS experiments were conducted by Christian Gräf and Christopher Mullen. The PA95-based high voltage amplifier and digital signalling design was made in cooperation with Ken Strain. The whitening filter frequency response was calculated with *LISO* by Gerhard Heinzl. Steven O’Shea, Colin Craig and Nicholas Scott provided assistance with the assembly of the amplifier.

Chapter 7 The interferometer simulations with *Finesse* were designed in collaboration with Vaishali Adya, Stefan Hild and Harald Lück, with technical support from Daniel Brown and Andreas Freise. The simulations with *Optickle* were conducted with the help of Tobin Fricke.

Other credits All of the plots in this thesis were produced with *Python* using the *Matplotlib*, *Numpy* and *Scipy* packages. Other diagrams were produced using *Inkscape* and *TikZ*. All of these tools are provided open source by a generous community of volunteers. Data and simulations were, in places, generated with *Matlab*, from Mathworks. The optical component graphics used in numerous figures throughout this work were provided by Alexander Franzen.

For Briagha and Felix.

Chapter 1

Gravitational waves

1.1 Event GW150914

Around 1.3 billion years ago a pair of black holes, one with 36 solar masses and the other with 29 solar masses, merged into a single black hole with 62 solar masses [1]. The missing energy equivalent to 3 solar masses was radiated away in the form of gravitational waves.

At 09:50:45 UTC on 14th September 2015, gravitational waves from the event reached the LIGO Livingston detector, perturbing the mirrors by 10^{-18} m and creating a signal large enough for the electronics controlling the interferometer to detect the ripple in spacetime more than 23 times above the background noise. Seven milliseconds later, the same wavefront passed the LIGO Hanford detector and moved the mirrors in the opposite direction. Meanwhile, computerised data analysis pipelines searching for coincident signals in each detector were running, and identified the event within a few minutes. Subsequent offline analysis using template banks representing waveforms produced by sources with different parameters were matched to the signal to calculate the most probable values. Given the significance of the signal above the detector's noise, and its fit to the templates, it was clear that the first detection of gravitational waves had been made.

The waveforms for event *GW150914*, as it has become known, are consistent with a binary black hole merger in that they were swept up in frequency before creating a loud “chirp” signal as shown in figure 1.1. The signal was only above each detector's background noise for the last few ms of this process, but it was consistent enough with theory that the collaboration could report a false alarm rate of less than 1 in 200 000. Not only did LIGO make the first observation of a gravitational wave, it also made the first detection of a binary black hole system and found the missing experimental proof of Einstein's theory of general relativity. The field opening up due to LIGO and the worldwide network of detectors in operation and under assembly, GEO 600, Virgo and KAGRA, represents an opportunity to

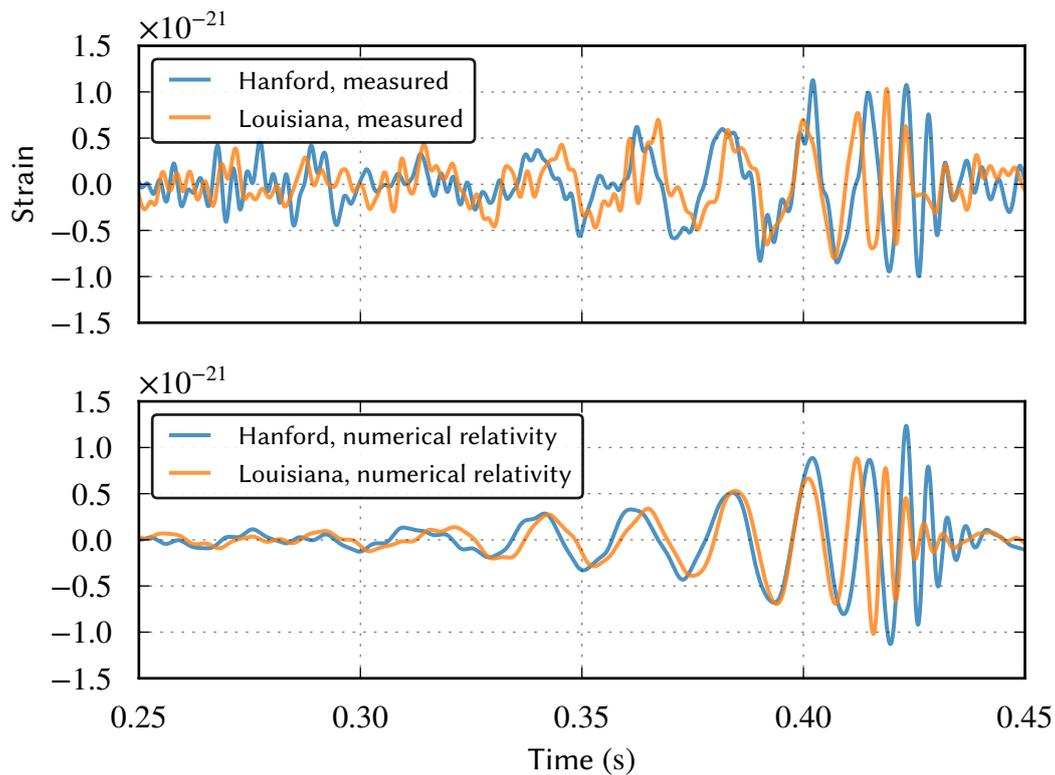


Figure 1.1: Measured strain from the LIGO Hanford and Livingston detectors around the time of event GW150914. Fast data pipelines work to find the best fitting template for the calibrated detector measurements from precomputed banks representing the various combinations of the system’s parameters. Beyond this initial identification, computationally intensive analysis of the data around the time of the candidate event is processed using a variety of techniques to identify the most likely source parameters, and these can later be matched to numerical relativity models. The numerical relativity models best fitting the analysed data are shown for each site in the lower panel. These curves were originally presented in ref. [1], and they have been passed through a 35 to 350 Hz band-pass filter to remove higher and lower frequency background noise unrelated to the signals.

study the universe in a completely new way.

1.2 Scientific outcomes from the first joint observation run

The first detection was made shortly before a science run in which the two LIGO detectors were kept at their most sensitive operating conditions as often as possible for a period of three months. In that time a second detection was made of a separate binary black hole merger, GW151226 [2], and taken alongside the first detection and another signal with lower significance it was possible for the LIGO Scientific Collaboration and Virgo collaboration to make some important scientific discoveries. Among the outcomes from the ob-

ervation run were that the first experimental tests of strong gravitational fields present during the events were still consistent with general relativity [3]; that “heavy” black holes exist in nature and that they probably formed in an environment with low metallicity [4]; and that the rate at which binary black holes merge in the known universe is constrained given the observations to between 9 and 240 per cubic gigaparsec per year [5], suggesting that the advanced detectors might see dozens more in subsequent science runs.

Beyond the detection of binary black hole mergers, other sources such as compact neutron star binaries have yet to be found but are predicted to exist producing signals within the sensitivity of the detector networks for future science runs [6]. Subsequent observations will hope to probe for new sources of gravitational waves, and with longer integration time it will be possible to set improved limits on the population of some of these astrophysical objects in our universe.

The long term goal of researchers in the field is to assemble a worldwide network of gravitational wave detectors to compliment the existing network of electromagnetic (EM) telescopes to facilitate *multi-messenger* astronomy. Such a network would be able to localise the sky position of incident gravitational waves well enough to allow for *EM follow-up* [6, 7], whereby the data from both optical telescopes and gravitational wave detectors can be combined to probe the astrophysics of the events in unprecedented detail and to identify and learn more about the nature and origin of their host galaxies.

1.3 General relativity and gravitational waves

A consequence of Einstein’s theory of general relativity, gravitational waves are produced by changes in the quadrupole moment of mass distributions such as the presence of non-spherical asymmetries in spinning objects or pairs of objects coalescing with elliptical orbits. The effect a gravitational wave has on spacetime as it propagates is to stretch it in one direction whilst contracting it in another. This *strain* can be expressed as a linear combination of “plus” and “cross” polarisation terms, shown for an initially circular ring of test particles on a 2D plane as a function of phase angle in figure 1.2.

Although in principle gravitational waves can be produced by all massive bodies, gravitational waves from Earth-bound objects, including the Earth itself, are not even remotely detectable. The strain in spacetime produced by such objects is so weak that there is no hope for us to make such a detection with any known technology. A good estimate for the strain, h_0 , produced by a pair of rotating objects is given by [8]

$$h_0 \lesssim \frac{2G (Mv^2)_{\text{nonspherical}}}{c^4 r}, \quad (1.1)$$

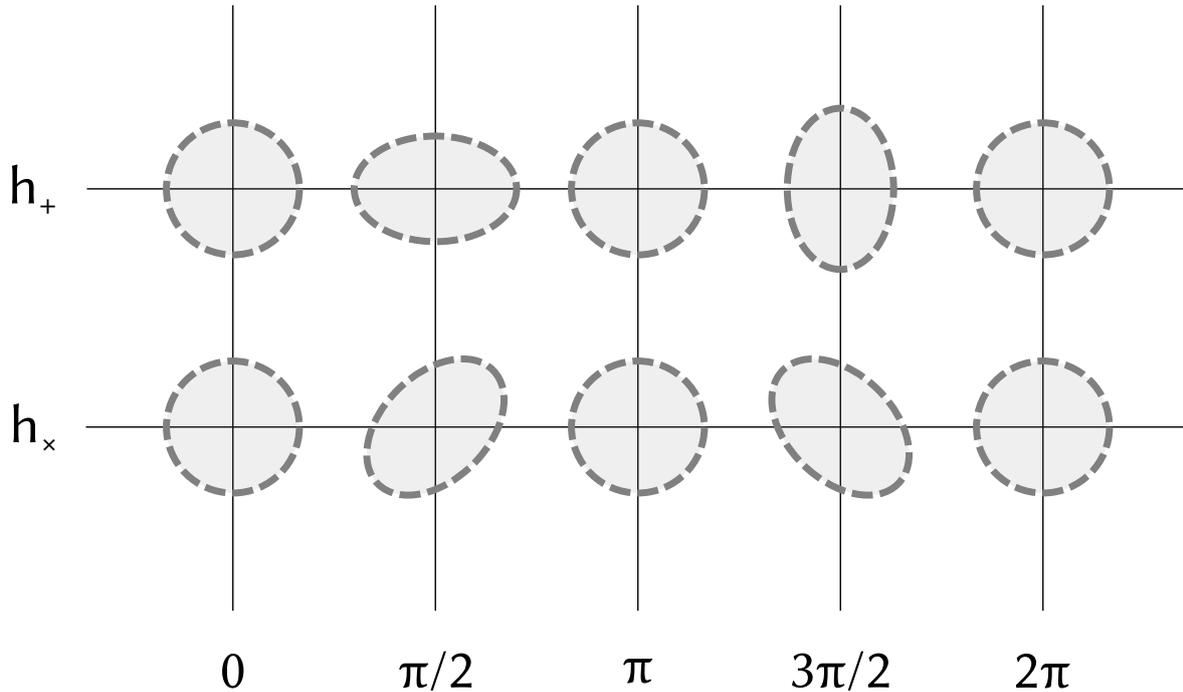


Figure 1.2: Plus and cross polarisations of a propagating gravitational wave. As the wave travels, shown in this depiction perpendicularly to the plane of this page, it stretches spacetime in one direction whilst contracting it in the other in an elliptic behaviour. A gravitational wave can be described as a linear combination of the two polarisations.

where G is the gravitational constant, $(Mv^2)_{\text{nonspherical}}$ is the kinetic energy associated with the non-spherical parts of the source required for the creation of gravitational waves, c is the speed of light and r is the distance between the source and the detector. To get an idea of what the strain would be for man-made sources, we can consider as in ref. [8] the case of two cars of mass $M = 10^3$ kg attached to opposite ends of a rod of length $d = 10$ m, spinning about its centre in a centrifuge at a frequency of $f = 10$ Hz. The tangential velocity of the cars will be around $2\pi f d \approx 600 \text{ m s}^{-1}$, about as fast as a modern fighter jet. Placing the detector one wavelength away, and using equation (1.1), the strain turns out to be around 4×10^{-43} . To be able to detect such a strain the current most sensitive detectors, Advanced LIGO, would require an improvement in sensitivity of 20 orders of magnitude, which is clearly ludicrous.

A pair of solar-mass objects orbiting each other at 100 Hz within 50 Mly produces a strain of only one part in 10^{21} , which is an amount only now detectable after decades of detector development. It is only the waves produced by the most violent redistribution of matter in the heaviest, most compact systems in the universe which we have any chance of detecting: binary black holes, compact binary neutron stars and core-collapse supernovae amongst others. Even then, gravitational radiation is only produced by the presence of a changing quadrupole moment and so only a subset of sources that happen to be in coalescence or contain surface asymmetries produce waves we have the ability to detect.

1.4 Development of the gravitational wave detector

Using the *local Lorenz* gauge an incident gravitational wave can be described as a change in the distance separating two reference points in spacetime, and so a measurement of the length between pairs of test masses placed at the different points on the edges of the ellipses shown in figure 1.2 can be made to infer the presence of passing gravitational waves. Given the behaviour of the propagating waves, the primary degree of freedom they excite in such an apparatus is the differential mode of the distance separating the test masses, $L_{(-)}$, which can be defined in terms of the position of the test masses x_A and x_B as

$$L_{(-)} = \frac{x_A - x_B}{2}. \quad (1.2)$$

The strain of an incident gravitational wave, h_0 , can be determined from the measured differential change in length between the test masses given the distance nominally separating the test masses, L :

$$h_0 = \frac{L_{(-)}}{L}. \quad (1.3)$$

1.4.1 Resonant bars

The first attempts to detect gravitational waves began with Joseph Weber's studies in the 1960s with his *Weber bar* [9]. This was a device developed to act as a direct strain meter, with incident gravitational waves exciting the separation of the material along the length of the bar. Piezoelectric sensors placed on the surface of an aluminium cylinder convert changes in length into electrical signals. While the expected change in length of such a cylinder from gravitational radiation would in most cases be tiny, the resonant frequency of the cylinder, typically in the kilohertz range, acts to enhance the amplitude of the length change at nearby frequencies. The sensitivity of such a bar as a function of frequency is determined in part by its quality factor (Q), with a necessary trade-off being made between peak sensitivity (high Q) and detection bandwidth (low Q). As sources of gravitational radiation are almost universally weak, the only reasonable hope of making such a detection is to choose a high Q material and hope for a signal at a favourable frequency.

Despite improvements over the following decades, the peak sensitivity of state-of-the-art resonant bar detectors was surpassed by *interferometric* gravitational wave detectors in 2003 [10] after it was shown that *second generation* detectors improving upon the initial designs would offer superior sensitivity across a much wider bandwidth [11]. The interferometer was first suggested as a means for gravitational wave detection shortly after the introduction of the Weber bar [12], but efforts to build prototypes and understand the significant sources of noise only gained momentum in the 1970s [13, 14].

1.4.2 The gravitational wave interferometer

The effect the strain shown in equation (1.3) has on the round trip time τ measured by light travelling between two test masses is, taking for example the plus polarisation component h_+ ,

$$\tau \approx \frac{2L}{c_0} - \frac{1}{2} \int_{t-\frac{2L}{c_0}}^t h_+(t) dt. \quad (1.4)$$

Given constant speed of light, c_0 , and that the gravitational wave h_0 is a sinusoidal function with angular frequency ω_g , this leads to a phase change $\delta\phi_{\text{GW}}$ as seen by the light with angular frequency ω_0 :

$$\delta\phi_{\text{GW}}(t) \approx h_0 \frac{\omega_0}{\omega_g} \sin\left(\omega_g \frac{\tau}{2}\right) \cos\left(\omega_g \left(t - \frac{\tau}{2}\right)\right). \quad (1.5)$$

This shows that the measurement of length between test masses with light offers the possibility to detect gravitational waves as phase fluctuations at the frequency of the signal. Gravitational wave induced phase modulation can also be described with the *transverse traceless* gauge as a change in the refractive index of the space between fixed test masses. The effect is equivalent [15], and so a change in length between the test masses can therefore be represented as a change in the frequency of the light, $\Delta\omega$, with respect to the light's nominal frequency:

$$h_0 = \frac{L_{(-)}}{L} = \frac{\Delta\omega}{\omega_0}. \quad (1.6)$$

From equation (1.5) the gravitational wave's *modulation depth* (see appendix A.3) can be approximated to

$$h_0 \frac{\omega_0}{\omega_g} \sin\left(\frac{\omega_g L}{c_0}\right), \quad (1.7)$$

showing that, in principle, the longer the distance separating the test masses the more modulation an incident wave will impart to the light and the stronger the signal will be. This is true up until the point at which the sine wave is maximum, $\frac{\pi}{2}$ (or, as with the design of radio antennae, the point at which the length is one quarter of the incident wavelength), and so to be optimally sensitive to a 100 Hz gravitational wave the distance between the test masses must be around 750 km which is clearly impractical for ground-based experiments. In section 2.3 we will discuss techniques to avoid the need for such long baselines.

Figure 1.3 shows the Michelson interferometer topology that all current detectors are based on. Coherent light from a laser is incident at the input to a beam splitter whereupon it is coupled into two arms with highly reflective mirrors at each end. Perpendicular arms are most sensitive to the differential change in length in spacetime created by gravitational waves as shown in figure 1.2. The light recombining at the beam splitter contains phase

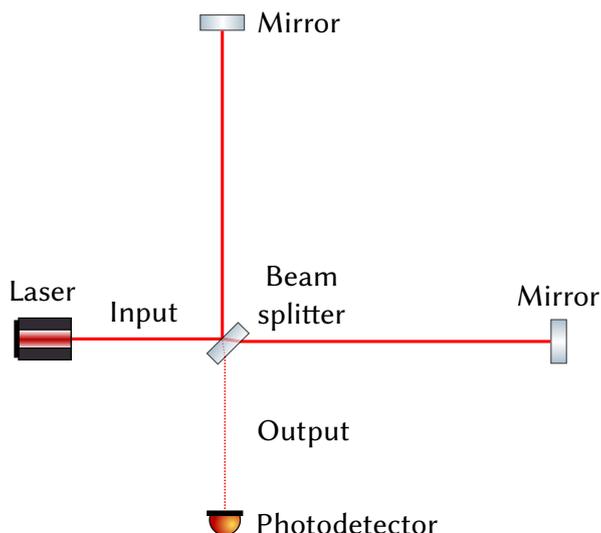


Figure 1.3: The simple Michelson interferometer used since the famous Michelson and Morley experiments of the 1880s. This version uses a laser as a source of monochromatic input light.

fluctuations from the motion of each test mass with respect to the beam splitter. If the arm lengths are arranged in such a way as to cancel the round-trip phase accumulation at the beam splitter in the absence of gravitational wave signal, the phase change between the arms due to a signal will appear at the beam splitter's *output* port where it can be measured by a photodetector.

Although this simple picture provides the foundation for the behaviour of the Michelson interferometer as a gravitational wave detector, the operation of a real detector is more complex and requires its own discussion. Chapter 2 will introduce the sensitivity improvements made to the Michelson interferometer and the research into the main sources of noise affecting its sensitivity over the course of the last 40 years of development.

1.5 Current and future interferometric detectors

As of the time of writing the Advanced LIGO detectors in the USA are online and commissioners are working towards reaching the design sensitivity. Advanced Virgo, situated in Italy, is due to begin science operations towards the end of 2016, with the KAGRA detector in Japan due to follow in 2019—these are the *second generation* detectors. GEO 600 in Germany has been operational in the years since the initial detectors stopped for upgrades, and is now transitioning into a detector-scale prototype facility. Planning is also under way to build an Advanced LIGO detector in India. The eventual network of detectors is shown in figure 1.4¹.

Beyond this generation, plans are afoot to build facilities which will push the sensitivity

¹LIGO India's exact location is yet to be decided.

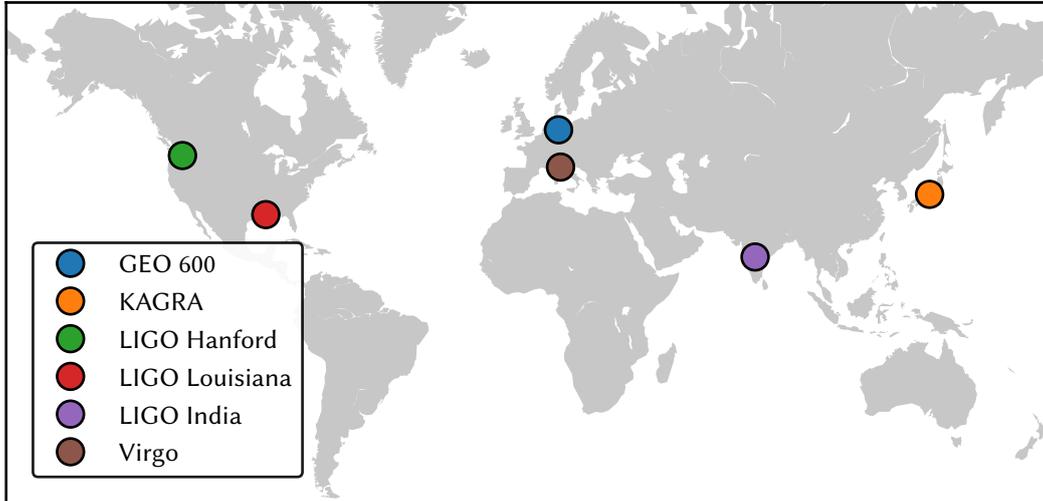


Figure 1.4: Worldwide interferometric gravitational wave detector network. GEO 600, LIGO Hanford and LIGO Livingston are operational, while Virgo and KAGRA are being commissioned and LIGO India is under construction. The locations of the Einstein Telescope and LIGO Cosmic Explorer are as yet undecided. The long baselines between sites assists with sky localisation.

of the ground-based gravitational wave detector to the limit, with the so-called *third generation* detectors. A European collaboration is working towards the *Einstein Telescope* [16] and the LIGO Scientific Collaboration is working towards *LIGO Cosmic Explorer* [17, 18]. Efforts are also under way to complement these detectors with a space-based counterpart with significantly enhanced low frequency sensitivity, *eLISA* [19]. Together, the network of ground- and space-based detectors will have unprecedented sensitivity from frequencies of mHz to kHz, providing an ability to study the universe in unparalleled fidelity.

1.6 Thesis structure

This thesis outlines work conducted with the goal of improving the sensitivity of future ground-based gravitational wave detectors.

Chapter 2 introduces some theoretical foundations and motivation for the work presented in the rest of this thesis. Chapter 3 presents an investigation into *waveguide* mirrors which offer a large potential improvement in Brownian thermal noise over conventional dielectric

mirrors used in existing detectors. One downside is the potential presence of a coupling effect between transverse motion and reflection phase. This chapter presents an experiment conducted to measure this coupling in order to give a clearer picture of this mirror's potential use in future gravitational wave facilities.

Chapters 4 to 6 present experimental research into a new type of gravitational wave interferometer: the Sagnac speed meter. Chapter 4 introduces the concept in more detail and presents an overview of an ongoing proof-of-principle experiment taking place in Glasgow. Chapter 5 highlights an important technical problem with the Sagnac speed meter configuration which is not present with current detectors: that the controller cannot determine the displacement of the cavity mirrors at low frequencies, leading to loss of sensitivity. A solution to the problem is presented through the modelling of the complete control system using knowledge of the response and noise of the apparatus as well as estimates for noise in the experiment as fully assembled, backed by numerical simulations. Finally, chapter 6 outlines the architecture and construction of experimental apparatus to test a new actuator design to be used in the Sagnac speed meter experiment: a plate capacitor electrostatic drive. Designs and tests of a high-voltage amplifier to create the required test mass actuation are presented.

The main body of the work concludes with chapter 7 where the current state of the sensing and control design for the low frequency interferometer as part of the planned Einstein Telescope facility is presented. This interferometer is to be primarily sensitive to frequencies below 10 Hz where existing detectors are dominated by seismic noise. Here, the sensing, controls and actuators found in the current generation of detectors are revisited through the use of numerical simulations.

Finally, the appendices provide additional information for the enthusiastic reader to support the main work. Appendix A provides a mathematical description of a basic interferometer and derives some useful figures of merit used throughout the work to describe interferometers. Appendix B discusses some aspects of controls to complement the main text. Appendix C discusses the differences between the two main numerical simulation tools used for the presented work in chapters 4, 5 and 7. A conclusion is provided in chapter 8.

Chapter 2

Sensitivity and noise in gravitational wave interferometers

To achieve maximum sensitivity in an interferometric gravitational wave detector to a particular type of signal the parameters of the optics, arm lengths and light fields must be considered alongside the characteristics of the signals and noise and the controllability and robustness of the resultant design. This chapter describes some of the considerations to be made in the design of detectors to provide the basis on which the rest of this thesis will build. Section 2.1 details the state in which an interferometer must be brought in order to be sensitive to gravitational waves and the means of keeping it there; section 2.2 introduces the limiting noise sources in ground-based gravitational wave detectors and the physical processes at play; section 2.3 discusses ways to improve the sensitivity of interferometers in the frequency bands of interest; and section 2.4 introduces concepts in order to reduce the most challenging noise source arising from the quantum nature of light.

2.1 Interferometer foundations

The effect that the output light from an interferometer has on a sensor (e.g. a photodetector) as some variable is modulated is termed its *response*. As discussed in chapter 1 the most important response to consider in gravitational wave interferometry from an astrophysical perspective is that of the differential motion of the arms (DARM) to the sensor at the output port. The response has a dependence on the input light power but it varies as a function of frequency due to the presence of additional *cavities* used to enhance or suppress the response in a given frequency band.

2.1.1 Measurement of interferometer length fluctuations

The complex-valued electric field amplitude of an electromagnetic wave propagating in time and space, E , can be expressed as

$$E = E_0 e^{i(\omega t - kx)}, \quad (2.1)$$

where i is the imaginary unit, ω is the wave's angular frequency, t is time, $k = \frac{2\pi}{\lambda}$ is the wave number and x is the coordinate in the direction in which E is measured. An arbitrary phase offset defined with respect to some point is contained within the complex-valued maximum field amplitude, E_0 .

Typically the underlying amplitude of a particular interferometer signal can only be inferred from the light power measured by a sensor. A simple example is the measurement of mirror displacement in a Michelson interferometer via the photocurrent output of a photodetector. The measured power P in this case would be

$$P = E^* E, \quad (2.2)$$

where $*$ denotes the complex conjugate.

Equation (2.1) can be simplified to a sinusoidal function with real maximum field amplitude E'_0 and phase offset ϕ :

$$E' = E'_0 \cos(\omega t - kx + \phi), \quad (2.3)$$

and in this way we can express the measured power as the square of the real field amplitude, i.e. $P = E'^2$.

Assuming that laser light with amplitude described by equation (2.3) is incident upon the beam splitter shown in the Michelson interferometer in figure 1.3, the light returning to the beam splitter having reflected from the north and east arms, n and e , respectively, would be

$$E'_n = -\frac{E'_0}{\sqrt{2}} \cos(\omega t - 2kL_n) \quad (2.4)$$

$$E'_e = \frac{E'_0}{\sqrt{2}} \cos(\omega t - 2kL_e), \quad (2.5)$$

where L_n and L_e are the two arm lengths. We employ a particular reflection phase convention such that a negative coefficient is gained on the light reflected from one side of the beam splitter, to conserve energy (see appendix C.2.1.1). We can also express L_n and L_e in

terms of the average arm length $L = \frac{L_n + L_e}{2}$ and differential length $\delta L = L_n - L_e$:

$$E'_n = -\frac{E'_0}{\sqrt{2}} \cos\left(\omega t - 2k\left(L + \frac{\delta L}{2}\right)\right) \quad (2.6)$$

$$E'_e = \frac{E'_0}{\sqrt{2}} \cos\left(\omega t - 2k\left(L - \frac{\delta L}{2}\right)\right). \quad (2.7)$$

The superpositions of the light from the arms leaving the beam splitter towards the input laser, E'_{in} , and the light leaving at the output port, E'_{out} , are then

$$\begin{aligned} E'_{\text{in}} &= \frac{E'_e - E'_n}{\sqrt{2}} \\ &= E'_0 \cos(\omega t - 2kL) \cos(k\delta L) \end{aligned} \quad (2.8)$$

$$\begin{aligned} E'_{\text{out}} &= \frac{E'_e + E'_n}{\sqrt{2}} \\ &= -E'_0 \sin(\omega t - 2kL) \sin(k\delta L). \end{aligned} \quad (2.9)$$

A real photodetector is not quick enough to measure changes in intensity at the frequency of the light. Instead, it sees the time averaged square of the field. The photodetector power at the output as a function of δL , P_{out} , is

$$P_{\text{out}}(\delta L) = \frac{P_0}{2} (1 - \cos(2k\delta L)), \quad (2.10)$$

where P_0 is the power of the incident laser light, showing that the signal from the differential arm length is encoded in the power of the light present at the output port. Note that implicit in this derivation is the assumption that the arms are both perfectly reflective. When the optics within the interferometer have different reflectivity, the calculation becomes more complicated and it is sometimes more practical to use a simulation tool, as discussed in appendix C.

2.1.2 Optimal operating point

The phase change created by the difference in the lengths of the arms shown in equation (2.10) as $k\delta L$ can be expressed as a combination of a *static* tuning α and the phase change created by incident gravitational waves, $\delta\phi_{\text{GW}}$, i.e.

$$k\delta L = \alpha + \delta\phi_{\text{GW}}. \quad (2.11)$$

The static tuning α is the differential arm phase at which the interferometer is nominally

held. In experiments where sensitivity can be sacrificed for simplicity, often it is practical to keep the interferometer in the state commonly referred to as “half way up the fringe”. Here, the interferometer’s arms are nominally tuned 45° out of phase such that the output signal oscillates about the midpoint between crest and trough of the superposition waveform at the output. As the gradient is steepest at this point, any small changes to the relative arm length of the Michelson interferometer result in a significant difference in power at the photodetector. This operating point, however, is not optimal in terms of *sensitivity* to arm length fluctuations. The static light power at this operating point contributes significant *shot noise* at the output. Its power spectral density is defined as [20]

$$S_{\text{shot}} = 2hf_0P_{\text{out}}, \quad (2.12)$$

where h is Planck’s constant and f_0 is the light frequency. The optimally sensitive operating point is therefore not simply one which maximises the signal gradient, but rather one which maximises the *signal-to-noise ratio* (SNR). It turns out that the reduced signal in the case of the operating point close to the *dark fringe*, where light from the two arms interferes destructively, is more than compensated for by the lack of shot noise such that the overall sensitivity is better. Interferometer operation near the dark fringe is the basis of *dc readout*, described in section 2.1.3.2, which is the standard measurement technique for all current generation detectors.

2.1.3 Readout

Note that the light power at the output shown in equation (2.10) as a sinusoidal function of the change in arm length is symmetrical and so displacements $\pm\delta L$ yield identical changes in light power. As the interferometer must be held at the dark fringe in order to maximise sensitivity, the controller requires a *bipolar* error signal providing a different response for motion in different directions. The purpose of the *readout* technique is to facilitate a bipolar error signal. Another benefit certain types of readout can provide is access to the displacement information in a particular *quadrature* of the output light. Signal (and noise) can in general be encoded in both the amplitude and phase of the light, representing the light’s quadratures. When the ratio between optical power and mirror mass is high this information is primarily contained within the phase quadrature, and when significant optomechanical interactions are present either with lighter mirrors or higher laser power the information can be encoded as a linear combination of the phase and amplitude quadratures. Some readout techniques facilitate arbitrary readout quadratures where the signal can be maximised with respect to the noise, while others have a quadrature fixed by the interferometer parameters.

There are two common types of readout technique. *Heterodyne* readout involves the use of a second light frequency used as a *local oscillator* for the primary light frequency that is resonant within the interferometer. This was one of the first techniques used to control laser interferometric gravitational wave detectors [21], but due to the presence of *cyclostationary* noise [22] and challenges related to the creation of a stable local oscillator frequency this has largely been superseded by *homodyne* techniques as the DARM sensor. Homodyne readout involves the use of the carrier as both a signal field and local oscillator, and this leads to some cancellation of noise sources common to the carrier at the expense of additional technical complexity. These techniques are discussed in greater detail in the following subsections.

2.1.3.1 Heterodyne readout

When light with multiple frequency components is incident upon a photodetector the resulting electrical signal shows the *beat signal* between the two components. Assuming that a photodetector has an incident electric field amplitude composed of two frequency components, we get [23]

$$E' = E'_0 \cos(\omega_1 t) + E'_0 \cos(\omega_2 t), \quad (2.13)$$

where ω_1 and ω_2 are the two frequencies and t is time. The photodetector measures the power of the field, P :

$$P = E'^2 = E_0'^2 (\cos^2(\omega_1 t) + \cos^2(\omega_2 t) + \cos((\omega_1 + \omega_2)t) + \cos((\omega_1 - \omega_2)t)). \quad (2.14)$$

If the difference frequency $\omega_1 - \omega_2$ in equation (2.14) is small, it can be measured by the photodetector. Most heterodyne techniques involve the phase modulation of a single carrier which creates a series of sidebands offset in frequency (see figure A.4) from the carrier that beat together at the photodetector. Different resonant conditions for the sidebands with respect to the carrier allow some to act as phase discriminants for others, and with suitable *demodulation* at the photodetector these can be used to sense displacement in the interferometer arms.

2.1.3.2 Homodyne readout

One way in which to picture homodyne readout is as a heterodyne readout with $\omega_1 = \omega_2$. It is possible to create a homodyne local oscillator by using a second laser with identical frequency to the first, though it is usually beneficial to use the same laser to benefit from coherent noise cancellation.

The first large scale application of homodyne readout in gravitational wave detectors was *dc*

readout [24], where a detuning (*dc offset*) is intentionally created within the interferometer's arms to allow for some of the carrier light to appear at the output port where it acts as a local oscillator to the rest of the carrier that contains the gravitational wave signal. This technique has the benefit that the local oscillator is filtered by the interferometer which suppresses certain types of noise, but it involves the intentional introduction of a classical light field at the output port. Another homodyne technique, *balanced homodyne detection*, involves the picking off of a fraction of the interferometer's input for use as a local oscillator. In this case, signal encoded in the light leaving the interferometer can be mixed with the local oscillator without the need for a dc offset. The dc readout technique is used in current generation detectors but for future interferometers it is possible that balanced homodyne detection will become the norm [25].

The technical implementation of dc readout is discussed in more detail in chapter 7, and balanced homodyne readout forms the basis for the experiment introduced in chapter 4.

2.1.4 Control

In order for the interferometer to be kept at its operating point the readout signal representing the positions of the mirrors (*test masses*) must be fed back to actuators. These actuators typically take the form of voice coils, piezoelectric stacks and other mechanical transducers. The ubiquitous technique for the control of the positions of mirrors is *linear negative feedback*, where the readout signal is passed through a *servo* which applies frequency-dependent filtering to enhance or suppress particular components and invert the signal before it is sent to the actuators. If the control system is designed to react quickly to test mass motion, the interferometer can be held almost exactly at the operating point where the error signal from the readout is nulled. Equation (2.10) shows that small arm displacements lead to linear changes in the output power, as is the case for other readout techniques. Effective control of the interferometer holds it within this linear region to ensure that the readout is maximally sensitive to the displacement of the arms.

Control strategies are discussed in greater detail in chapters 3, 5 and 7. Appendix B also introduces some background concepts useful for the understanding of the control strategies presented in these chapters.

2.2 Measurement noise in interferometers

The “signal” in an interferometer is the collection of electrical oscillations representing the particular variable of interest which, in most cases, represents the motion of the test masses in the arms. “Noise”, on the other hand, refers to the unwanted oscillations that appear in

the measurement independent of such a variable. The sensitivity of an interferometer is represented by the magnitude of the signal with respect to the noise, the SNR, introduced in the context of the operating point in section 2.1.2.

Gravitational wave interferometers are limited by a plethora of noise sources across the spectrum. The knowledge of the limiting noise sources gained from the science runs undertaken by the initial generation of interferometric detectors (LIGO, Virgo, GEO 600 and TAMA300) has fed in to the design of the current second generation.

The creation of *noise budgets* from theoretical descriptions and measurements of sources is a useful way to examine how noise influences the sensitivity of an experiment. The noise budget for Advanced LIGO's design configuration is shown in figure 2.1. At its most sensitive frequencies, Advanced LIGO is limited by *quantum* and *thermal* noise, while at lower frequencies the motion of the ground from seismic sources sets the limit. Careful design involving specially selected materials and techniques has reduced thermal noise arising from the mirror coatings and suspensions and technical noise associated with electronics and facilities. Quantum noise sets the fundamental limit given the available light power and mirror masses utilised in the detectors. In order to improve the sensitivity beyond the limit set by quantum noise, approaches that involve changing the nature of the quantum interactions within the interferometer have to be implemented. Important noise sources useful for the rest of this thesis are discussed throughout this section.

2.2.1 Noise arising from loss and uncertainty

Quantum theory showed that the universe contains a continuous spectrum of quantum fluctuations at all frequencies. Virtual photons are constantly created and annihilated in all space, albeit with an average energy of zero, producing the measurement uncertainty predicted by quantum mechanics. Collections of mirrors within interferometers create local filters of this quantum spectrum which allow a subset of vacuum modes to circulate. Virtual photons are able to enter the interferometer via its *loss points*, where light can escape the interferometer, just as virtual photons created within the interferometer are allowed to leave. As the vacuum fluctuations are uncorrelated with the motion of the test masses, non-unity reflectivity of optics, scattering and other photon loss effects within an interferometer lead to the intrusion of vacuum noise.

In lasers, a pumped electromagnetic field creates a state which can be used as the input light for an interferometer. This typically involves pumping the field into the *coherent* state in which the average laser amplitude and phase quadratures are matched, and noise arises from the presence of virtual photons with arbitrary amplitude and phase in the pumped field leading to an uncertainty in the number of photons output from the laser.

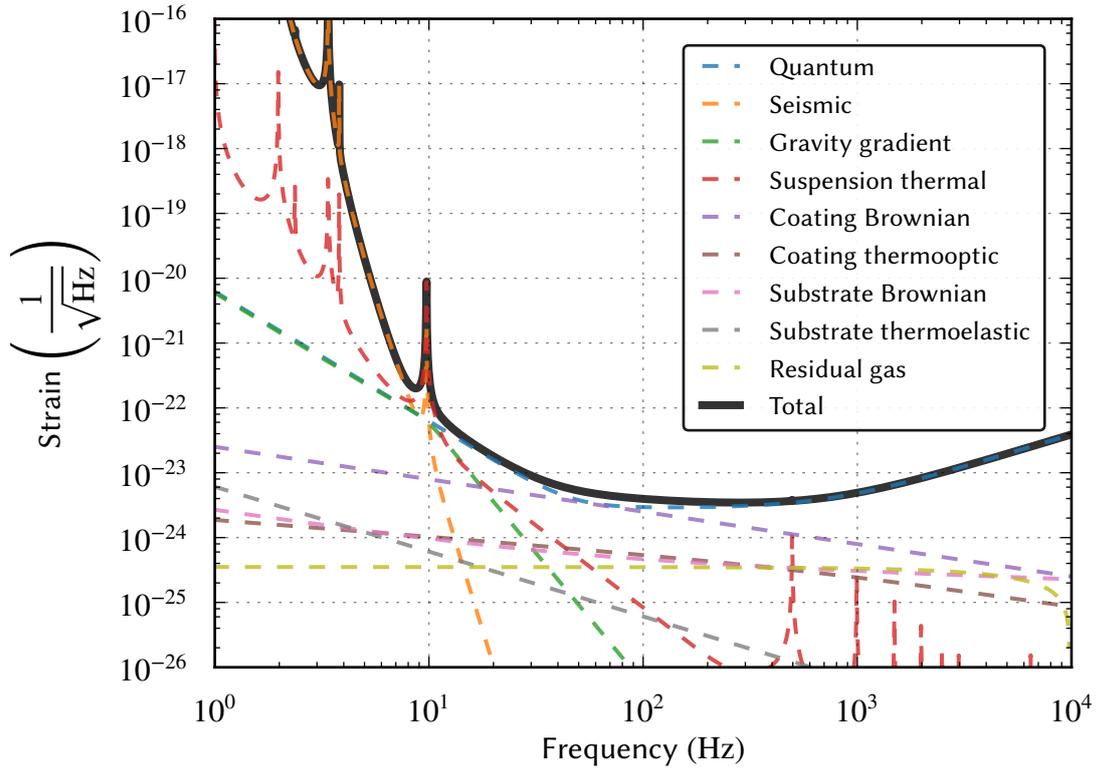


Figure 2.1: Advanced LIGO noise budget calculated with GWINC [26]. Greater sensitivity to gravitational waves is achieved by having lower residual strain noise. The incoherent sum of the noise sources leads to the overall sensitivity of the interferometer, and this is shown in black. All of the noise sources shown have some frequency dependence, and optimal sensitivity in a detector is reached by designing the experiment in such a way as to minimise the noise sources in the frequency band of interest. The creation of budgets like this from theoretical descriptions of noise sources is a useful way in which to understand how they affect the sensitivity.

Noise in interferometers does not arise solely from the quantum vacuum, however. In general, noise was shown by Callen and Welton to arise from loss processes quantified by their *fluctuation dissipation theorem* developed in the early 20th century [27], which showed that the noise power spectral density created due to fluctuations,

$$S_{\text{fluc}} \propto \frac{T}{Q}, \quad (2.15)$$

is related not only to the energy quantified by the temperature T , but also to the quality factor Q related to the lossiness of the material.

The effect of noise on the interferometer can be calculated by quantifying the magnitude of noise entering at a loss point and propagating this noise to the signal detection point where it can potentially mask the signal. The noise at a photodetector is then the sum of noise propagated from each point of loss to the readout point. The way in which some forms of noise can enter a Michelson interferometer and propagate to the output port is shown in

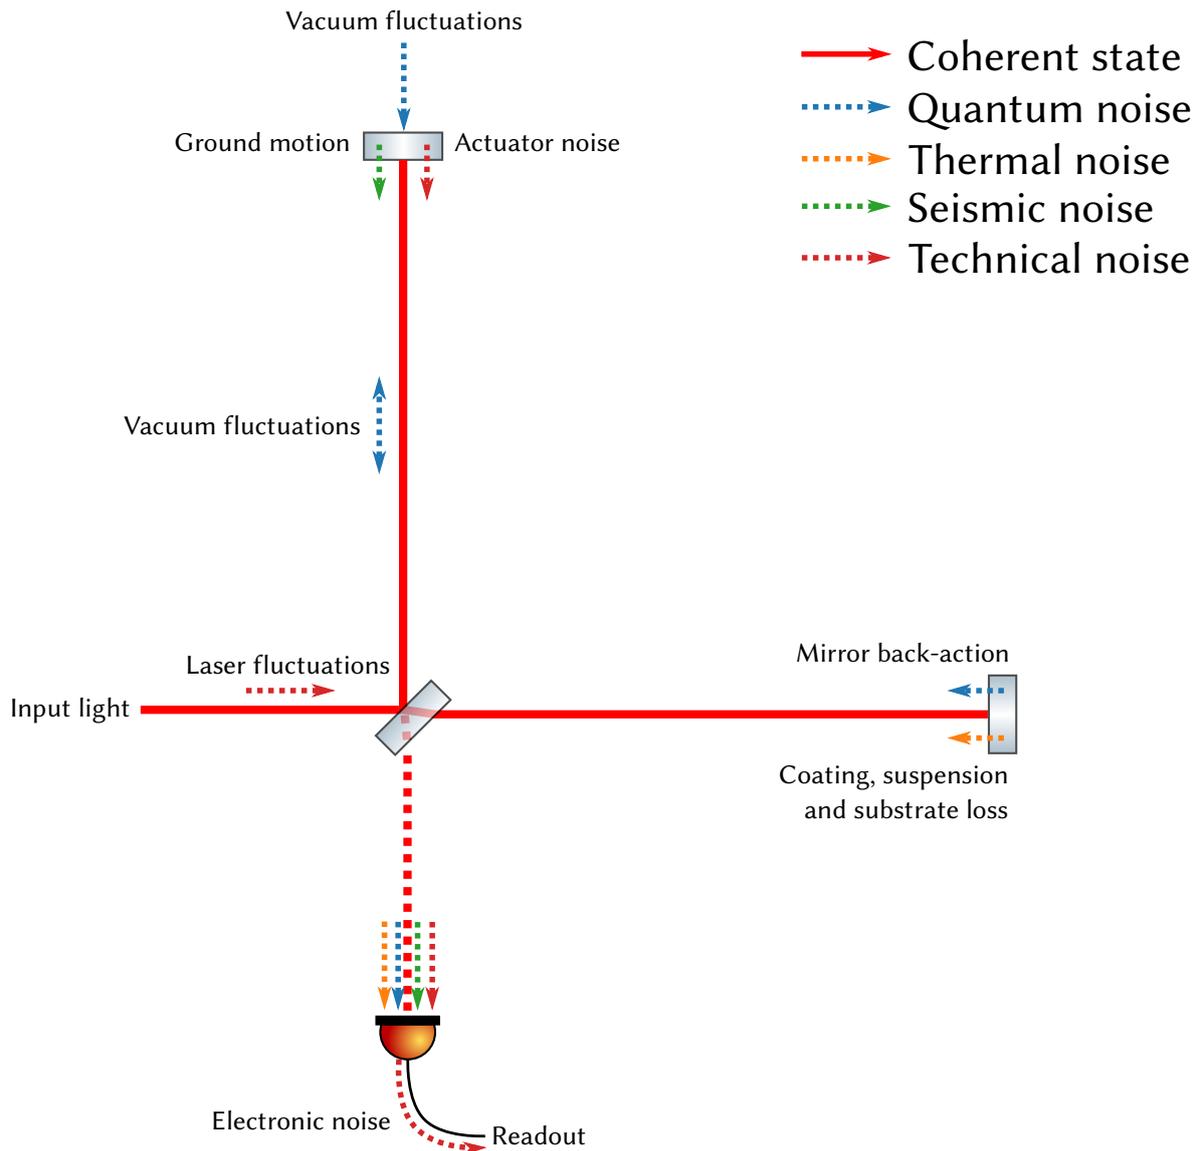


Figure 2.2: Some entry points for noise in a Michelson interferometer.

figure 2.2.

2.2.2 Thermal noise

Thermal noise arises from loss in materials used to reflect and focus light and to suspend test masses, where photons given a phase change due to thermal excitations are able to propagate to the sensors.

Thermal noise is quantified by a material's *loss angle*, which is the imaginary part of the Young's modulus relating applied stress to the corresponding strain of the material. Material with a high loss angle results in an applied stress creating an associated strain at a different time, and during this time the incident light can accumulate noise via thermal fluctuations of the material. The most significant thermal noise contributions in current

generation gravitational wave detectors arise from the test mass optical coatings and suspensions.

2.2.2.1 Coating thermal noise

In the conceptual design for the first generation of gravitational wave detectors such as GEO-600 [21] the designers were not aware that thermal noise associated with the reflective coatings on mirrors would play a significant role in the sensitivity of the interferometers. For a long time it was known that thermal noise would contribute to the sensitivity of the detectors, particularly from the bulk material forming the test masses, but it soon became clear as the detectors were being commissioned that thermal noise arising from the reflective mirror coatings would dominate the thermal noise associated with the test masses in the frequency band of interest despite forming only a tiny fraction of the test masses by volume. Investigations conducted by Harry et al. [28, 29], among others, concluded that mechanical loss present in the numerous dielectric coating stacks on the test masses required for high reflectivity led to Brownian noise creating a limit to the sensitivity of detectors across a wide range of frequencies. Contributions from thermoelastic noise, arising from the thermal expansion coefficient of the materials of the coatings [30], and thermorefractive noise, arising from the change in refractive index caused by fluctuations in the material's temperature [31], produce further noise which will become more important as coatings with improved Brownian noise are developed.

Over the past two decades, efforts have been made to both quantify and reduce coating thermal noise. Particular interest is being paid to the study of coatings for cryogenically cooled mirrors, such as the sapphire (Al_2O_3) test masses to be used in KAGRA [32]. A loss peak in the mirror material silica (SiO_2), for detectors until recently ubiquitous, occurs at low temperature. This makes the material unsuitable for cryogenic use, as mechanical loss will couple into the light within the interferometer and make its way to the detection port. Other materials such as sapphire do not feature this loss peak and provide lower thermal noise than silica at room temperature for a given mirror design. Coating noise is also proportional to temperature, so cryogenically cooled mirrors can offer better performance. Additionally, crystalline coatings made from compounds such as AlGaAs can offer future detectors a coating thermal noise reduction of up to 3 over the current state of the art [33] if technical challenges in their manufacturing can be overcome.

The dominant contribution to coating noise in current generation gravitational wave detectors, Brownian noise, has a power spectral density given by [28]

$$S_{\text{coating}} = \frac{2k_B T}{\pi^{3/2} f} \frac{1}{wY} \left(\phi_{\text{sub}} + \frac{1}{\sqrt{\pi}} \frac{d}{w} \left(\frac{Y'}{Y} \phi_{\text{para}} + \frac{Y}{Y'} \phi_{\text{perp}} \right) \right), \quad (2.16)$$

for Boltzmann constant k_B , temperature T , frequency f , beam size w , Young's modulus Y , loss angle ϕ and coating thickness d . The Young's moduli are split into components representing the coatings and substrate, Y and Y' , and the loss angles are split into parallel and perpendicular components in the coatings, ϕ_{para} and ϕ_{perp} , and substrate, ϕ_{sub} , respectively. The measurement and interaction between these components is an active area of research. Figure 2.1 shows coating Brownian noise jointly dominating the noise in Advanced LIGO at frequencies around 70 Hz.

A mirror topology which avoids the use of many alternating coating layers can potentially offer an improvement in noise performance. Mirrors employing grating structures can resonantly reflect light with less coating material than similarly performing dielectric mirrors [34], though at the expense of additional technical complexity in their utility in gravitational wave detectors [35]. Chapter 3 discusses a form of grating mirror for use in interferometers.

2.2.2.2 Suspension thermal noise

The test masses in audio-band gravitational wave detectors must be suspended from pendulum systems to filter ground vibrations, and current generation observatories (with the notable exception of KAGRA) utilise fused silica fibres, a technique pioneered for GEO 600 [36]. The reason for the use of this material is that the thermal noise present within the previously used steel loops was high enough to impart significant displacement noise to the test mass in the gravitational wave channel, with the noise becoming dominant at frequencies around 100 Hz where the interferometer would otherwise be most sensitive [37]. Due to its high quality factor, fused silica has reduced mechanical loss and therefore lower noise. Figure 2.1 shows that suspension thermal noise is no longer a dominant noise source, unlike in Initial LIGO.

As KAGRA will be a cryogenic detector, it does not gain the same noise benefit from using fused silica. Instead, it will use crystalline sapphire which offers similar noise performance at low temperatures.

At higher frequencies, suspension *violin modes* have a significant influence on the measured noise [38]. A violin mode with high quality factor can resonantly enhance noise such that it dominates all other sources in a narrow band at frequencies starting around a few hundred Hz¹. This is reduced through the use of heavier test masses, which push the violin mode frequencies higher, away from the detection band, and with special monitoring techniques [39].

¹Figure 2.1 appears to show that violin modes are not dominant, however the narrow linewidth of the noise is such that the resolution is insufficient to show the effect.

2.2.3 Quantum noise and the Standard Quantum Limit

Arising from the Heisenberg Uncertainty Principle, the quantum noise present within a classical interferometer² limits its sensitivity.

Classical laser light in the coherent state, approximating what a standard laser will output, contains equal fractional amplitude and phase uncertainties. The phase fluctuations appearing at the sensor used to measure the output of the interferometer and amplitude fluctuations interacting with the interferometer's test masses create noise at the measurement ports³. A fundamental limit to the sensitivity of classical interferometers arises from the combination of these two effects; this is described in more detail in section 2.2.3.3. The following subsections summarise results from ref. [40].

2.2.3.1 Quantum shot noise

As described in section 2.2.1, open ports in the interferometer allow vacuum noise to enter, and when this noise is measured by a photodetector it appears as *quantum shot noise*. The phase fluctuations upon the light produce a varying photocurrent due to the stochastic arrival of photons at the sensor. The displacement-equivalent power spectral density of shot noise in an interferometer is

$$\tilde{x}_{\text{shot}}^2 = \frac{\hbar c^2}{P\omega_0}, \quad (2.17)$$

in units of m^2/Hz , for power P and laser angular frequency ω_0 . As this noise arises from spontaneous creation and annihilation of photons in space, it is a statistical random process and so the spectral density has equal power at all frequencies; it is *white*. The strain-equivalent power spectral density is equation (2.17) normalised to the arm length L :

$$\tilde{h}_{\text{shot}}^2 = \frac{\hbar c^2}{P\omega_0 L^2}. \quad (2.18)$$

Since it scales with input power, the detrimental effect on the sensitivity due to phase uncertainty is mitigated by an increase in the classical light power injected into the interferometer.

2.2.3.2 Quantum radiation pressure noise

Despite being massless, photons impart momentum to mirrors upon reflection inversely proportional to their wavelength. The strongest effect this has on an interferometer is

²Note the misnomer: a *classical* interferometer can still be limited by *quantum* noise. The name refers to the readout technique, namely the measurement of classical light intensity to determine displacement.

³A derivation of both effects can be found in, for example, ref. [40].

via *dc radiation pressure*, which arises from the classical light power circulating within the interferometer. In a suspended interferometer this radiation pressure effect extends the microscopic arm cavity length, with the equilibrium point being defined by the equivalence of the radiation pressure force to the suspension's restoring force.

Quantum radiation pressure, on the other hand, arises from the fluctuating momentum imparted onto the test masses by fluctuations in the number of photons present within the interferometer from the laser and loss points. As with quantum shot noise this effect is related to the input power of the interferometer, but in this case fluctuations in the number of input photons creates amplitude noise that is transformed into equivalent strain noise via the dynamics of the mirror. Amplitude fluctuations upon the light beat with the classical field, creating a force noise. This fluctuating force changes the position of the mirror microscopically via its mechanical susceptibility and this appears as phase noise at the output port. As the spectrum of noise from virtual photons is white the energy imparted to the mirror is the same at all frequencies. The mechanical susceptibility of a suspended mirror follows an inverse square law in frequency above the resonant frequency, and so in terms of strain this noise source is most important at low frequencies. The radiation pressure noise power spectral density is given in this case by

$$\tilde{x}_{\text{rp}}^2 = \frac{P\hbar\omega_0}{c^2 m^2 \omega^4}, \quad (2.19)$$

with reduced mirror mass m and angular frequency of mirror oscillation $\omega = 2\pi f$. The reduced mirror mass is the effective mass of the mechanical mode, given in the case of a Fabry-Perot Michelson interferometer as

$$m = \frac{m_1 m_2}{m_1 + m_2}, \quad (2.20)$$

where m_1 and m_2 denote the individual cavity test masses.

The strain-equivalent power spectral density is

$$\tilde{h}_{\text{rp}}^2 = \frac{P\hbar\omega_0}{c^2 m^2 \omega^4 L^2}. \quad (2.21)$$

The strain amplitude noise, \tilde{h}_{rp} , is proportional to $\frac{1}{\omega^2}$ as expected from a free mass.

2.2.3.3 The Standard Quantum Limit

Note that equation (2.19) is proportional to power while equation (2.17) is inversely proportional to power. This implies the existence of a lower bound on the achievable sensitivity at a given observation frequency f in the case of uncorrelated shot and radiation pressure

quantum noise sources. This bound, known as the *standard quantum limit* (SQL) [41], is a direct consequence of the Heisenberg Uncertainty Principle in a continuous measurement of a test mass.

The SQL is the point at which the sum power spectral density of shot and radiation pressure noise is minimised, and this occurs when the individual components are equal. For each laser power there exists a single frequency at which the SQL can be reached. The SQL forms a sensitivity limit with amplitude spectral density proportional to $\frac{1}{f}$ which can only be surpassed with special, *sub-SQL* techniques. The presence of cavities in the arms of a Michelson interferometer (formed by placing an additional, partially reflecting mirror in each arm) can enhance the power available to be able to reach the SQL. In terms of the strain-equivalent power spectral density, the SQL is specified for two free test masses separated by a distance L by [42]

$$\tilde{h}_{SQL}^2 = \frac{8\hbar}{m\omega^2 L^2}, \quad (2.22)$$

with units of Hz^{-1} .

The strain-equivalent power spectral density noise for a Michelson interferometer with arm cavities can be written with respect to the SQL [43] as

$$S_{MI}^h = \frac{\tilde{h}_{SQL}^2}{2} \left(\frac{1}{\kappa} + \kappa \right), \quad (2.23)$$

where the SQL is reached only at a single frequency. The term κ is the (dimensionless) *opto-mechanical coupling factor* [43]:

$$\kappa = \frac{P_0}{P_{SQL}} \frac{2\gamma^4}{\omega^2 (\gamma^2 + \omega^2)}, \quad (2.24)$$

with P_0 the laser power at the test masses, P_{SQL} the laser power required to reach the SQL at the cavity pole frequency and γ the arm cavity half-bandwidth. P_{SQL} is given as [43]

$$P_{SQL} = \frac{mL^2\gamma^4}{4\omega_0}. \quad (2.25)$$

The effect of κ is described in more detail in section 4.1.1.

The SQL is a locus defined at all frequencies, while the spectral density of a quantum noise limited interferometer touches the SQL at only one frequency. By injecting more photons into the interferometer to carry more information regarding the motion of the mirrors, we see a smaller shot noise spectral density while we see a larger radiation pressure noise spectral density [44]. This situation is illustrated in figure 2.3 for different input powers.

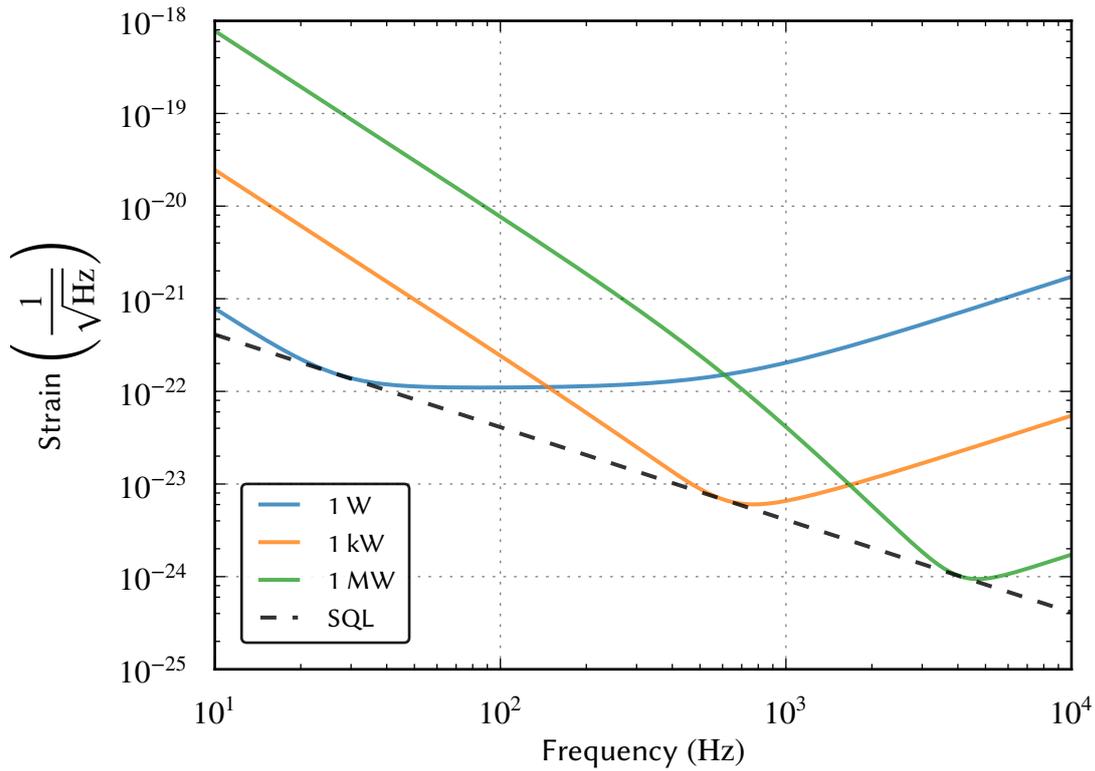


Figure 2.3: The SQL for a Michelson interferometer with arm cavities of length 1 km, mirrors with reduced mass 50 kg and optimal frequency 100 Hz, along with quantum noise limited sensitivity curves for three different intracavity powers. The effect of the cavity pole frequency is visible in the case of the blue curve. The higher the intracavity power, the higher the strain sensitivity can be pushed, but at the expense of higher radiation pressure noise and thus higher optimal frequency for a given interferometer configuration. Quantum non-demolition techniques can be used to surpass the SQL (see section 2.4).

An important distinction to make here is that the SQL is defined for *uncorrelated* shot and radiation pressure noise. Techniques exist in theory and practice to reduce overall noise by introducing correlations between the two noise components with so-called *quantum non-demolition* interferometry, and this is discussed in greater detail in section 2.4 and chapter 4.

2.2.4 Other fundamental noise

2.2.4.1 Seismic noise

The Earth's surface vibrates with a large amplitude and low frequency [16]. At around $10 \mu\text{Hz}$ tidal forces due to the gravitational interaction between the Earth and Moon⁴ dominate the spectrum producing displacements of up to $100 \mu\text{m}$ [45]. At around 0.15 Hz the swell of the ocean can be measured almost anywhere on the Earth, even far from coasts.

⁴ $10 \mu\text{Hz}$ is about one cycle per day, the same as the Earth's rotation.

These effects produce a large amount of displacement noise at low frequencies which must be filtered.

As seismic noise is large in amplitude, it is able to move test masses in interferometers far enough that they no longer fulfil the resonant condition and lose light power. In almost all audio-band interferometric experiments a large degree of isolation must be utilised to mitigate this seismic noise. In Advanced LIGO, active platforms sitting atop passive damping materials are used to reduce this noise. Test masses are also suspended from many pendulum stages to isolate higher frequencies such that by 10 Hz the ground motion is suppressed by more than 10 orders of magnitude [46].

Homogeneous, vertical surfaces do not couple vertical seismic noise into the gravitational wave channel horizontal to each test mass. Real suspended optics, however, contain imperfections in their manufacturing and couple a small amount of vertical motion into the horizontal direction. In addition, the curvature of the Earth over distances like the 4 km arms in Advanced LIGO mean that the local gravitational fields at the ETMs are not entirely aligned to those of the ITMs, and so to achieve cavity resonance the operating point requires a slight off-horizontal tilt which creates seismic noise coupling. In Advanced LIGO the requirement for vertical to horizontal coupling is to be below 1×10^{-3} .

2.2.4.2 Gravity-gradient noise

Changes in the density of the ground near the test masses created by seismic noise can couple to the gravitational wave channel via *gravity-gradient* noise, also referred to as *Newtonian* noise. No experiment has successfully been able to decouple this subtle effect from other sources of noise, but it is believed from extensive modelling effort that this noise source will represent a problem particularly for low frequency detectors such as ET-LF [16, 47]. Simulations have shown promise in subtraction of gravity-gradient noise inferred from a series of auxiliary witness sensors [48] as well as a benefit to shaping the profile of the ground near test masses [49].

Gravity gradient noise will be discussed in the context of ET-LF in chapter 7.

2.2.5 Technical noise

2.2.5.1 Laser frequency and intensity noise

A perfect laser would provide output at a single, well defined frequency. In reality such lasers do not exist and their outputs contain spectral impurities. As the laser wavelength is the “metre stick” by which we make displacement measurements in interferometers, it is very important to ensure that the laser’s wavelength, and therefore frequency, is well

defined. Frequency stabilisation control loops involving optics and electronics are usually necessary in high precision interferometric experiments.

Laser frequency noise affects the phase of the output light by creating beats between waves with different frequencies, created via thermal effects in the laser material. This can be expressed as a time-varying shift $\phi(t)$ in the underlying wave's phase. This phase transforms into frequency noise via the relation

$$\Delta f = \frac{1}{2\pi} \frac{d\phi}{dt}. \quad (2.26)$$

The spectral density of frequency noise can be calculated from the autocorrelation between a frequency fluctuation at time t and another at time $t + \Delta t$, but a simpler method is to realise that the laser is used to measure the cavity length by relating its change in frequency to the change in length via equation (1.6). Multiplying the relative frequency noise by the difference in arm lengths δL gives a first order estimate of the displacement noise created by fluctuations in the laser:

$$\tilde{x}_{\text{freq. noise}} = \delta L \frac{\Delta f}{f}. \quad (2.27)$$

The laser's intensity fluctuates due to similar mechanisms. Thermally driven misalignments within the laser can lead to scattering and the production of higher order modes, which reduce the intensity of the fundamental mode. This has a similar effect on the output as frequency noise, coupling relative intensity noise $\frac{\Delta P}{P}$ directly to the output via the microscopic offset from the dark fringe condition δl :

$$\tilde{x}_{\text{int. noise}} = \delta l \frac{\Delta P}{P}. \quad (2.28)$$

Equations (2.27) and (2.28) show that the level to which the dark fringe condition is satisfied determines the laser noise witnessed at the output. This is because of the cancellation at the beam splitter from noise in the two arms. Laser noise propagates to the beam splitter, where it is split between the arms. Matching the arms macroscopically cancels frequency noise and matching the arms microscopically cancels intensity noise.

2.2.5.2 Electronic noise

Johnson-Nyquist noise arises from loss within electronic conductors. The noise scales with resistance and is characterised in units of V^2/Hz by the equation

$$\tilde{v}_{\text{john. noise}}^2 = 4k_B T R, \quad (2.29)$$

where $\tilde{v}_{\text{john. noise}}^2$ is the voltage noise power spectral density, k_B is the Boltzmann constant and R is the electrical resistance. The Johnson-Nyquist noise from a resistor in the $\text{k}\Omega$ to $\text{M}\Omega$ range is comparable to the noise of some low-noise operational amplifiers at room temperature, and so care must be taken in the choice of passive and active components in the design of electronics to avoid introducing excess fluctuations.

Other electronic noise can arise in integrated circuits used as part of readout electronics in detectors. Current and voltage noise present at the input and output of devices such as operational amplifiers (op-amps) can become larger than the signals being amplified without careful selection of the device for the intended application. This is examined in more detail in sections 5.2.6 and 6.3.

2.2.5.3 Quantisation noise

The conversion of analogue signals to digital and vice versa for sensing and control involves the use of analogue-to-digital and digital-to-analogue converters (ADCs and DACs, respectively). Noise in ADCs and DACs arises from *quantisation error* ϵ_{ADC} , which is the mismatch between the underlying signal input or output and the level determined by the ADC or DAC [50]. By averaging over successive cycles, ADCs and DACs can make good approximations of the underlying signals, reducing the quantisation error in the case of the ADC to the interval $\epsilon_{\text{ADC}} \in \left(-\frac{\Delta}{2}, \frac{\Delta}{2}\right]$ [50], where Δ is the smallest voltage reference in an ADC with 2^b codes:

$$\begin{aligned}\Delta &= \frac{V_{\text{max}} - V_{\text{min}}}{2^b} \\ &= \frac{V_{\text{range}}}{2^b}.\end{aligned}\tag{2.30}$$

Noise from ADCs and DACs is minimised through selection of hardware with appropriate dynamic range for the signals to be sensed. The use of *whitening* techniques can also prevent the prevalence of quantisation noise; this is discussed in greater detail in section 5.2.8.1.

2.3 Sensitivity of the Michelson interferometer

The phase change due to gravitational waves appearing at the output is proportional to the power in the arms, and so greater input power leads to greater response at the output (given the caveats regarding quantum noise as discussed in section 2.2.3).

As shown in section 1.4.2, the arm length of a Michelson interferometer to provide optimal modulation upon the light field from a passing gravitational wave can be many hundreds of km for audio band signals. Furthermore, given the standard wavelength $\lambda_0 = 1064 \text{ nm}$

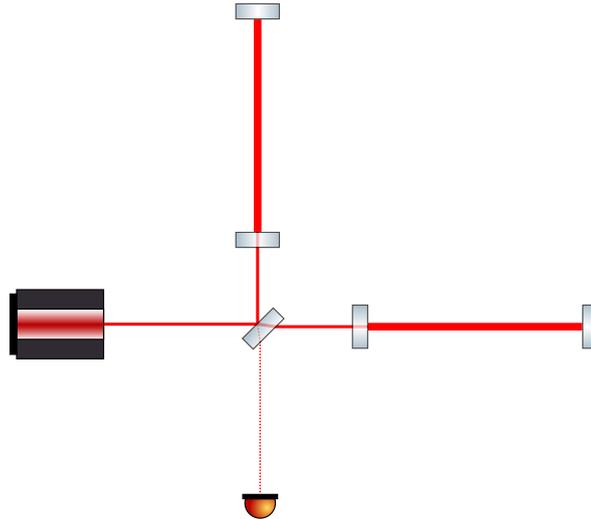


Figure 2.4: Fabry-Perot Michelson interferometer topology.

for which low noise lasers exist, and a gravitational wave strain $h_0 = 10^{-21}$ similar to that of GW150914, the modulation index will be of the order $\frac{\omega_0}{\omega_g} \approx 10^{12}$ and so the phase of the light will have to be measured at the output port to a precision of around $h_0 \frac{\omega_0}{\omega_g} \approx 10^{-9}$ rad, a difficult feat.

When the interferometer is held close to the dark fringe the light from each arm containing common phase changes exits the beam splitter back towards the input, and so this light would otherwise be lost.

Improvements to the Michelson interferometer design have been made over the past decades in order to address these issues, and these are discussed in the following subsections.

2.3.1 Fabry-Perot arm cavities

One way to simultaneously reduce the phase measurement requirement and the *effective* arm length is to use Fabry-Perot cavities. Fabry-Perot cavities increase a photon's path length by reflecting it many times between two partially transmissive mirrors. By placing Fabry-Perot cavities in the arms of a Michelson interferometer, as shown in figure 2.4, the response can be enhanced in a particular frequency band defined by the cavity parameters. The effect of the cavity on the sensitivity can be characterised by the *finesse* as discussed in appendix A.2.1. Increased cavity finesse leads to a greater number of stored photons, allowing for greater response to incident gravitational waves, but over a narrower bandwidth than the simple Michelson interferometer. The reflectivity of the mirrors in Fabry-Perot cavities must be chosen to allow for sufficient sensitivity in the desired band; the objective is not simply to maximise the light storage time.

The arm cavities within a Fabry-Perot Michelson interferometer must be held at the op-

erating point just as with a Michelson interferometer to maintain maximum sensitivity. Angular misalignments allow higher order modes of the light field to resonate which can complicate the longitudinal control of the interferometer and can introduce additional noise coupling from mirror surface defects.

2.3.2 Power recycling

As shown in section 2.1.2 gravitational wave interferometers are typically held close to the dark fringe where the carrier light is rejected by the beam splitter back towards the input laser. It is typical to place a Faraday isolator in the input path to prevent the interferometer from sampling the positions of the input optics used to steer the laser light towards the beam splitter, and so this light is dumped. Once lost this light is not available to sample the positions of the test masses.

To compensate for interferometer light lost towards the input port it is possible to increase the input laser power, but in general appropriate input lasers are already used at the maximum output power that satisfies an experiment's laser noise requirement, and this doesn't solve the underlying loss mechanism: some light will still be dumped by the Faraday isolator. Another approach is to place a *power recycling mirror* at the input to the interferometer which reflects the returning light back into the interferometer by forming a cavity between the recycling mirror and the arms, effectively increasing the power stored there. Used in combination with Fabry-Perot arm cavities this technique can achieve enhanced sensitivity over the standard Michelson interferometer. The power recycling mirror can be calibrated to enhance the carrier power in a band wider than the intended detector bandwidth, and the Fabry-Perot mirrors can be calibrated to set the detector bandwidth. The first generation Initial LIGO and Initial Virgo detectors were power-recycled Fabry-Perot Michelson interferometers.

2.3.3 Signal recycling

Signal recycling is a similar concept to power recycling, whereby an additional mirror is placed within the interferometer to selectively enhance light in a particular frequency band [51]. In this case the *signal recycling mirror* is placed at the output port of the beam splitter to create an additional cavity between the output and the arms. This mirror enhances the signal power at the expense of the bandwidth of the arms, as opposed to the carrier enhanced by the use of a power recycling mirror. The signal recycling mirror's transmissivity can be set to determine the frequency range over which this enhancement occurs, and the position of the signal recycling mirror can be tuned to focus this enhancement in either a narrow or broad frequency band [52]. This is discussed in more detail in

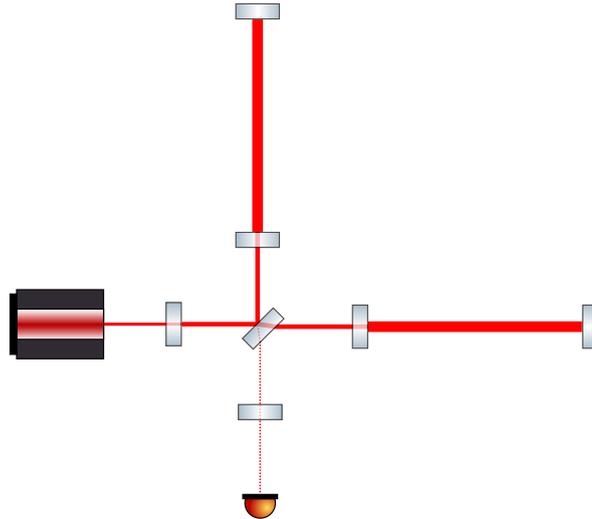


Figure 2.5: Dual-recycled Fabry-Perot Michelson interferometer.

section 7.2.

2.3.4 Dual recycling with Fabry-Perot arm cavities

The natural combination of power and signal recycling with the Fabry-Perot Michelson interferometer leads to the *dual-recycled Fabry-Perot Michelson interferometer* shown in figure 2.5. This is the topology that provides the greatest sensitivity in a given band of interest, either broadband or narrowband depending on the tuning of the signal recycling cavity, for a given laser power, arm length and mirror mass; it is therefore the topology employed in current generation detectors. The use of *dual recycling* was initially demonstrated in both table-top and suspended prototype experiments [53–55], and later a full-scale dual-recycled Michelson interferometer detector was demonstrated at GEO 600 [56, 57]. The Advanced LIGO interferometers were the first to fully implement the dual-recycled Fabry-Perot Michelson interferometer topology in detectors capable of sensing gravitational waves.

2.4 Surpassing the Standard Quantum Limit

Predictions for the population of sources within the range of the advanced detectors show that it is beneficial to improve the sensitivity at low frequencies [58]. The sensitivity of a Michelson interferometer at low frequencies can be increased through the use of heavier masses as shown by equation (2.22), scaling proportionally to \sqrt{m} . The use of mirrors larger and heavier than the 40 kg mirrors used in Advanced LIGO is a considerable technical challenge. The availability of test mass material of suitable quality in such dimensions is

not clear, as is the ability for the suspension systems to isolate noise from such large masses.

To improve sensitivity at higher frequencies, equation (2.17) shows that laser power can be increased. As with heavier mirrors, this presents technical challenges in laser stability [59], the control of *parametric instabilities* [60] and the thermal effects associated with absorption in materials [61].

To bypass the problems associated with the use of heavier mirrors and more powerful lasers, a number of techniques have been proposed in the literature to increase the sensitivity of interferometers beyond the SQL through the use of *quantum non-demolition* (QND) [62]. These include the modification of the optics of the interferometer [43], such as through the injection of squeezed light [44], variational readout [63, 64] or speed meters [65]; and the creation of new light-mirror interactions to increase the response of the interferometer to differential motion of the test masses [66].

2.4.1 Squeezing

The use of squeezing is an attempt to reduce the quantum noise at the output of the interferometer by injecting vacuum light with *correlated* noise. By choosing a suitable injection *quadrature*, it is possible to remove some of the quantum noise impinging upon the signal at a frequency of interest, instead moving the noise terms into the orthogonal, unobserved readout quadrature. Squeezing is particularly favourable in combination with dc readout, a combination currently implemented in GEO 600 [67, 68].

To reduce the effect of frequency-dependent *ponderomotive* squeezing arising from the mechanical susceptibility of the test masses, and to achieve broadband reduction of quantum noise, it is necessary to inject the squeezed light via filter cavities to provide a frequency-dependent phase shift to the vacuum field [43]. These cavities are typically high finesse, which makes the squeezed light particularly susceptible to filter cavity loss [69].

Squeezing has been demonstrated in GEO 600 with high duty cycle [70] and is a planned upgrade for Advanced LIGO in the near future [71]. The proposed designs for the Einstein Telescope and LIGO Cosmic Explorer assume 10 dB effective squeezing.

2.4.2 Variational readout

Instead of modifying the input noise at the output port of the interferometer with frequency-dependent squeezing, variational readout achieves sub-SQL sensitivity through the use of a homodyne detector with a homodyne angle chosen to create coherent cancellation of quantum noise between the local oscillator and the signal. Frequency dependent varia-

tional readout can be achieved in a similar way to squeezing: the output light can be passed through filter cavities in the same way as squeezed input.

Variational readout can in theory be combined with squeezing either with a fixed squeezing angle [72] or through a complicated frequency dependence of both squeezing and homodyne phase filter cavities [73]. The use of homodyne readout in gravitational wave detectors, however, is not considered mature enough for upgrades to existing or future facilities, primarily due to the stability requirements for the local oscillator field [74].

2.4.3 Light-mirror interactions

Optical springs [75–79], *optical inertia* [80, 81] and *intracavity schemes* [82–84] have been proposed for use in gravitational wave detectors to improve sensitivity beyond the SQL through modification of the mechanical response of the interferometer’s mirrors with light using optomechanical interaction.

The creation of optical springs requires complicated control arrangements. The use of two optical springs to remove the instabilities created by a single spring can relax some of the control requirements [78] but full studies of the effect of noise and the sensitivity on this type of interferometer are not at a stage to be able to predict their use in future detectors.

2.4.4 Modification of the interferometer design

First proposed in 1990 [65], the measurement of the speed of test masses instead of displacement can lead to a reduction in quantum noise. Proposals for experiments to measure speed were made later and involved the use of an additional optical cavity at the output port of a Michelson interferometer [85, 86], termed a *sloshing* cavity, creating an interaction between the main interferometer and the sloshing cavity that samples the test mass coordinates in a way that resembles speed.

In 2003, Chen showed that the Sagnac interferometer contained the necessary characteristics of a speed meter [87] and estimated the sensitivity that such an interferometer might achieve when the corner mirrors are replaced with arm cavities to resemble a Fabry-Perot Michelson interferometer. This work was later expanded to include the effect of losses [88, 89].

The Sagnac speed meter is being considered as a potential upgrade for the Einstein Telescope beyond its initial configuration [90, 91], albeit following a polarising topology with linear arm cavities [88] due to the sensitivity degradation from back-scattering in triangular arm cavities [92].

2.5 The future of ground-based gravitational wave interferometry

Plans are in place for upgrades to Advanced LIGO after the science run in 2016, when squeezed light injection will be implemented. In the medium term, the interferometer may be adapted to run with cryogenic optics to provide sensitivity at lower frequencies. The Advanced LIGO and Advanced Virgo detectors already push their current facilities to their limits, however, and so in the long term the goal is to build new facilities with significantly improved sensitivity. A conceptual design study for the Einstein Telescope was completed in 2011 [16], a new European facility in a triangular, 10 km arm configuration, and similar studies for a new 40 km LIGO facility are ongoing [17, 18]. These facilities are planned for the late-2020s to early-2030s, and the ongoing research and development work will help to determine the technologies that become part of these detectors. The reduction of quantum and thermal noise and the control of such interferometers will be crucial areas of investigation, and some potential solutions are presented in the rest of this work.

Chapter 3

Measurement of transverse-to-longitudinal phase coupling in a waveguide mirror

The following chapter has been adapted from Upper limit to the transverse to longitudinal motion coupling of a waveguide mirror [35], published in Classical and Quantum Gravity in 2015. The article was entirely written by the author and is suitable for inclusion, expanded as appropriate, within this thesis. The results and conclusions presented are identical.

3.1 Thermal noise in advanced detectors

At their most sensitive frequencies, the second generation detectors are expected to be limited by Brownian thermal noise arising from the reflective coatings on the detectors' test masses [28, 93–95]. In order to help mitigate this limitation beyond the next generation of detectors, efforts are under way to develop mirror coatings with lower thermal noise [96, 97].

In the case of Advanced LIGO, each end test mass (ETM) consists of a substrate with 19 pairs of sub-wavelength coatings which produce a transmission of 5 ppm for 1064 nm light with very little loss [98], with each coating layer contributing to the overall thermal noise [28, 95]. The approach taken by Levin to calculate the thermal noise of mirrors [93] shows that mechanical loss at the front surface of a mirror contributes more to the Brownian noise level than loss from an equivalent volume in the substrate. Additionally, typical coating materials tend to exhibit mechanical loss orders of magnitude higher than typical substrate materials [28, 95]. For these reasons particular attention is being given to the reduction of coating thermal noise to improve the sensitivity of future detectors.

One strategy, to be applied for example in KAGRA [99], is to cool the mirrors to cryogenic temperatures. While this can potentially reduce the thermal noise of the mirrors [100], the utilisation of cryogenic mirrors requires new infrastructure, different choices of mirror substrate and coating materials and poses the challenge of heat extraction from the mirror without spoiling its seismic isolation and thermal noise performance. Efforts in the application of cryogenics are also under way to identify suitable substrate and coating materials for ET-LF, the low frequency interferometer as part of the proposed Einstein Telescope [47, 101–103], and the proposal for a cryogenic upgrade to LIGO, Voyager [104].

3.2 Waveguide and grating mirrors

Apart from using different coating materials [33, 105] or different beam shapes [106–108] such as with LG33 modes [109], another potential approach is to utilise waveguide mirrors (WGMs) [110–113]. These mirrors can possess high reflectivity at a wavelength determined by their structure. In contrast to conventional dielectric mirrors, mirrors possessing waveguide coatings can exhibit high reflectivity without requiring multiple stacks [114]. A waveguide coating instead presents incident light with a periodic grating structure of high refractive index material n_H on top of a substrate with low refractive index n_L (see figure 3.1). Light is forced into a single reflective diffraction order, the 0th. In transmission, only the 0th and 1st diffraction orders are allowed as long as the condition in equation (3.1) for the grating period, p , and the light’s wavelength in vacuum, λ , is fulfilled [110]. The light diffracted into the 1st order undergoes total internal reflection at the substrate boundary where it excites resonant waveguide modes. With suitable parameters the light at the waveguide boundary to the incident beam will contain a 180° phase shift with respect to the 0th order transmitted light, causing destructive interference such that most of the incident light is reflected [115].

$$\frac{\lambda}{n_H} < p < \frac{\lambda}{n_L} \quad (3.1)$$

We have shown in ref. [117] that a suitably optimised WGM can provide a reduction in coating thermal noise amplitude of a factor of 10 at cryogenic temperature compared to the ETM employed in Advanced LIGO. Figure 3.2 shows Brownian thermal noise modelled for different numbers of bilayers (following ref. [28]) for the Advanced LIGO ETM alongside the grating thermal noise result from Heinert et al.. For each additional bilayer in the dielectric stack, the mirror’s transmissivity decreases but its Brownian thermal noise increases. The grating, however, requires only a change in grating parameters to produce a specific transmissivity, and this does not have a strong dependence on Brownian thermal noise.

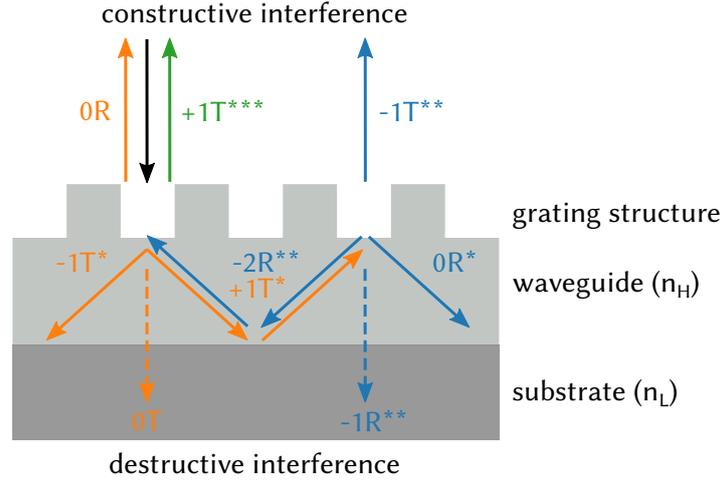


Figure 3.1: Propagation of light within a waveguide mirror (a reproduction of [116, figure 1]). The grating and waveguide layers have refractive index n_H , and sit atop a substrate of refractive index n_L . The incident light (black) diffracts into the first orders ± 1 (orange). At the substrate the light reflects and at the grating the light either interferes with the incident light to produce reflection or reflects back into the waveguide (blue). The numbers next to the arrows represent the diffraction order, T represents transmission through the grating, R represents reflection from the grating and each asterisk represents a single diffraction. The grating parameters can be tuned such that the phase difference at the front (waveguide) surface produces high reflection. In realisations of waveguide mirrors such as this, a thin etch-stop layer is placed between the grating and waveguide layers to assist fabrication [113].

3.2.1 Transverse to longitudinal coupling in grating mirrors

Previous efforts to demonstrate grating structures as alternatives to dielectric mirrors have identified phase noise in the light reflected from the grating not otherwise present in dielectric mirrors [118, 119]. This effect arises from transverse motion of grating mirrors with respect to the incident light. Incident light at angle α is reflected into the m^{th} diffraction order, exiting at angle β_m (see figure 3.3). The change in path length δl_L between the reflected and incident light is then

$$\delta l_L = \zeta_a + \zeta_b = \delta y (\sin \alpha + \sin \beta_m), \quad (3.2)$$

where ζ_a and ζ_b represent the relative optical path length of each depicted ray. The phase modulation induced in the reflected light by periodic, transverse motion of the WGM is proportional to the period with a 90° phase lead over the transverse motion [120]. The frequency noise added to the reflected light can be enough to mitigate the improvement in coating thermal noise, as witnessed in a study of 2nd order Littrow gratings [120], where the level of coupling between longitudinal and transverse motion was found to be 1:100. Although WGMs also possess gratings, the resonant waveguide structure can be tuned to remove this effect as shown in simulations by Brown *et al.* [116].

There are two mechanisms by which grating mirrors can couple transverse motion into

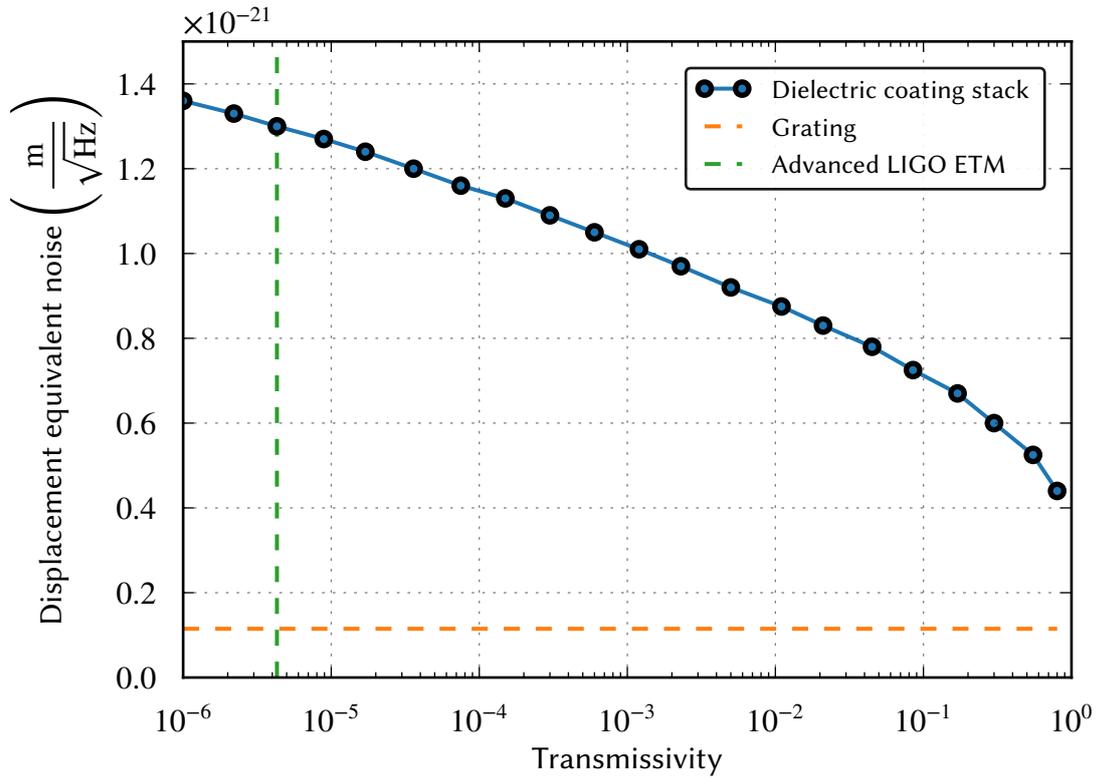


Figure 3.2: Reproduction of results from Heinert et al. [117] showing Brownian thermal noise in an Advanced LIGO style ETM at room temperature versus that of a WGM at cryogenic temperature, as a function of transmissivity. The markers in the coating curve represent the number of quarter-wavelength bilayers forming the dielectric stack. Each additional bilayer in the coating stack produces lower overall transmissivity, but also increases Brownian thermal noise. The grating's Brownian thermal noise contribution is independent of transmissivity. The transmissivity of Advanced LIGO's ETM is shown as a vertical, dashed line.

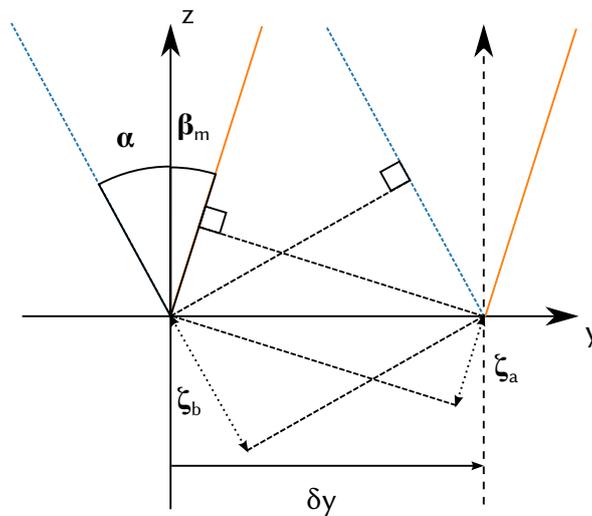


Figure 3.3: Optical path length changes ζ_a and ζ_b due to transverse motion of a Littrow grating, based on ref. [119, figure 4]. Incident light diffracted into a different order undergoes a path length change $\delta l_L = \zeta_a + \zeta_b$.

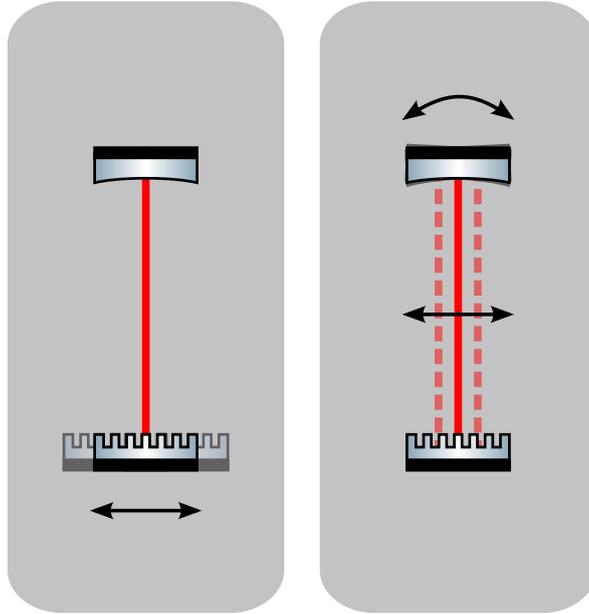


Figure 3.4: Two ways in which light can be scanned across the surface of the WGM. The left panel shows the effect of WGM motion with respect to a static beam, while the right panel shows the effect of light beam motion (due to rotation of the cavity mirror opposite the WGM) with respect to a static WGM. In the latter picture, the cavity’s eigenmode is translated transversely to the optical axis as one mirror is rotated.

longitudinal phase changes (see figure 3.4). The first is through transverse motion of the grating, which can in principle be minimised with appropriate suspension design. The second mechanism is the coupling of changes in the opposite cavity mirror’s alignment into the spot position on the grating mirror. This effect is of particular importance to gravitational wave interferometers, where longer arm lengths can increase its detrimental impact. The latter mechanism is the primary focus of this work.

In order to quantify its transverse coupling, a WGM was produced in collaboration with Friedrich-Schiller University Jena, Germany; see table 3.1 for its properties. It was designed for light of wavelength 1064 nm, and consists of an etched grating structure on top of a waveguide layer, both tantala (Ta_2O_5), on a silica substrate. This chapter details an experiment carried out to measure its transverse coupling level.

3.3 Experiment

The fabricated WGM was used as the input coupler for a Fabry-Perot cavity, held on resonance using the Pound-Drever-Hall (PDH) technique [121]. The error signal provided by the PDH technique represents changes in cavity length, and this can be fed back to the laser’s frequency via a frequency stabilisation servo. This technique is described in more detail in section 3.3.4.1.

Parameter	Value
Materials	SiO ₂ , Ta ₂ O ₅ , Al ₂ O ₃
Design λ	1064 nm
Grating depth	390 nm
Waveguide depth	80 nm
Etch stop depth	20 nm
Grating period	688 nm
Fill factor	0.38
Reflectivity	96%

Table 3.1: Design parameters of the WGM produced by Friedrich-Schiller Jena for the experiment to measure transverse to longitudinal coupling. It is similar to the one used in ref. [113], with increased reflective surface area.

3.3.1 Effect of waveguide mirror rotation on cavity length

A non-zero WGM transverse to longitudinal coupling, ω_1 , produces a phase shift on the reflected light. This manifests itself as an effective change in cavity length, δl_W , as the laser light is scanned across its grooves by the rotation of the ETM:

$$\delta l_W = \theta \kappa \omega_1, \quad (3.3)$$

where θ is the ETM's (small) rotation angle and κ is the cavity's coefficient of ETM rotation to transverse WGM spot displacement.

Additional cavity length changes from mirror rotation are created by two geometrical effects as shown in figure 3.5. The first effect, δl_s , is due to the position of the beam with respect to the centre of the mirror's surface. For a rotation θ , a beam offset from the centre of the mirror by a displacement y will receive a change in (longitudinal) path length of

$$\delta l_s = y \tan \theta \approx y\theta, \quad (3.4)$$

for small angles. The second effect, δl_d , is due to the depth d of the mirror, proportional to the rotation angle θ . The position of the centre of the mirror with respect to the zero rotation case, y_d , is then

$$y_d = \frac{d}{2} \tan \frac{\theta}{2} \approx \frac{d}{4} \theta, \quad (3.5)$$

and the change in path length this causes is

$$\delta l_d = y_d \tan \theta \approx \frac{d}{4} \theta^2. \quad (3.6)$$

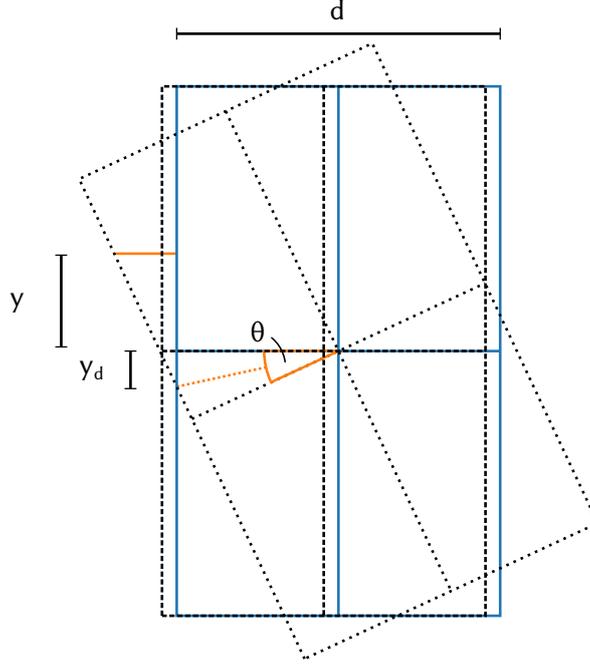


Figure 3.5: Geometrical ETM longitudinal coupling. For a given rotation θ and spot centre position offset y , the (longitudinal) position change in the surface of the mirror (shown in blue) as seen by the reflected light is approximately $y\theta + \frac{d}{4}\theta^2$. The straight, solid red line in the figure shows this longitudinal change.

The total longitudinal effect δl_E caused by the rotation of the ETM is therefore:

$$\delta l_E = \delta l_s + \delta l_d \approx y\theta + \frac{d}{4}\theta^2, \quad (3.7)$$

and the total length signal from all mirror effects will then be:

$$\delta l(\theta) = \delta l_W + \delta l_E \approx \theta\kappa\omega_1 + y\theta + \frac{d}{4}\theta^2. \quad (3.8)$$

3.3.2 Cavity length signals

Considering the ETM's dimensions and mass, it is possible to calculate the cavity length change due to the two geometrical effects shown in equation (3.7) for a given rotation. From the cavity length change it is then possible to infer the WGM's transverse to longitudinal coupling level using equation (3.8). The phase effect associated with transverse to longitudinal coupling is expected to be independent of the average spot position, whereas there is a phase change about the ETM's centre of rotation. It is expected therefore that a spot position will exist, for a non-zero WGM transverse coupling level, offset from the ETM's centre of rotation, for which there is a cavity error signal minimum. This effect arises as a result of the coherent cancellation of δl_W and δl_E (see figure 3.6). The spot position corresponding

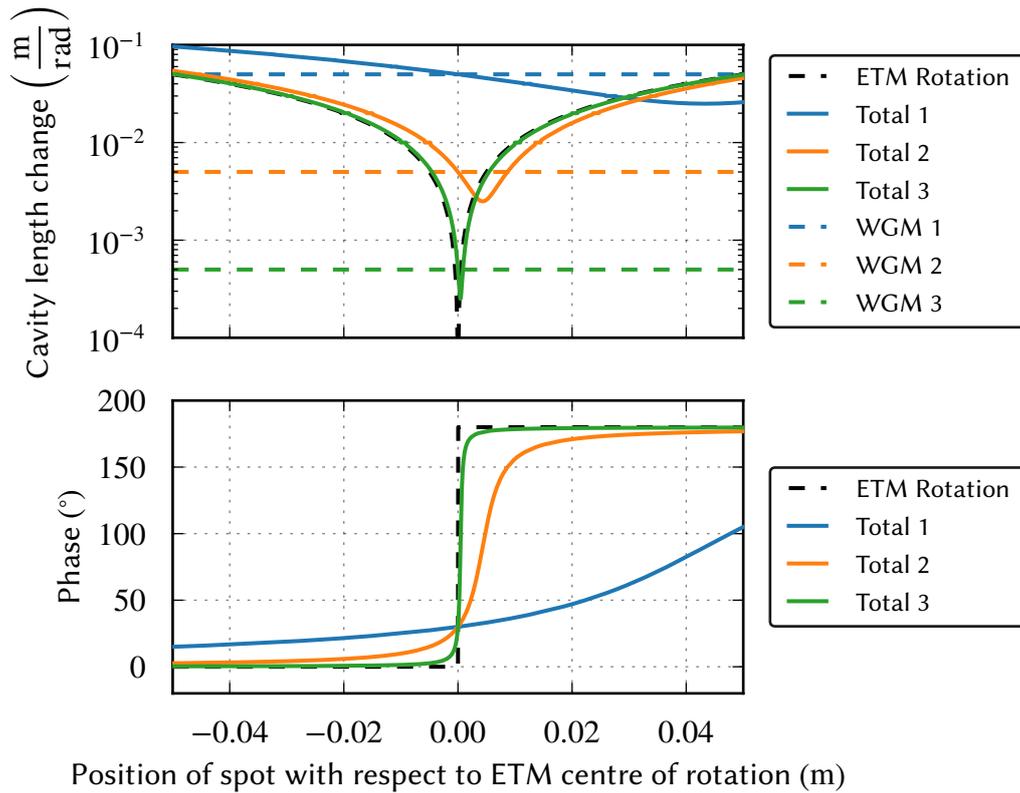


Figure 3.6: Simulations of indicative cavity longitudinal error signals during ETM rotation for different levels of WGM coupling. The signals are functions of the transverse position of the reflected light relative to the ETM’s centre of rotation, the angle of rotation, the mirror depth and the WGM’s coupling level. The rotation to longitudinal coupling of the ETM (black dashed line) combines with the transverse to longitudinal coupling of the WGM (green, orange and blue dashed lines) to produce cavity length changes (green, orange and blue solid lines). In this example configuration, the ETM rotation is 1×10^{-7} rad, the ETM’s depth is 0.1 m and the corresponding WGM coupling levels are 1:370 (blue), 1:3700 (orange) and 1:37000 (green).

to the error signal minimum allows the WGM’s transverse to longitudinal coupling level to be inferred.

Examples of WGM coupling levels yielding cavity length changes smaller than (green), larger than (blue) and roughly equivalent to (orange) the ETM’s effects are shown in figure 3.6. For cases where the WGM’s coupling level yields a significant cavity length change with respect to that of the ETM’s rotation, coherent cancellation creates a trough significantly offset from the ETM’s centre of rotation.

3.3.3 Experiment infrastructure

The Glasgow 10 m prototype facility provides a test bed in which the WGM’s transverse to longitudinal coupling can be quantified. The prototype is housed in a class 1000 clean room and consists of an input bench at atmospheric pressure and a vacuum envelope able

to reach pressures of order 10^{-3} Pa. The envelope consists of nine 1 m diameter steel tanks, each connected by steel tubes, arranged into two parallel arms of length 10 m, with a shorter arm for input optics situated between them.

In the experiment, 1064 nm laser light was passed through a single-mode fibre to provide spatial filtering, and an electro-optic modulator (EOM) to impose RF sidebands on the light required to produce an error signal with the PDH technique. The light was then coupled into the vacuum system via a periscope. A control system senses the motion of the cavity with the RF photodiode near tank 1 and provides corrective actuation on the laser crystal's piezoelectric transducer (PZT) and temperature via a frequency stabilisation servo and associated electronics (see section 3.3.4). This configuration is shown in figure 3.7.

Tanks 2 and 3 house a beam splitter and steering mirror, respectively, attached to double stage suspensions. In tanks 4 and 5 were sets of two triple suspension chains based on the GEO 600 design [122]. A viewport present to the rear of tank 5, and to the side of tank 1, allowed for light to exit the vacuum envelope for the purposes of sensing and control.

The WGM was attached to an aluminium block of mass 2.7 kg and suspended from tank 4's cascaded (triple) pendulum, forming the cavity's ITM. A silica test mass, also 2.7 kg, with a 40 ppm transmission coating, was used as the ETM, suspended from a similar triple pendulum in tank 5. On the rear surface of the ETM were three magnets for the purpose of actuation, the positions of which are shown in figure 3.8. With optimal alignment the mirrors formed an overcoupled cavity with finesse 155.

A three-stage reaction chain was placed behind the triple pendulum of the ETM to provide voice coil actuation upon the magnets on the ETM's rear surface. The upper and intermediate stages were identical to those of the chain carrying the ETM, however—for the purposes of another experiment, not reported here—the lower stage was split into two parts separately suspended from the intermediate stage. The part closer to the ETM was a 1.8 kg aluminium block that carried the voice coils. The other part was a 0.9 kg aluminium block required to balance the suspension.

3.3.4 Measuring cavity length changes

An RF photodetector was placed at the viewport on tank 1, where it could view the light reflected from the cavity. By using the PDH technique (see section 3.3.4.1), the signal from this photodetector provided an error signal for the cavity length. This signal was fed back to the laser via the frequency stabilisation servo to maintain cavity resonance. The servo's high frequency feedback signal—a voltage applied across the laser's PZT—provided a means of calibrating cavity length changes at frequencies greater than 12 Hz. Using the PZT's actuator coefficient, $1.35 \text{ MHz}/V_{\text{rms}}$, the cavity length change δl per error signal volt could

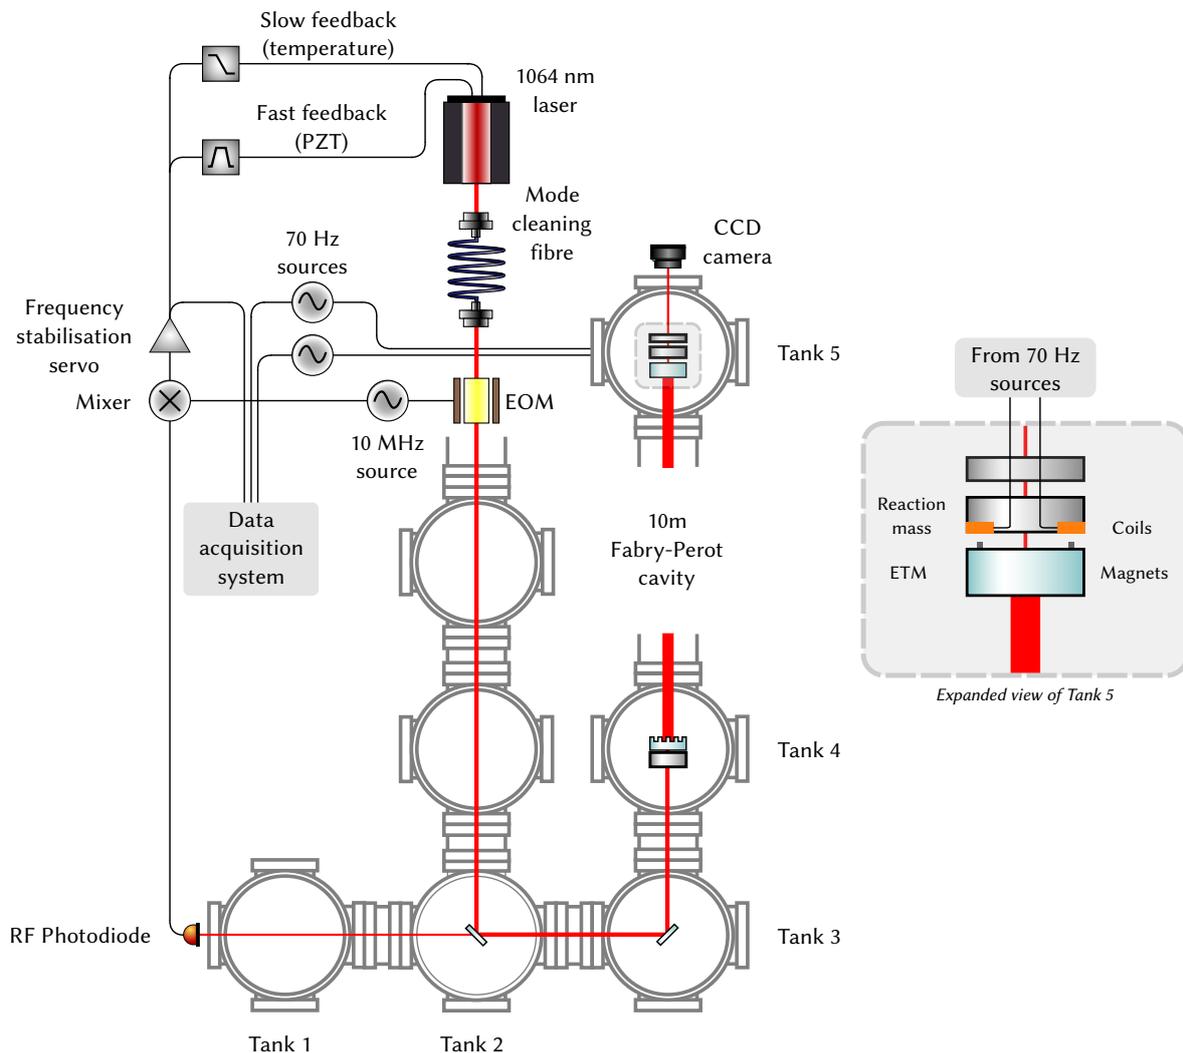


Figure 3.7: The experimental setup in the prototype facility. The laser light is passed through input optics (not shown), a mode cleaning fibre and an EOM before being coupled into the vacuum system via a periscope. It then travels to tank 2 where it is reflected off a beam splitter and directed into one of the arms of the prototype by a steering mirror in tank 3. The two mirrors in tanks 4 and 5 form a Fabry-Perot cavity. The cavity mirrors are suspended from triple stage suspensions, and the beam splitter and steering mirror are both suspended from double suspensions.

The ETM is rotated in yaw using the 70 Hz source. It is fed to a coil driver where it is coupled into tank 5 via a vacuum feedthrough. Coil formers on the front edges of the reaction mass contain wound copper wire connected to the vacuum feedthrough. Magnets are attached to the back of the ETM. The reaction mass is behind the ETM, containing a hole in its centre to allow light to exit the vacuum tank where it can be viewed with the CCD camera. A larger version of the contents of tank 5 can be viewed in the panel to the right of the figure.

The cavity is held on resonance by the frequency stabilisation servo. This feeds back to the light's frequency via the laser crystal's temperature below 12 Hz and its PZT above 12 Hz up to a unity gain frequency of 16 kHz (see section 3.3.4.2).

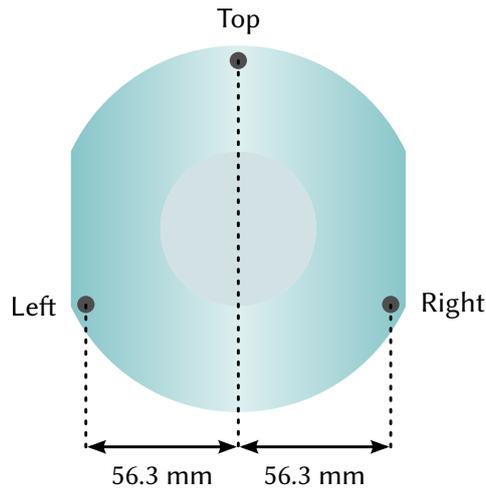


Figure 3.8: The positions of the magnets on the rear surface of the ETM. The designations used in this article are shown next to each magnet. The top magnet is positioned at the centre of yaw, near the top of the mass. The left and right magnets are positioned 56.3 mm either side of the centre of yaw. Coils on the ETM’s reaction mass (not shown) are positioned coaxially behind each magnet.

be calculated to be $133 \text{ nm}/V_{\text{peak}}$.

3.3.4.1 The Pound-Drever-Hall technique

Several techniques exist to control cavity length fluctuations, the earliest of which was through the use of *dc locking*. This involves keeping the power on a photodetector constant by feeding back an error signal to either the laser’s frequency or to actuators on the cavity mirrors. The error signal arises from a change in measured photodetector power from some set point. This set point is in the case of dc locking necessarily away from the dark fringe, because at that point a cavity length increase or decrease results in an identical drop in power incident upon the photodetector and thus cannot be used as a discriminant for cavity length. Instead, the set point must be somewhere on the slope leading to or from the peak, where a length increase has opposite sign compared to a length decrease. As the set point is not situated at the operating point producing maximum cavity power build-up, and thus sensitivity, this technique has since been superseded by a range of techniques that result in superior performance. Dither locking [123], which uses a dither signal applied to a mirror to act as a local oscillator to the carrier from which an error signal can be derived, was used for the output mode cleaner of Enhanced LIGO [124]. Another technique, tilt locking [125], uses the beat between the first and fundamental spacial modes of the carrier. We will however focus on a fourth technique, Pound-Drever-Hall (PDH) [121, 126], which is particularly suited for Fabry-Perot experiments requiring good sensitivity across a wide bandwidth.

As described in chapter 2, phase modulation cannot be detected by standard photodetector electronics. Although modulation due to the presence of phase sidebands leads to a change

Parameter	Description
Cavity input power	Approx. 150 mW
ETM transmissivity	40 ppm
ETM radius of curvature	15 m
ETM spot size	2.138 mm
ITM transmissivity	4 %
ITM radius of curvature	∞
ITM spot size	1.554 mm
Cavity length	9.81 m
Cavity finesse	155
Cavity g-factor	0.347
Beam waist size	1.554 mm
Beam waist position	At ITM
Sideband frequency	10 MHz

Table 3.2: Cavity parameters.

in light power, this change happens on time scales too short for a photodiode to register a change in its electronic signal: any stray capacitance in its material or transmission lines filters the phase modulation in the same way as a low pass filter, averaging the modulation to zero. Figure A.2, however, shows that the phase of the error signal from a cavity detuned from resonance appears to be a good error signal: it is bipolar about the optimal operating point. The PDH technique provides access to this phase information.

The key features of PDH are highlighted below. More detailed descriptions can be found in, for example, refs. [23] and [126]. Mirror motion imparts phase modulation upon the carrier as shown in appendix A.3. Restating equation (A.21) we see that phase modulation upon the carrier produces upper and lower sidebands with frequencies $\omega_0 \pm \omega$, where ω represents the phase modulation frequency:

$$E = E_0 e^{i\omega_0 t} \left(1 - \frac{m^2}{4} + i \frac{m}{2} (e^{-i\omega t} + e^{i\omega t}) \right), \quad (3.9)$$

assuming sub-wavelength motion.

Phase modulation can also be intentionally imparted upon the carrier through the use of an EOM, as depicted in figure 3.7. With the PDH technique, the EOM is placed in the path of the cavity's input light where it imparts strong phase modulation to the carrier at RF frequencies. The choice of this frequency band is motivated by the availability of low cost and low noise electronics, the lack of 1064 nm laser frequency noise, and the avoidance of the audio band where experiments in the field of ground-based gravitational wave interferometry typically desire high sensitivity. The RF sidebands produced by the EOM must be chosen to be greater than the cavity's FWHM (see appendix A.2.1) to prevent them from entering the cavity. As the carrier light resonates within the cavity and reflects back to-

wards the laser, any phase modulation imparted to the carrier by mirror motion beats with the RF sidebands that do not enter the cavity, with the difference in phase showing up as signal sidebands upon the RF sidebands. The signal sidebands can be recovered from the field through demodulation at the RF frequency.

In a typical PDH setup, a frequency generator is fed both to the EOM and to a mixer connected to the output of an RF photodetector placed in reflection of the cavity (see figure 3.7). This ensures that the same frequency used to create the RF sidebands is used to demodulate the superposition of fields reflecting from the cavity. Mixing the oscillator's signal with the reflected light is equivalent to multiplying the reflected field by a factor of $\sin \omega t$ or $\cos \omega t$, which yields a signal with frequency components proportional to the cavity motion at those frequencies. This signal is bipolar, with cavity mirror motion in one direction yielding a different sign to motion in the other direction. This *error* signal can be fed back to the cavity's actuators to hold the cavity resonant.

3.3.4.2 Cavity control

The reflected light from the cavity was sensed with the RF photodiode placed near tank 1. This was mixed in order to demodulate the field and recover the cavity error signal, and this was coupled into an analogue servo containing filters designed to keep the Fabry-Perot cavity on resonance. This analogue servo design resembles that of the servo presented in ref. [127], and contains feedback paths to the laser's temperature and PZT, able to correct the laser's frequency at low and high frequencies, respectively, thus providing a way to maintain the cavity resonance condition. Estimated open loop gains of both the temperature and PZT feedback are shown in figure 3.9. The PZT feedback is flat below 30 Hz, where the temperature feedback servo performs the majority of the actuation. Above this point, a slope proportional to $\frac{1}{f^3}$ removes feedback at higher frequencies to provide greater control bandwidth for corrections at low frequencies where it is most needed due to seismic noise (see appendix B.4.2). A differentiator is present above 3 kHz to correct the phase of the feedback signal to allow it to cross below the unity gain point without creating unstable signals with equal magnitude but opposite sign to the measured motion, as discussed in appendix B.4.3. A copy of the feedback to the PZT is sent to CDS where a series of filters produce the temperature feedback signal. Gain is applied to the low frequency components to ensure the laser temperature's signal is strong where its response is strong. Additional resonant gain is applied at the suspension resonance frequency of 0.6 Hz to prevent it from ringing. One further filter is used to provide a stable crossover with the PZT feedback at around 12 Hz. Above 12 Hz the response of the temperature feedback is greatly suppressed due to material's time constant. The unity gain point of the combined servo is around 16 kHz.

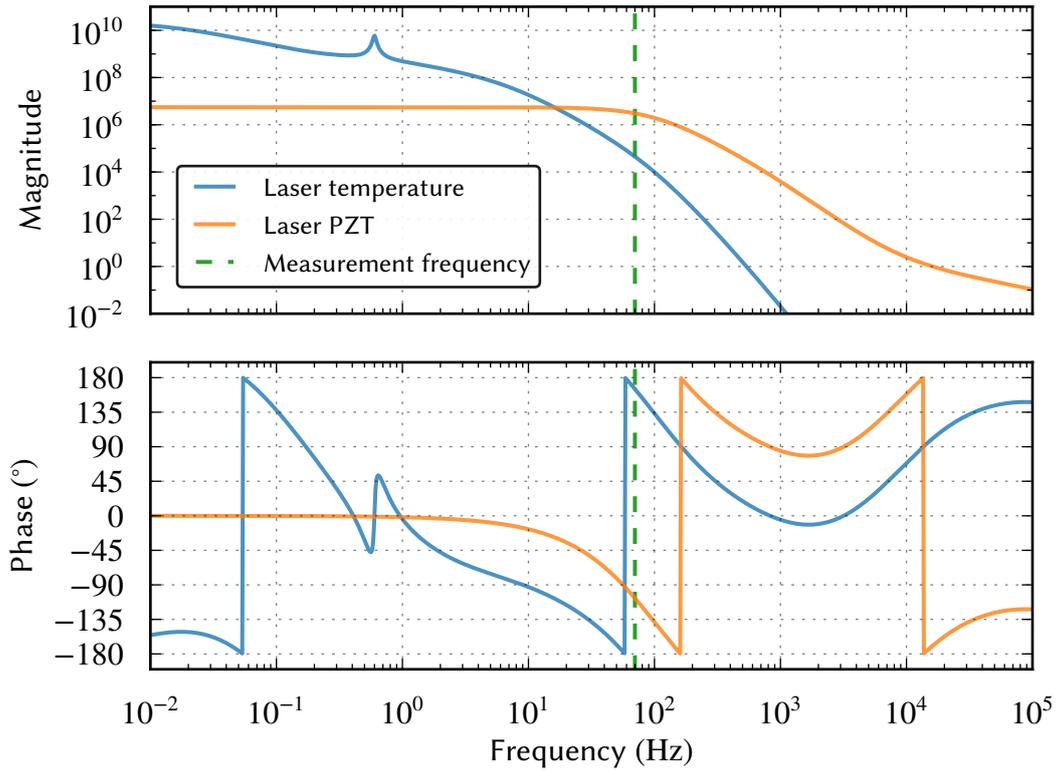


Figure 3.9: Estimated open loop gain of the analogue servo used to feed the cavity error signal back to the laser’s temperature and PZT. The laser temperature actuator provides the majority of the feedback below 12 Hz, where it has a strong response. Its effect is significantly reduced above this point, where the time constant of the material becomes significant. The PZT actuator provides actuation at higher frequencies, where it can provide smaller but faster corrections up to many tens of kHz. Various filters are required to produce stable temperature-to-PZT crossover and unity gain points. The frequency used for measurements discussed later is shown as a dashed line.

3.4 Measurements and analysis

From the orientation of the WGM’s gratings, it was expected that actuation of the ETM in yaw, which would scan the cavity light across the WGM’s surface transverse to the direction of its grooves, would exhibit WGM transverse to longitudinal coupling if present.

3.4.1 Voice coil balancing

For the purposes of actuation upon the ETM, two sinusoidal signals V_L and V_R (corresponding to the left and right voice coils on the ETM’s reaction mass, respectively) were produced using separate, phase locked signal generators. A signal frequency of 70 Hz was chosen so as to be above the suspensions’ pole frequencies, where the mirror’s dynamics approximate that of a free mass, but low enough to provide an adequate rms test mass rotation. The signals V_L and V_R , with suitable balancing (see below), could then be actuated in- or

out-of-phase to produce longitudinal or yaw actuation upon the ETM, respectively.

When V_L and V_R were equal in magnitude but out-of-phase, the ETM's movement contained a linear combination of rotational and longitudinal components due to force imbalances between the voice coils. To ensure that actuation upon the ETM contained only a yaw component, the cavity's longitudinal error signal was minimised during out-of-phase actuation by changing the gain of V_L . This balanced the magnitude of the torque applied by each actuator to the left and right sides of the ETM. Any WGM transverse to longitudinal coupling present would act with phase orthogonal to this voice coil actuation and would thus be unchanged by the torque balancing.

Pitch actuation upon the ETM, which would scan the cavity light in a direction parallel to the WGM's grooves, was not expected to contribute to the cavity's error signal via the WGM's coupling. However, unintended pitch actuation upon the ETM would couple into the cavity's length via the same geometrical mechanism as yaw shown in equation (3.7). To minimise the ETM's pitch component during actuation in yaw, the cavity's error signal was minimised by applying an offset voltage to the top coil. In practice, minimal pitch coupling was achieved when the offset signal was zero.

3.4.2 Actuator calibration

To calibrate the cavity's longitudinal response to voice coil actuation, the voice coils were actuated with the balanced V_L and V_R signals in-phase for a period of 120 s. The total force applied to the ETM by the voice coils is related to the applied voice coil signal frequency f , the ETM's mass m and the cavity length change δl as defined in equation (3.8):

$$F = 4\pi^2 f^2 m \delta l. \quad (3.10)$$

3.4.3 Measurement of transverse to longitudinal coupling

Four spot positions corresponding to y in equation (3.4) were chosen across the surface of the ETM. The input beam was aligned to the cavity axis corresponding to each spot position using the beam splitter and steering mirror nearest to the ITM, and the cavity mirrors were aligned to create a fundamental mode resonance. The voice coil signals V_L and V_R were set out-of-phase to produce motion on the ETM in yaw. The magnitudes of V_L and V_R were not altered between the longitudinal calibration and this yaw actuation, so it was expected that the previously outlined minimisation of yaw to tilt actuation would also result in minimal longitudinal to tilt actuation. The cavity length signal was recorded for a period of 300 s.

For each nominal spot position an additional measurement was taken with V_L set to ± 0.1 V from its balanced setting for a period of 60 s. This allowed two additional data points to be

obtained for each spot position. By calculating the gradient (cavity length change per spot position with respect to the centre of yaw) of the central and inner-left spot positions, it was possible to assign an effective spot position for each of the offset points.

The spot positions used to obtain cavity error signals are shown with respect to the centre of the ETM's reflective surface in table 3.3. The spot positions were subject to two sources of error: the measurement of the spot positions with respect to the centre, and the error in the ETM's centre of rotation due to misalignment between the voice coils and their corresponding magnets. The spot position error was assumed to be ± 1 mm from visual inspection of the suspensions, measured with the CCD camera placed in transmission of the ETM, using the known width of the ETM's reaction mass as a calibration.

3.4.3.1 Background noise subtraction

To remove the noise floor from the data, measured cavity length signals from frequencies f in the ranges $65 \text{ Hz} \leq f \leq 68 \text{ Hz}$ and $72 \text{ Hz} \leq f \leq 75 \text{ Hz}$ were averaged and the resultant figure subtracted in quadrature from the peak signal at $f = 70 \text{ Hz}$. The 2 Hz gap between the measurement frequency and the bound of each noise floor estimate was used due to the finite resolution of the length signal peaks.

3.4.3.2 Effect of voice coil misalignment

Misalignment between the voice coils and magnets can lead to unintended torque and longitudinal actuation, confusing the calibration. Literature appears to be sparse on this matter, and preliminary analysis involving finite element simulation packages were unfruitful. Instead, to evaluate whether this effect could be significant, a small experiment was configured as shown in figure 3.10. A rod was placed above a magnet with the voice coil attached to its end, with both magnet and voice coil having the same dimensions as those of the main experiment. The magnet was glued to a thick perspex disc attached to the base edge of an upturned plastic cup to allow the force applied to the magnet to rigidly couple to the base of the cup. The cup was itself placed upon scales with $1 \mu\text{g}$ precision and a translation stage with $25 \mu\text{m}$ precision.

With the front edge of the voice coil separated from the base of the magnet by 7.9 mm—close to that of the main experiment—a series of force measurements were taken. A constant current source of 50 mA was applied through the coil whilst incrementing the translation stage in steps of 0.1 mm. Figure 3.11 shows the results from this experiment. Individual data points are dominated by hysteresis of the measurement apparatus but taken as a collection of points appear to follow a quadratic scaling law. With a simple 2nd order quadratic regression, the results show that the force drop due to voice coil and magnet misalignment has an

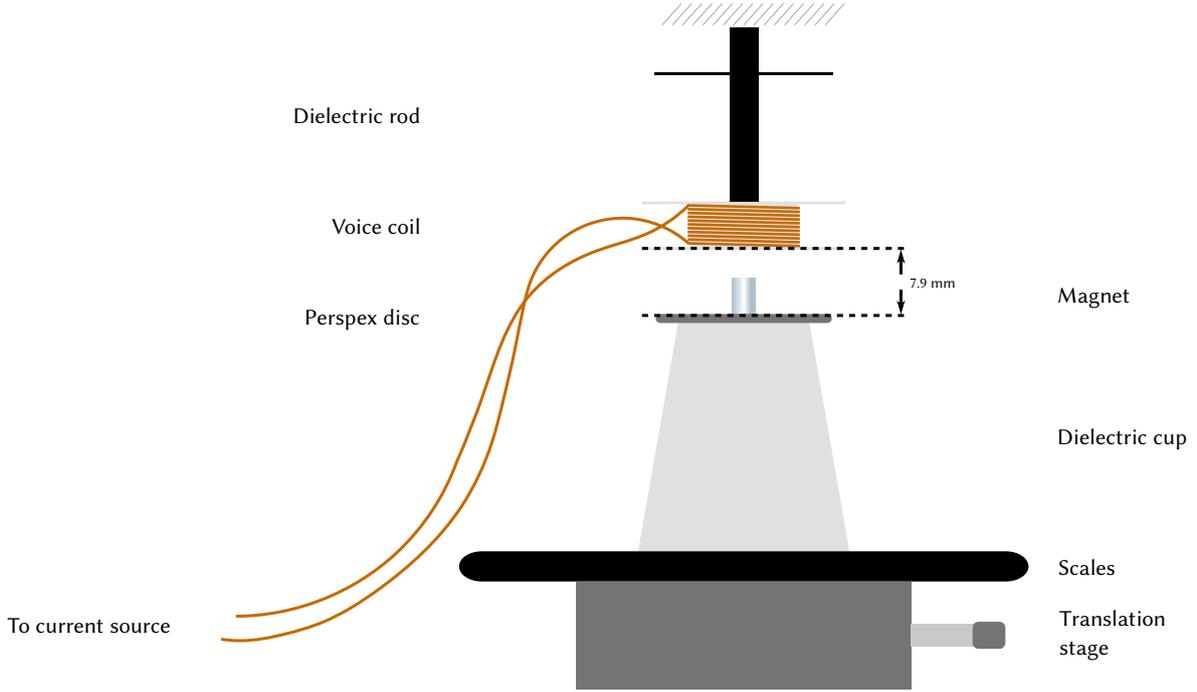


Figure 3.10: Experiment to measure the effect of misaligned voice coil actuation.

Spot position [mm]		
-0.1 V	0.0 V	0.1 V
-12.9	-12.5	-12.1
-5.4	-5.0	-4.6
-0.4	0.0	0.4
12.1	12.5	12.9

Table 3.3: Spot positions on the ETM for the far left, inner left, central and right positions, respectively. The positions are shown in groups of three corresponding to the offset applied to V_L . All spot positions have an error of ± 1 mm.

upper limit of 0.11% given the error from determining the alignment by visual inspection. This corresponds to a negligible error of less than $\pm 30 \mu\text{m}$ in the results, showing that the dominant source of error in the experiment comes from the spot positions.

Knowledge of the distance of the ETM's voice coils from the centre of rotation, y_c ; the ETM's moment of inertia, I ; the coil driving frequency, f ; and the force calibration from Equation 3.10, allowed the rotation angle to be obtained geometrically:

$$\theta = \frac{F y_c}{4\pi^2 f^2 I}. \quad (3.11)$$

The numerical simulation tool *Finesse* (see appendix C.1.1) was then used to calculate κ for the cavity parameters shown in table 3.2. This was determined to be 18.5 m rad^{-1} . The WGM's transverse displacement was then the product of κ and θ .

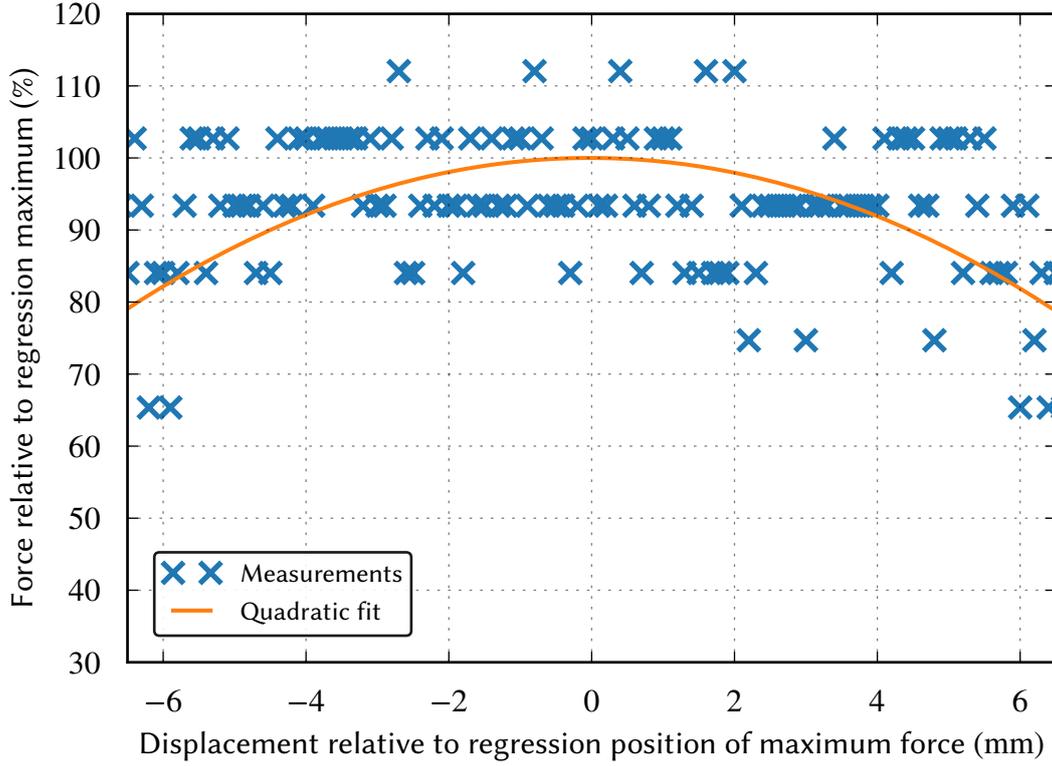


Figure 3.11: Change in force as a function of transverse displacement from voice coil axis. A quadratic fit has been applied to the data and the axes shown are with respect to the position and magnitude of the maximum fitted force, following the assumption that this position is nearest to the optimal alignment. This fit shows that, within the alignment error of the voice coils and magnets (0.5 mm), the maximum drop in force is negligible.

3.4.4 Analysis of the coupling level

Using the known contribution to the cavity length signal from the rotation of the ETM, δl_E , and the cavity length signals δl measured during the experiment, the WGM's most probable coupling level could be calculated statistically using Bayes' theorem. For this experiment, Bayes' theorem can be expressed mathematically as

$$p(\vec{\omega}|\mathcal{D}) \propto p(\mathcal{D}|\vec{\omega}) p(\vec{\omega}), \quad (3.12)$$

where $p(\vec{\omega}|\mathcal{D})$ is the probability density distribution of the experimental parameters, $\vec{\omega}$, given the observed data, \mathcal{D} (the *posterior*); $p(\mathcal{D}|\vec{\omega})$ is the likelihood and $p(\vec{\omega})$ is the probability distribution of the experimental parameters. The observed data \mathcal{D} are the measured cavity error signals for each of the spot positions.

In this analysis we are primarily interested in estimates of the model parameters. We are therefore free to ignore the constant evidence factor $p(\mathcal{D})$ present in Bayes' theorem when calculating the posterior. In the future it may be of interest to compare different models

for the coupling level (or lack thereof), in which case the evidence could be calculated to obtain a model odds ratio.

3.4.4.1 Model and parameters

To obtain a posterior for the WGM's coupling level, it was necessary to build a model and state prior belief of the parameters' probability distributions.

In the model, the ETM's geometrical longitudinal effect at arbitrary spot position y (equation (3.7)) for the rotation and mirror depth used in the experiment was combined coherently with a specified level of WGM transverse to longitudinal coupling, ω_1 . It was then possible to predict the total change in cavity length δl as a function of spot position y , given the fixed parameters θ , κ and d , using equations (3.3) and (3.7):

$$\begin{aligned} \delta l (\vec{\omega}, y, \theta, \kappa, d) &= \delta l_W (\theta, \kappa, \omega_1) + \delta l_E (y, \theta, d) \\ &\approx \theta \kappa \omega_1 + y \theta + \frac{d}{4} \theta^2. \end{aligned} \tag{3.13}$$

The effect of *beam smearing* was also considered. The suspended optics contain residual displacement noise, leading to a broadening of the trough at which the ETM's longitudinal coupling and any WGM coupling cancel (see figure 3.6). To model this effect, the assumption was made that the motion of the spots on the ETM followed a Gaussian distribution about their nominally measured position. Eight-hundred small 'offset distances' δy were applied uniformly to the spot positions, drawn from a randomly generated Gaussian distribution. The number of offset distances was chosen as a compromise between adequate statistical significance and computation time. Calculating the cavity length change as a function of spot position for each of these offset positions, and combining them in an uncorrelated sum, allowed an average, 'smeared' signal to be modelled which more closely resembled the measurements. The standard deviation of the Gaussian distribution was an additional parameter, ω_2 , provided as an input to the model.

The summing of signals introduced by the modelling of beam smearing led to an artificial increase in the magnitude of the model's predicted cavity length signals. To compensate for this effect, a further parameter was introduced: a multiplicative scaling factor, ω_3 , applied uniformly to the model's predicted cavity length signals. This factor also had the additional effect of compensating for the uncertainty in the calibrated cavity length signals. By marginalising over a suitable distribution of scaling factors, it was possible to account for this uncertainty in the analysis of the WGM's coupling level. The model used in the

Parameter	Symbol	Prior Distribution	Dimensions
WGM transverse to longitudinal coupling	ω_1	Uniform, $\left[0, \frac{1}{1000}\right]$	$\frac{m \text{ (longitudinal)}}{m \text{ (transverse)}}$
Spot smearing noise standard deviation	ω_2	Uniform, $\left[0, 3 \times 10^{-3}\right]$	m (transverse)
Multiplicative scaling factor	ω_3	Uniform, $\left[0, \frac{1}{10}\right]$	

Table 3.4: The distributions assumed for each of the free parameters in the model, along with their dimensions, prior to the computation of the posterior.

analysis to predict the smeared, scaled cavity length change, $\delta l'$, was then

$$\delta l'(\vec{\omega}, y, \theta, \kappa, d) = \omega_3 \sqrt{\sum_{i=1}^{800} \delta l(\vec{\omega}, y + \delta y_i, \theta, \kappa, d)^2}, \quad (3.14)$$

where δy_i is the i^{th} offset distance, drawn from a Gaussian distribution with standard deviation ω_2 .

3.4.4.2 Likelihood

The likelihood function assumed for the model was a Gaussian distribution,

$$p(\vec{\omega}|\mathcal{D}) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^N \frac{(D_i - \delta l'(\vec{\omega}, y_i, \theta, \kappa, d))^2}{\sigma^2}\right), \quad (3.15)$$

where N is the number of spot positions and σ^2 is the (identical) variance of each of the measured spot positions.

3.4.4.3 Priors

Bayes' theorem requires an assumption of probability distributions (*priors*) for each of the free parameters prior to the consideration of the measured data. The assumptions made for each free parameter in the model can be found in table 3.4. The upper bound on coupling was assumed to be a factor 10 better than the grating mirror measured in ref. [120], given the indication from ref. [116] that no coupling is present. The bounds on the scaling factor and spot smearing standard deviation were chosen from earlier observations of the behaviour of the signals during the experiment. All priors were assumed to be uniform.

3.4.4.4 Algorithm

A form¹ of the Metropolis-Hastings Markov-Chain Monte-Carlo (MCMC) algorithm [128] was applied to the model to marginalise over the three parameters. The outputs of the MCMC are a chain of samples (values at each parameter) that are drawn from the posterior distribution. A histogram of samples for a given parameter gives the marginal posterior distribution for that parameter from which the mean and standard deviation can be calculated.

To ensure the convergence of the MCMC on the posterior, a ‘burn-in’ period of 100 000 iterations was performed. The convergence was verified manually following completion. A further 100 000 iterations were then used to sample from the posterior and this second set is the one that we used for our results.

3.5 Results

From the parameter marginalisation it was possible to produce a posterior probability density distribution for the coupling level as shown in figure 3.12. The coupling level predicted from the distribution is bounded between 0 and 1:17000 with 95% confidence, with a mean coupling level of 1:27600. The probability density distributions for the scaling and standard deviation parameters are shown in figure 3.13. The scaling posterior distribution indicates a mean value of 29.3×10^{-3} with standard deviation 0.94×10^{-3} . The posterior distribution for the beam smearing parameter indicates a range of possible values between 0 and 1.3×10^{-3} m. All of the posterior distributions lie well within their prior ranges (see table 3.4), showing that the marginalisation did not indicate that a better fitting set of parameters could be found beyond the space defined by the priors.

The measured cavity length signals as well as the 95% upper limit and mean WGM coupling level predicted by the analysis are shown in figure 3.14. The phase discrepancy between the model and the measurements, as witnessed in this figure most profoundly for the spot positions around -5×10^{-3} m, is thought to be an artefact from the modelling of the beam smearing effect. The residual test mass motion that motivated the inclusion in the model of beam smearing may have contained some non-Gaussian behaviour; this is indicated by the shape of the beam smearing distribution shown in figure 3.13.

¹“Yet Another Matlab MCMC code” by Matthew Pitkin. Available as of time of writing at <https://github.com/mattpitkin/yamm>.

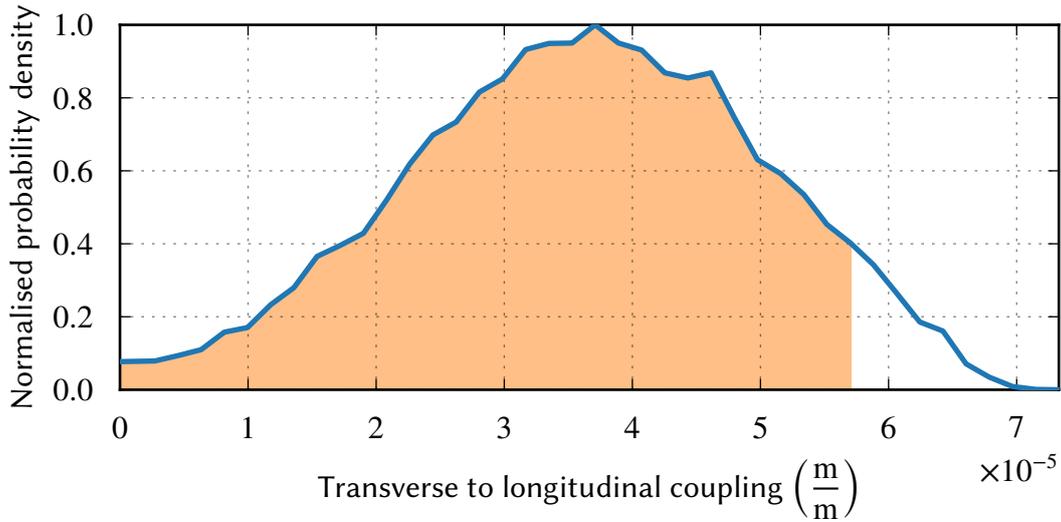


Figure 3.12: Posterior probability density distribution of WGM coupling levels (in units of meters longitudinal per metre transverse) yielded by statistical analysis of the data. The shaded region shows the coupling levels falling within the most probable 95% of the distribution.

3.6 Outlook

The upper limit on the predicted coupling level, 1:17000, represents a significant improvement over previously measured grating designs such as the 2nd order Littrow grating measured in ref. [120].

The evidence for transverse to longitudinal coupling is consistent with zero from figure 3.12, though it is not the most likely value from the experiment despite the literature suggesting otherwise. A future improvement to this experiment might be to perform local measurements of the rotation of the ETM with additional sensors, to avoid the calibration susceptible to measurement error outlined in section 3.4.3.

While the reduction in Brownian thermal noise in WGMs is a clear advantage to future gravitational wave detectors, other technologies at similarly early stages of technical readiness present similar improvements, such as the use of crystalline coatings [33]. The mirror coating technique that eventually becomes standard for gravitational wave observatories may depend on the direction that another field, fabrication engineering, takes in the near future.

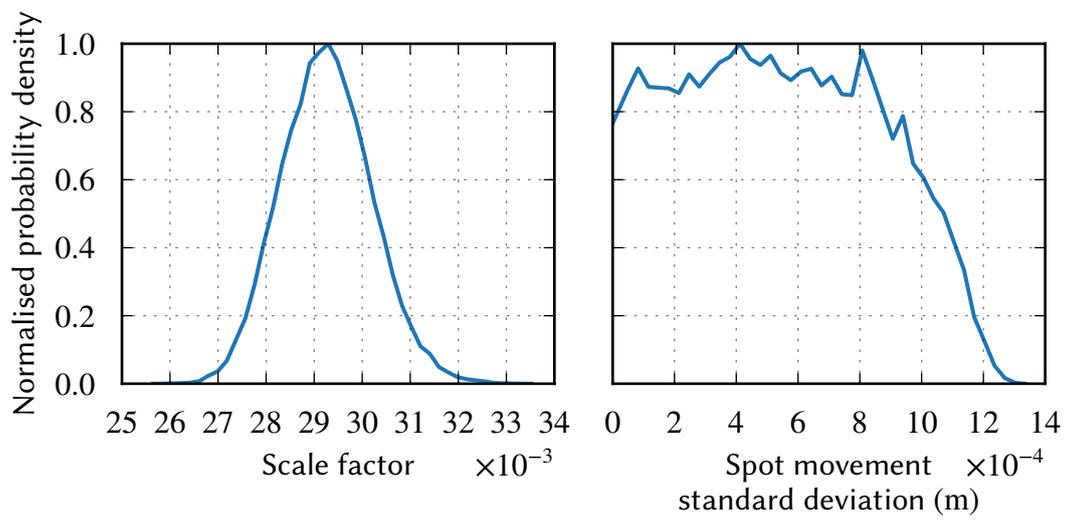


Figure 3.13: Posterior probability density distributions of scaling applied to the model's predicted longitudinal signal (left) and standard deviation assumed for the Gaussian distribution used to model beam smearing (right).

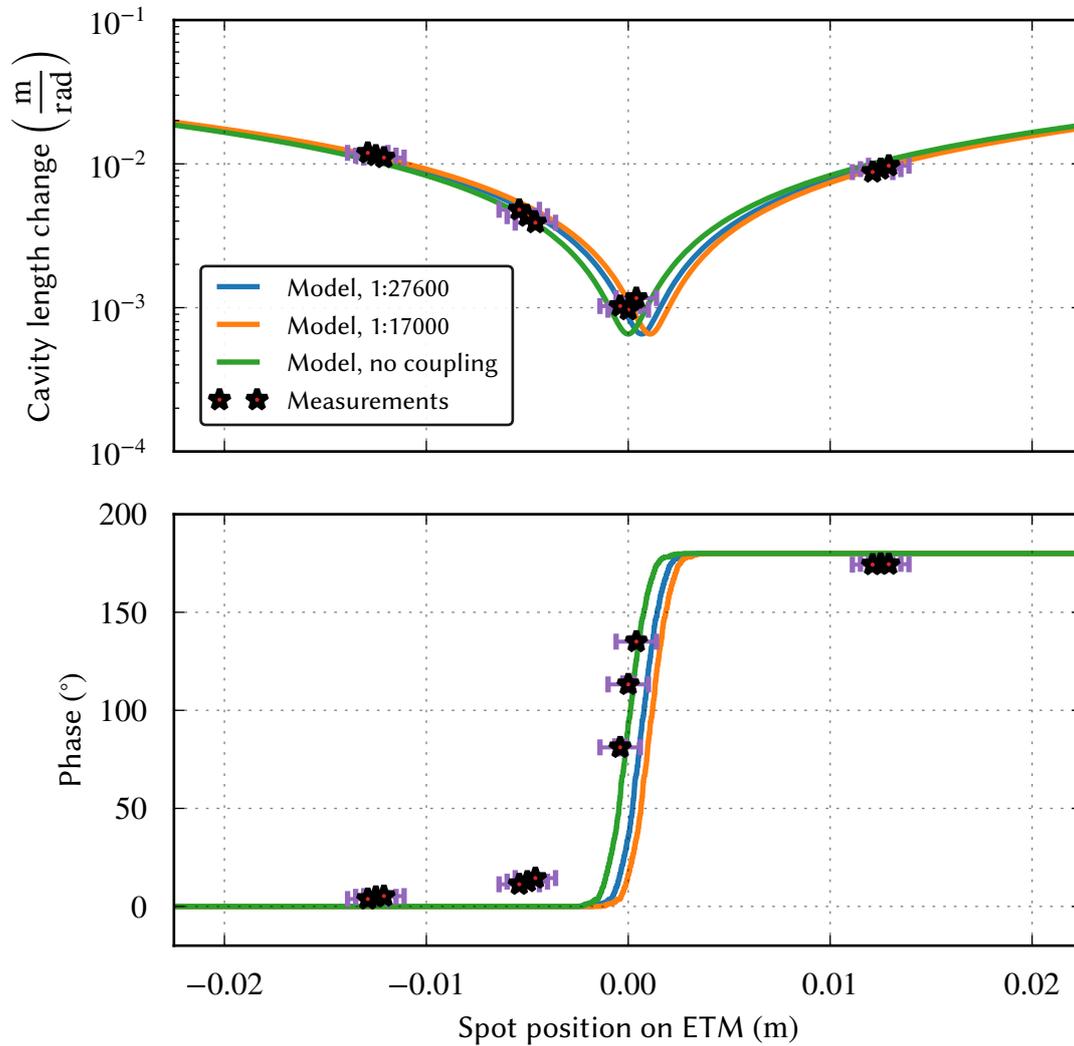


Figure 3.14: Measurements and simulations of the cavity length signal for spot positions with respect to the ETM's centre of yaw. The calibrated cavity length change per radian (vertical axis) from the measurements is shown (blue stars) alongside the model's simulated cavity length changes per radian for the mean (red), 95% upper limit (green) and zero (black) WGM coupling levels. The simulated plots use a scaling factor of 29.3×10^{-3} and a beam smearing standard deviation of 0.8×10^{-3} m.

Error bars are shown on the measured spot positions corresponding to their uncertainty. The errors in cavity length change are obtained from the noise floor surrounding each measurement. The noise floors were approximately constant for all measurements, with mean value 8×10^{-5} m rad $^{-1}$. Phase error bars are visible for the central values. The errors on each phase measurement, from left to right, are: ± 0.0188 , ± 0.0254 , ± 0.0283 , ± 0.1387 , ± 0.1721 , ± 0.2178 , ± 3.2726 , ± 3.2303 , ± 2.0603 , ± 0.0385 , ± 0.0342 and ± 0.0336 degrees.

Chapter 4

The Sagnac speed meter experiment: introduction and technical design

This chapter introduces the ongoing Sagnac speed meter experiment in Glasgow and serves as background for the work presented in chapters 5 and 6. The first half of this chapter discusses the measurement of test mass displacement and speed in the context of interferometers, and compares the sensitivity of the two measurements. The second half details the particular setup employed in the Glasgow experiment.

4.1 Position and speed meters

4.1.1 Sensitivity of a position meter

In an ordinary Fabry-Perot Michelson interferometer the motion of each cavity in the longitudinal direction either increases or decreases the round-trip phase of the light in that arm. The phase difference of the two recombined beams at the beam splitter then leads to a signal at the output port proportional to the differential phase which can be measured using a heterodyne or homodyne readout as discussed in section 2.1.3.

The presence of classical light power in the arms leads to dc radiation pressure which imparts a force upon the test masses. The reaction of the mirrors' restoring force, either from its pendulum in a suspended experiment or its mount on a table-top experiment, means that the interferometer can be held at the operating point as discussed in section 2.1.2 by making microscopic corrections to the position of the optic. The output signal can be calculated by considering *input-output relations* which define the effect that input light has at the output given the interferometer dynamics.

4.1.1.1 Input-output relations

We use the *two-photon formalism* [129, 130] in order to calculate the effect that the interferometer has on the light's amplitude and phase. This represents the input and output in terms of its *cosine* and *sine* quadratures, namely:

$$\vec{a} = \begin{pmatrix} a_c \\ a_s \end{pmatrix} \quad (4.1)$$

$$\vec{b} = \begin{pmatrix} b_c \\ b_s \end{pmatrix}, \quad (4.2)$$

where \vec{a} and \vec{b} are the input and output field vectors, respectively, both amplitude spectral density functions with units of $\sqrt{\text{Hz}^{-1}}$. The output for an interferometer can be expressed as the sum of the strain scaled by the interferometer dynamics and the noise present at the detector [89]:

$$\vec{b} = \frac{\tilde{h}}{\tilde{h}_{\text{SQL}}} \vec{R} + \mathbb{T} \vec{a}, \quad (4.3)$$

where \tilde{h} is the power spectral density of the strain applied to the test masses in units of Hz^{-1} , normalised to the amplitude spectral density of the SQL, \tilde{h}_{SQL} , as presented in equation (2.22); \vec{R} is the (dimensionless) response of the interferometer from strain to the output and \mathbb{T} represents the (dimensionless) transfer matrix of the input fields to the output fields.

The differential arm response, $\vec{R}_{(-)}$, is given by

$$\vec{R}_{(-)} = e^{i\beta_{\text{FP}}} \sqrt{2\kappa} \vec{H}, \quad (4.4)$$

where β_{FP} is the round-trip phase of the light in the arms, κ is the optomechanical coupling factor (as introduced in section 2.2.3.3) and \vec{H} represents the cosine and sine quadratures of the readout angle, ζ :

$$\vec{H} = \begin{pmatrix} \cos \zeta \\ \sin \zeta \end{pmatrix}. \quad (4.5)$$

The round-trip phase is defined as

$$\beta_{\text{FP}} = \arctan \frac{f}{\gamma_{\text{arm}}}, \quad (4.6)$$

where f is frequency and γ_{arm} is the Fabry-Perot cavity half-bandwidth (see appendix A.2.1).

The effect that the mirror dynamics have on the readout is governed by κ_{MI} defined for a Michelson interferometer in equation (2.24). This models the effect that a force applied to a mirror has on its position. Mirrors suspended from pendulum systems can be approximated

at high frequencies to be free masses, where the effect an applied force has on the position is diminished at higher frequencies, and in the case of a Fabry-Perot Michelson interferometer this force-to-displacement filtering effect scales as $\frac{1}{f^2}$ below the cavity's pole frequency, and $\frac{1}{f^4}$ above it.

The term \mathbb{T} can be further broken down:

$$\mathbb{T} = e^{2i\beta_{\text{FP}}} \begin{pmatrix} 1 & 0 \\ -\kappa & 1 \end{pmatrix}, \quad (4.7)$$

showing that the optomechanical coupling factor transforms the input \vec{a} in the cosine (amplitude) quadrature into the sine (phase) quadrature governed by the mirror dynamics on its way to the output. The matrix of noise power spectral densities for the quadratures of the light at the output port, \mathbb{S} , is given by averaging over all frequencies:

$$\mathbb{S} = \langle \vec{b} \cdot \vec{b}^\dagger \rangle. \quad (4.8)$$

Expanding the terms we get

$$\mathbb{S} = \begin{pmatrix} 1 & 0 \\ -\kappa & 1 \end{pmatrix} \langle \vec{a} \cdot \vec{a}^\dagger \rangle \begin{pmatrix} 1 & -\kappa \\ 0 & 1 \end{pmatrix}, \quad (4.9)$$

where h in equation (4.3) has been set to 0 to remove signal.

4.1.1.2 Sensitivity of a Fabry-Perot Michelson interferometer

In the case of dc readout as discussed in section 2.1.3.2 and used in current generation detectors, the readout angle $\zeta = 90^\circ$ represents the phase quadrature:

$$\vec{H}_{\text{dc}} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (4.10)$$

We can calculate the signal in the absence of noise at the dc readout of the Fabry-Perot Michelson interferometer, $O_{(-)}^{\text{dc}}$, as a function of differential arm cavity strain $\tilde{h}_{(-)}$ by rearranging equation (4.3):

$$O_{(-)}^{\text{dc}} = \frac{\tilde{h}_{(-)}}{\tilde{h}_{\text{SQL}}} e^{i\beta_{\text{FP}}} \sqrt{2\kappa_{\text{MI}}}. \quad (4.11)$$

This is shown in figure 4.1 for $\tilde{h} = 1$, arm length 1 km, mirror mass 40 kg, laser wavelength 1064 nm and cavity half-bandwidth 250 Hz.

Normal, unsqueezed vacuum noise at the input has equal noise contributions in the cosine

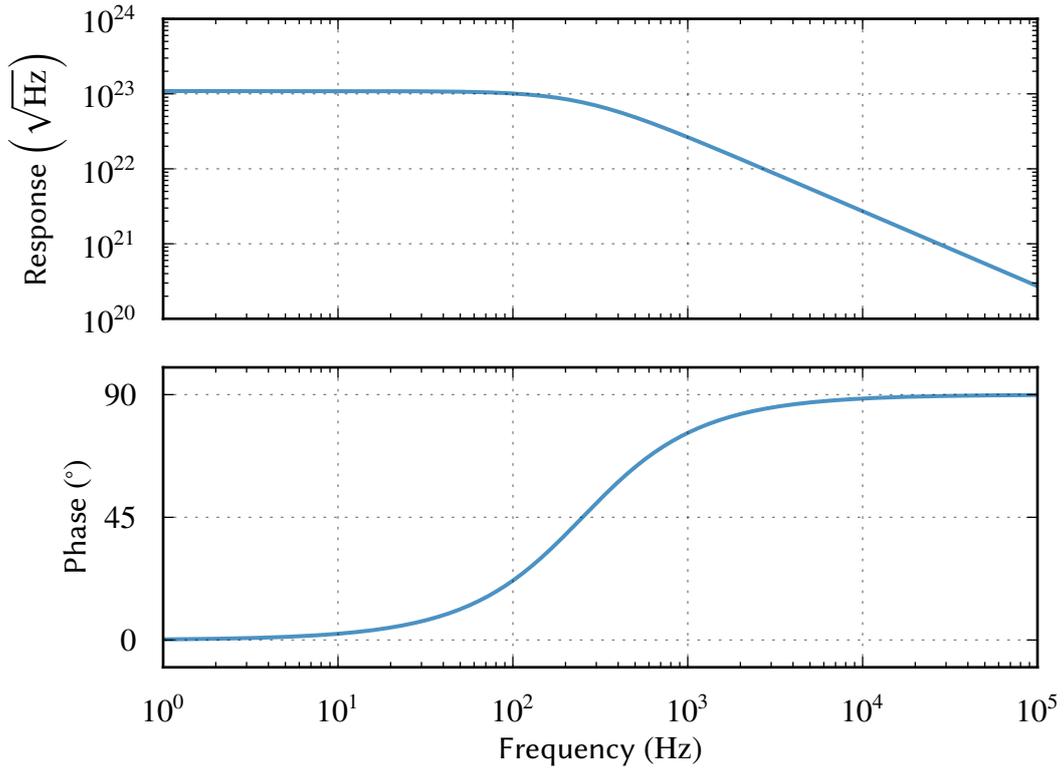


Figure 4.1: Response of a Fabry-Perot Michelson interferometer to differential arm cavity motion. This shows the signal that would appear at a photodetector placed at the output of the interferometer given unit differential motion of the cavity. The cavity provides a signal with the same response for motion at frequencies below the cavity's pole frequency. Beyond the pole, the response is reduced as the motion becomes faster than the cavity's storage time.

and sine quadratures, and so we can set it to the identity matrix:

$$\langle \vec{a}_{\text{vacuum}} \cdot \vec{a}_{\text{vacuum}}^\dagger \rangle = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (4.12)$$

The matrix of noise power spectral densities of the light quadratures at the output due to quantum noise, \mathbb{S}_{FPMI} , is then

$$\begin{aligned} \mathbb{S}_{\text{FPMI}} &= \begin{pmatrix} 1 & 0 \\ -\kappa_{\text{MI}} & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -\kappa_{\text{MI}} \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & -\kappa_{\text{MI}} \\ -\kappa_{\text{MI}} & 1 + \kappa_{\text{MI}}^2 \end{pmatrix}. \end{aligned} \quad (4.13)$$

The noise seen by the dc readout is determined by the homodyne angle:

$$\mathcal{S}_{\text{FPMI}}^{\text{dc}} = \vec{H}_{\text{dc}}^T \mathbb{S}_{\text{FPMI}} \vec{H}_{\text{dc}}. \quad (4.14)$$

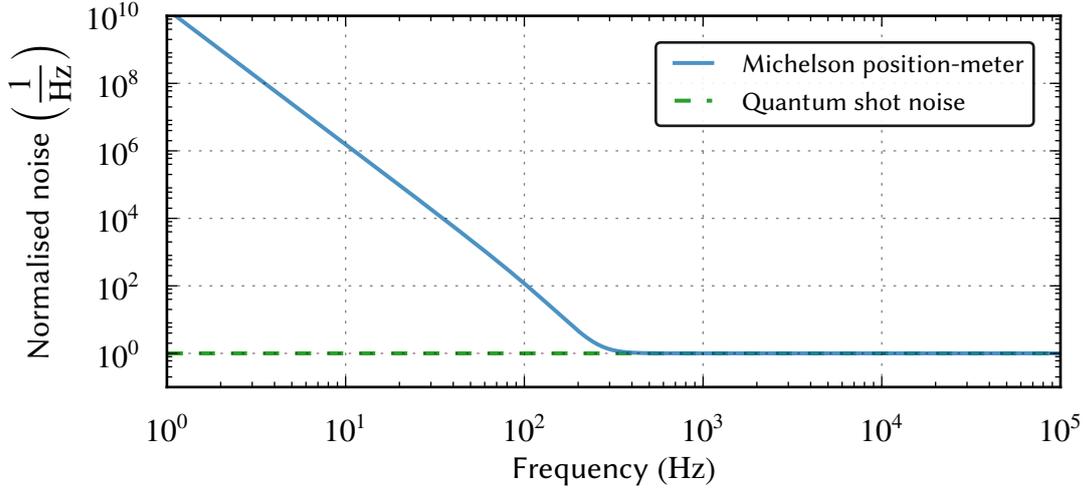


Figure 4.2: Quantum noise of a Fabry-Perot Michelson interferometer at the output port, normalised to quantum shot noise. This shows the noise present at the photodetector produced due to quantum noise entering the interferometer and interacting with the mechanics. At high frequencies, the noise in a Michelson interferometer is almost entirely due to the quantum shot noise on the sensor; at low frequencies the noise is dominated by the light reaching the sensor due to fluctuations in the positions of the test masses due to quantum radiation pressure noise.

For dc readout, the output noise spectral density is shown in figure 4.2. This is the combination of radiation pressure noise from the mirrors and shot noise on the sensor, and these two effects combine to produce the quantum noise spectral density. At high frequencies, the quantum noise is equal to the quantum vacuum noise input from equation (4.12) (corresponding to $\kappa_{\text{MI}} \approx 0$) which shows that the signal on the sensor is limited by noise propagating to the output with no significant effect from optomechanical interactions. Below the cavity pole, the vacuum fluctuations move the mirror by an amount governed by the mirror's optomechanical coupling and this random detuning of the cavity converts coherent carrier light into radiation pressure noise at the output.

The quantum noise power spectral density at the output as a function of $\tilde{h}_{(-)}$, $S_{\tilde{h}_{(-)}}$, is given by the ratio of the quantum noise power spectral density at the sensor to the response of the interferometer to that sensor for differential arm cavity strain, i.e.

$$S_{\tilde{h}_{(-)}} = \frac{S_{\text{FPMI}}^{\text{dc}}}{\left|O_{(-)}^{\text{dc}}\right|^2}, \quad (4.15)$$

in units of Hz^{-1} . This is a useful indication of the sensitivity at the output of the interferometer. The more common amplitude spectral density representation, $s_{\tilde{h}_{(-)}} = \sqrt{S_{\tilde{h}_{(-)}}$, is shown in figure 4.3.

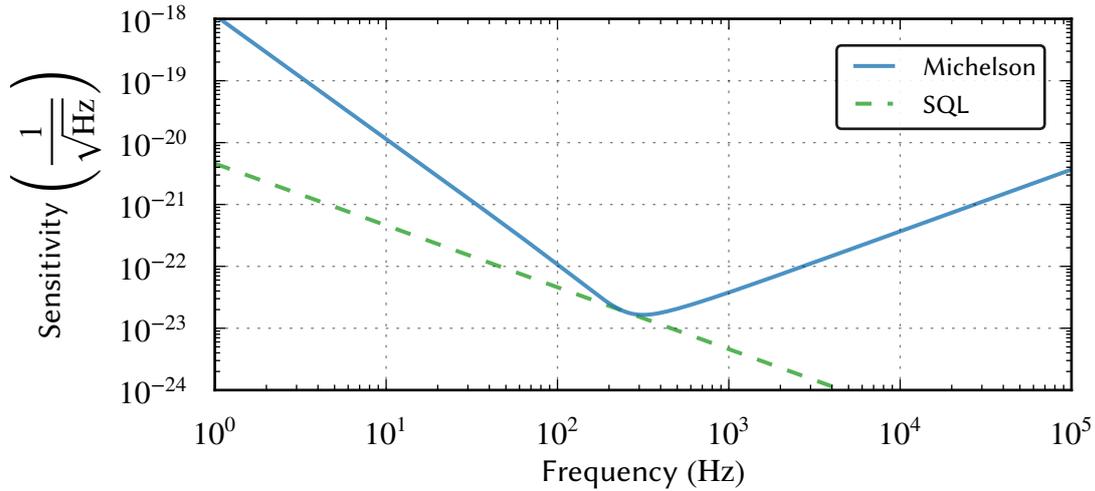


Figure 4.3: Quantum noise limited sensitivity of a Fabry-Perot Michelson interferometer at the output port to differential arm cavity motion. This is calculated by taking the quantum noise at the probe shown in figure 4.2 and dividing it by the response from differential arm cavity motion to the probe shown in figure 4.1. In this case the cavity power was chosen to touch the SQL at the cavity pole and this is shown in the figure. For smaller cavity power, the touching frequency moves down in frequency.

4.1.2 Sensitivity of a speed meter

Since the early 1990s it has been known that the measurement of momentum, known to be a quantum non-demolition (QND) observable, offers the ability to surpass the SQL in interferometric measurement [65]. Ideally, the back-action applied to test masses by a measurement of momentum—a consequence of the Heisenberg Uncertainty Principle—does not affect its future value and so momentum can in principle be measured to arbitrary precision. Velocity is an appropriate observable to measure momentum and also approximates a QND scheme due to its relation to momentum. Interferometers that measure velocity are called *speed meters*, and their principle of operation is as follows. Light from a laser enters the interferometer as it would for a position-meter, and accumulates a phase shift proportional to the propagation and the signal from any gravitational waves or disturbances in the positions of the test masses. As the light reflects from the test masses it imparts radiation pressure arising from its classical amplitude and the amplitude quantum vacuum fluctuations as discussed in section 4.1.1 for a Fabry-Perot Michelson interferometer. Within the interferometer there must be a mechanism to impart a phase shift equivalent to 180° to one light field to create a second light field that samples the same mode. Propagating through the interferometer, the radiation pressure imparted to the mirrors by one field is superimposed upon the radiation pressure imparted from the other, and as these effects are out of phase within the light travel time the radiation pressure force can be suppressed.

There are many different speed meter topologies in the literature [88, 90, 91, 131]. Speed

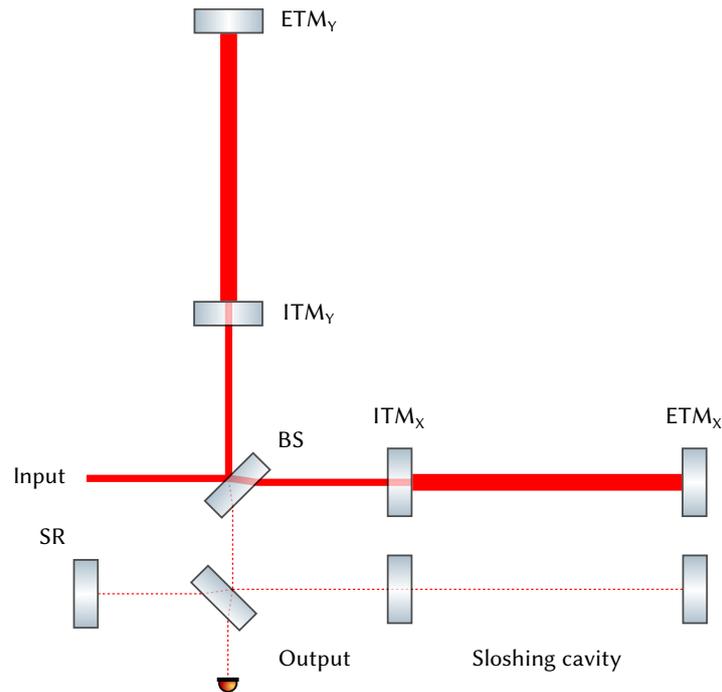


Figure 4.4: Layout of a Fabry-Perot Michelson interferometer with a sloshing cavity as presented in ref. [86]. The light leaving the Fabry-Perot Michelson interferometer is coupled into a sloshing cavity by a beam splitter where it receives a phase shift, and it re-enters the interferometer via the recycling mirror to the left of the sloshing beam splitter. The light incident upon the beam splitter then contains light that has sampled the mirrors at two points in time, leading to a speed meter effect.

meters are also being considered as alternatives to the proposed Michelson interferometers in the Einstein Telescope [132, 133]. We consider here two speed meter topologies to highlight the significantly different forms in which a speed meter interferometer can take.

4.1.2.1 The Michelson-type speed meter

Initial suggestions for the application of speed meter type interferometers in the field of gravitational wave detection were focused on a Fabry-Perot Michelson interferometer topology with the addition of a *sloshing cavity* at the output port [85, 86] as shown in figure 4.4. Here the 180° phase shift is imparted to the light by the addition of a beam splitter and sloshing cavity at the output of the interferometer. The light returning from the sloshing cavity is either re-injected into the interferometer or transmits through the beam splitter where it reflects from a signal recycling mirror (SR). The light at the output of this interferometer then contains reduced quantum radiation pressure noise.

4.1.2.2 The Sagnac-type speed meter

It was realised by Chen that the *zero-area Sagnac* interferometer topology is a speed meter [87]. This interferometer is arranged such that incident photons enter into two counter-propagating modes which sample the position of the test masses at different intervals. The Sagnac interferometer is sensitive to the rotation of the Earth via the area enclosed by its arms, and so to avoid this the propagation of the light is arranged in a zero-area configuration to cancel the rotation-induced phase accumulation from each arm. The remaining signal at the output contains information of the difference in round-trip phase of the two counter-propagating modes due to test mass motion. Given two test mass positions x_A and x_B in arms A and B , respectively, over a time interval of Δt each counter-propagating mode will measure phase changes $\delta\phi_A$ and $\delta\phi_B$ arising from motion of the arms less than the light propagation time [87]:

$$\delta\phi_A \propto \Delta x_A(t) + \Delta x_B(t + \Delta t) \quad (4.16)$$

$$\delta\phi_B \propto \Delta x_B(t) + \Delta x_A(t + \Delta t). \quad (4.17)$$

At the output port, the combined signal will then be the difference of phase,

$$\begin{aligned} \delta\phi_A - \delta\phi_B &\propto (\Delta x_A(t) - \Delta x_A(t + \Delta t)) - (\Delta x_B(t) - \Delta x_B(t + \Delta t)) \\ &\propto \Delta \dot{x}_A(t) - \Delta \dot{x}_B(t), \end{aligned} \quad (4.18)$$

which shows that the signal is proportional to the relative velocity of the test masses. The output port is automatically at the dark fringe for the carrier light as long as the motion of the test masses is slower than the light propagation time. The output is not dark for the signal sidebands, and as they contain components from the test masses sampled at different times the signal is proportional to test mass speed.

The layout of a Sagnac speed meter interferometer can be arranged in different forms [91], and we show one based on a zero-area Sagnac enhanced with ring cavities as arms in figure 4.5.

4.1.2.3 Input-output relations

The same approach to that for a Fabry-Perot Michelson interferometer in section 4.1.1 can be taken to calculate the response and noise of a speed meter, but with a value of κ modified for a speed-meter [87],

$$\kappa_{SM} = 4\kappa_{MI} \sin^2 \beta_{FP}, \quad (4.19)$$

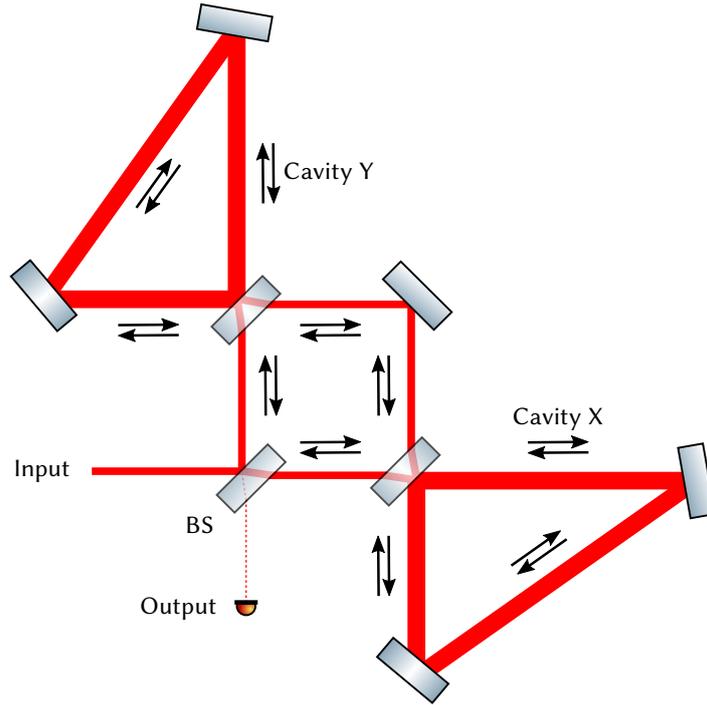


Figure 4.5: Layout of a zero-area Sagnac speed meter with ring cavities. The input light is split at the beam splitter where it forms two counter-propagating modes within the inner Sagnac mirrors, denoted by black arrows. At each ITM, the light is partially transmitted into the arm cavities, and upon exiting the cavities this light is either sent back to the beam splitter or to the next cavity. The recombined light at the beam splitter contains fields that have interacted with all of the mirrors and the difference in phase between the counter-propagating modes provides a signal proportional to relative test mass velocity.

and as the round-trip phase includes both arms and an extra reflection from or transmission through the beam splitter, it is also modified:

$$\beta_{\text{SM}} = 2\beta_{\text{FP}} + \frac{\pi}{2}. \quad (4.20)$$

The response of a Sagnac speed meter to differential arm cavity motion is shown in figure 4.6 for parameters identical to that of figure 4.1. Notice that below the cavity pole, the response vanishes towards dc, consistent with a speed measurement. The higher response above the cavity pole is a consequence of the fact that the light samples the interferometer in both directions. For fair comparisons to the Michelson interferometer the choice may be made to alter the input power and readout angle of one with respect to the other.

The corresponding quantum noise at the output port is shown in figure 4.7. Note that the noise is unity at high frequencies as with the Fabry-Perot Michelson interferometer, but is suppressed at low frequencies due to the cancellation of back-action due to radiation pressure from quantum vacuum fluctuations. The cancellation is not perfect due to the time delay between the two consecutive visits of the arm cavities by the counter-propagating

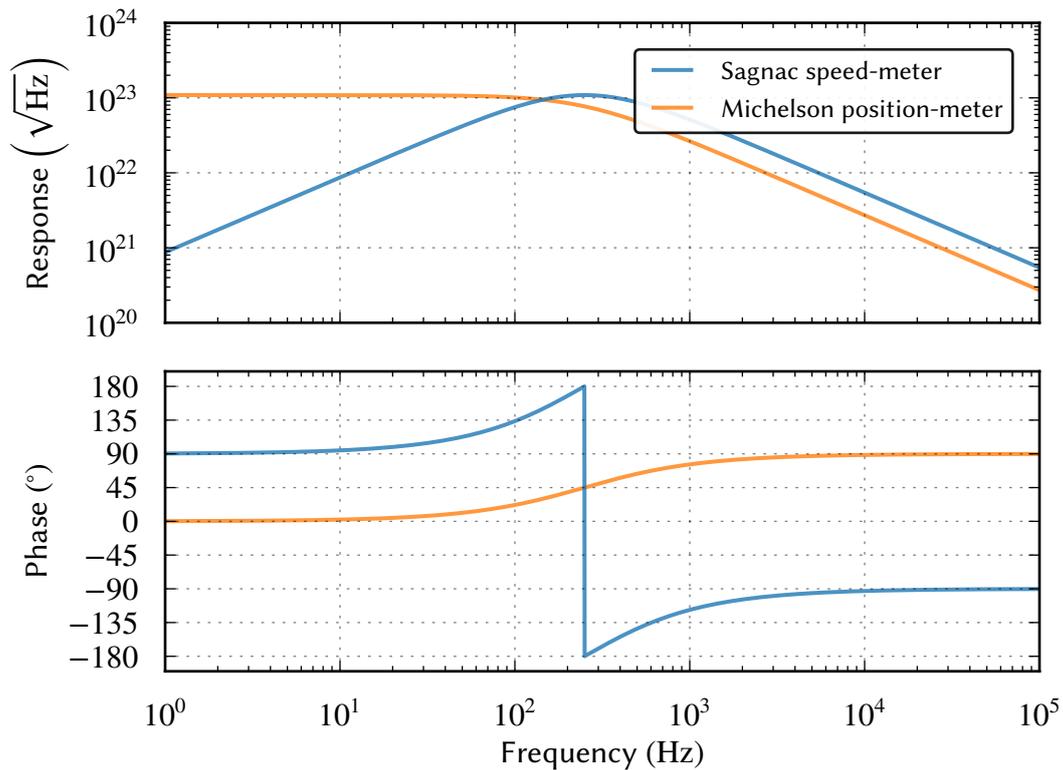


Figure 4.6: Response of a Sagnac speed meter to differential arm cavity motion. In contrast to the Michelson interferometer, the Sagnac speed meter has response proportional to frequency below the cavity pole.

modes.

While the response in a Sagnac speed meter is reduced at low frequencies, the quantum noise is further reduced and so the overall quantum noise limited sensitivity is improved over an equivalent Fabry-Perot Michelson interferometer in the absence of loss, as shown in figure 4.8. For lossy speed meters the sensitivity is degraded. In the next subsection we consider loss in the case of a Sagnac speed meter but loss in any QND interferometers significantly affects sensitivity.

4.1.2.4 Loss in Sagnac speed meters

The QND behaviour of the interferometer arises from the fact that the output port contains only commutative time-dependent momentum information. Time-independent position information can, however, enter the output port of the interferometer in the presence of certain types of loss [88]. For *symmetric* loss, such as from balanced but imperfect reflectivity of the ETMs or substrate absorption in the ITMs, incoherent vacuum fluctuations can enter the interferometer at a point after the light has been split into the counter-propagating modes and this affects sensitivity. It has also been shown that *asymmetric* loss results in a

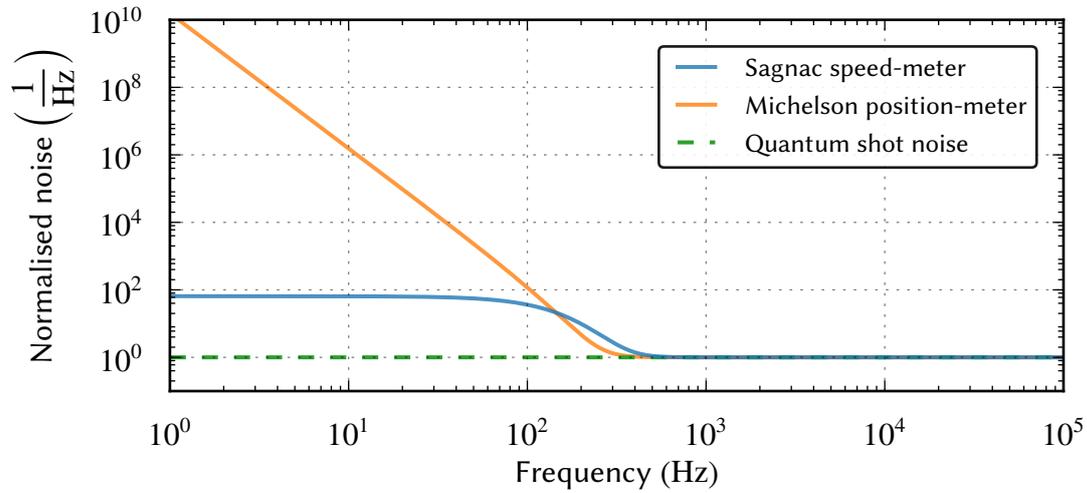


Figure 4.7: Quantum noise of a Sagnac speed meter at the output port, normalised to quantum shot noise. Like the Michelson interferometer, the high frequency noise contribution arises from quantum shot noise from incoherent vacuum fluctuations entering the interferometer. In contrast to the Michelson interferometer, the Sagnac speed meter has flat noise at low frequencies below a transition region, as the test mass noise fluctuations are cancelled by the counter-propagating modes in the instance where the quantum radiation pressure forces are balanced.

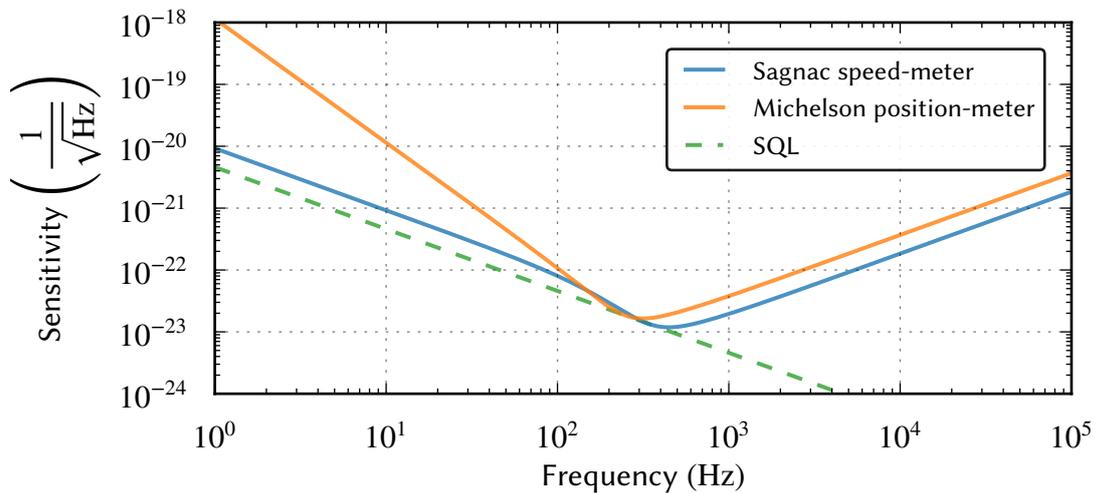


Figure 4.8: Quantum noise limited sensitivity of a Sagnac speed meter at the output port to differential arm cavity motion. In contrast to the Michelson interferometer, the Sagnac speed meter sensitivity at low frequencies follows the gradient of the SQL due to its reduced quantum noise. This improved sensitivity is in practice difficult to achieve as the presence of loss in the interferometer introduces a *Michelson-like* sensitivity slope at a frequency proportional to the level of loss.

greater decrease in sensitivity [89], considering effects from imperfect beam splitting and imbalanced ITM reflectivity in a Sagnac speed meter.

The optic in a Sagnac speed meter most susceptible to asymmetries is typically the beam splitter, as coatings typically cannot be manufactured with better than around 0.1 % tolerance in the amplitude reflectivity at the standard wavelength for detectors¹. Imperfect splitting leads to different power in the counter-propagating modes which leads to asymmetries. In the Sagnac speed meter some of the light that would otherwise have exited at the input port of the interferometer (towards the input laser) instead exits at the output port due to the imbalanced beam splitter, carrying time-independent signal and therefore damaging the sensitivity. This appears on displacement sensitivity curves as an additional $\frac{1}{f}$ slope at low frequencies such that it resembles the $\frac{1}{f^2}$ displacement sensitivity of a Michelson interferometer. The minimisation of loss is therefore critical in the design of a Sagnac speed meter experiment.

4.2 The Glasgow Sagnac speed meter experiment

The Sagnac interferometer has been demonstrated in table-top experiments [134] as well as lab-scale prototypes [135, 136] but so far the QND behaviour of the Sagnac speed meter topology has not been shown; nor has the topology been implemented as a suspended prototype with greatest sensitivity in the audio band to demonstrate its applicability for ground-based gravitational wave detectors. Here we present an ongoing proof-of-concept experiment based at the University of Glasgow to test audio-band reduction of quantum radiation pressure noise in a suspended Sagnac speed meter over an equivalent Michelson design [137].

The Glasgow Sagnac speed meter experiment is enhanced with the presence of triangular ring cavities to increase the sensitivity of the interferometer to differential motion of the test masses in the arms, and balanced homodyne readout will be utilised to sense the quantum correlations present at the output port. As quantum radiation pressure noise depends inversely upon the reduced mass of the arm cavities, in order for the interferometer to be dominated by quantum radiation pressure noise at low frequencies one of the *core* optics in each of the triangular arm cavities will be much lighter than the other two. The ETMs in this case will be around 100 g while the ITMs will be around 1 g. The intracavity power will be around 4 kW which is high enough to create significant radiation induced force measurable above other noise sources at frequencies between 100 and 700 Hz.

The intended optical layout is shown in figure 4.9. The input light is coupled by the main

¹It turns out that the Sagnac speed meter offers an excellent means of measuring a beam splitter's asymmetry.

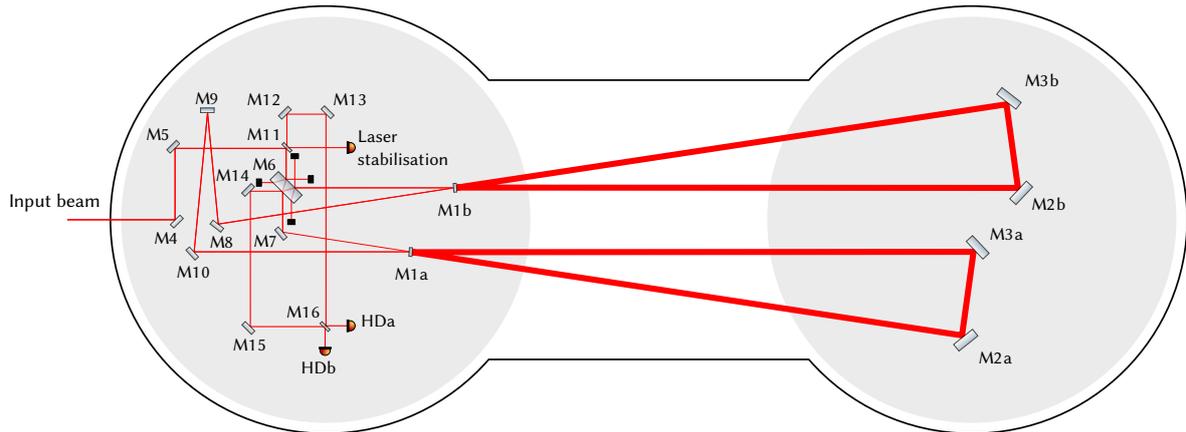


Figure 4.9: Sagnac speed meter experiment layout. The in-vacuum part of the experiment will be situated in two 1 m diameter tanks joined with a connecting tube. The suspended optics will be placed on breadboards atop passive isolation stacks, joined together with a bridge structure. Viewports are situated on both tanks at each side to facilitate in-air sensing. The vacuum system is capable of reaching pressures below 10^{-6} mbar to suppress the impact of residual gas noise.

beam splitter M_6 into counter-propagating modes in the *inner Sagnac*, i.e. the cavity formed by the mirrors M_6 , M_7 , M_1^A , M_{10} , M_9 , M_8 and M_1^B . At ITMs M_1^A and M_1^B the light is partially transmitted into arm cavities *A* and *B*, respectively, where the light is resonantly enhanced by the highly reflective ETMs M_2^A and M_3^A and M_2^B and M_3^B . Upon exiting the cavities, the light again propagates through the inner Sagnac where parts recombine at M_6 and parts enter the opposite arm cavity. The light leaving M_6 towards M_{14} contains the signal encoded as quantum correlations on the light, and this is enhanced at M_{16} by the local oscillator provided by the light leaving M_6 towards M_{12} . The balanced homodyne detectors (BHDs) HD_A and HD_B measure the transmitted light, and the signals primarily contain information regarding the relative velocity of the arm cavities.

The interferometer is to be situated within an ultra-high vacuum system formed from two adjoined cylindrical tanks with pumps capable of reaching pressures below 10^{-6} mbar. Each tank contains a breadboard for the attachment of components, and this breadboard is itself isolated from ground motion by a series of passive damping stacks. The breadboards are rigidly connected via a bridge structure to ensure that residual platform motion is common to all suspended optics. The optics will be suspended from pendulum systems, with the most important test masses suspended from multiple stages to provide additional isolation from seismic noise. The parameters for the optics, laser injection, suspension systems and materials can be found in refs. [137], [89] and [138] in chronological order.

4.2.1 Balanced homodyne detection

The sensor for the gravitational wave channel, the differential arm cavity degree of freedom, will be balanced homodyne readout, in contrast to the *de-facto* standard in gravita-

tional wave observatories, dc readout. It is difficult to change the dc readout quadrature to optimise the sensitivity in the presence of imprecisely known loss; the homodyne angle is fixed by the propagation length between the detuned arm cavities and the output port. With balanced homodyne readout arbitrary homodyne angles can be chosen by microscopically tuning of the relative phase of the (separate) local oscillator and signal paths.

Balanced homodyne readout involves making a subtraction of two signals measured in this case by HD_A and HD_B , observing light combined from a local oscillator, a , and the signal output of the Sagnac speed meter, b . The local oscillator field should not contain signal, and so this will be taken from the light reflected from the interferometer back towards the input, via M_{12} and M_{13} . The signal power measured at each photodetector output c and d then contains [74]:

$$\begin{aligned} c^\dagger c &= \frac{1}{2} (a^\dagger a + a^\dagger b e^{-i\phi} + ab^\dagger e^{i\phi} + b^\dagger b) \\ d^\dagger d &= \frac{1}{2} (a^\dagger a + a^\dagger b e^{i\phi} + ab^\dagger e^{-i\phi} + b^\dagger b), \end{aligned} \quad (4.21)$$

where ϕ is the homodyne angle. The mixing of the two signals at M_{16} results in the dc part of one field beating with the ac part of the other field. Subtracting the signals on the two balanced homodyne photodetectors yields a photocurrent, I_{BHD} :

$$\begin{aligned} I_{\text{BHD}} &= c^\dagger c - d^\dagger d \\ &= a^\dagger b e^{-i\phi} - a^\dagger b e^{i\phi} + ab^\dagger e^{i\phi} - ab^\dagger e^{-i\phi}, \end{aligned} \quad (4.22)$$

where we assume that the signal and local oscillator are split equally between the two detectors by the balanced homodyne beam splitter. This shows that, as long as the beam splitter has matched transmissivity and reflectivity, a delicate subtraction of the two photocurrents from HD_A and HD_B results in an error signal containing only the ac parts of the signal corresponding to the motion of the test masses amplified by the local oscillator field, and a small signal from the local oscillator path enhanced by the classical light at the output port.

In the presence of imbalanced beam splitting the laser noise and residual carrier light in the signal path results in additional noise photocurrent and so it is important to use a beam splitter with balanced reflectivity and transmissivity and for the intensity noise of the laser to be controlled to a high degree [74].

4.2.1.1 Quantum noise limited sensitivity of the main readout

Without including the effect of asymmetric loss, the quantum noise limited sensitivity of the BHD readout in the Sagnac speed meter experiment is shown in figure 4.10, calculated with

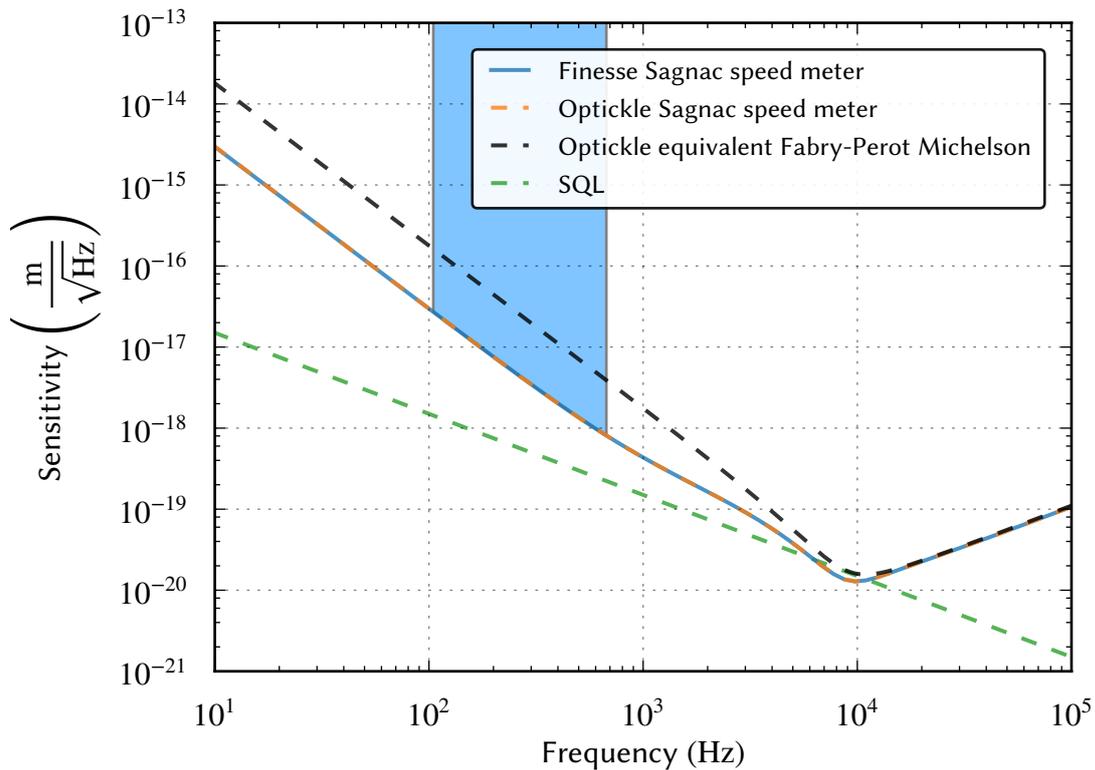


Figure 4.10: Predicted quantum noise limited sensitivity of the Sagnac speed meter experiment calculated with Optickle and Finesse. Also shown is the equivalent Fabry-Perot Michelson interferometer configuration, calculated with Optickle, and the SQL given the effective mass of the interferometer test masses. The shaded blue region shows the intended measurement band, where reduced quantum radiation pressure noise should be visible below the expected noise from the equivalent Fabry-Perot Michelson interferometer.

both Finesse and Optickle (see appendices C.1.1 and C.1.2). Also shown is the sensitivity of an equivalent Fabry-Perot Michelson interferometer using the same parameters as for the Sagnac speed meter, but with input power scaled by a factor of approximately 2.5 to match its high frequency sensitivity. The homodyne angles of the Sagnac speed meter and Michelson interferometer are 45° and 90° , respectively, to optimise both interferometers for high frequency sensitivity fairly. The intended measurement band is between 100 Hz and 700 Hz so a reduction of a factor of around 3 to 5 is in theory possible as long as other sources of noise are kept sufficiently low. A comprehensive consideration of the noise budget is given in chapter 5.

4.2.2 Suspensions

The *auxiliary* optics M_4 , M_5 , M_7 - M_{10} and M_{12} - M_{15} will be suspended from two stage pendulum systems. These steering optics do not have as stringent residual motion requirements as the test masses and so these suspensions can in comparison have a relatively simple design.

The core optics will be suspended from two different systems: the 100 g ETMs from triple stage suspensions based on the design for the ETM suspensions of the AEI 10 m prototype in Hanover, Germany, and the 1 g test masses from bespoke quadruple stage suspensions. The beam splitter M_6 will require another design loosely based on the auxiliary suspensions but with greater filtering. Each optic class requires a separate suspension design due to the differences in geometry, the requirements for residual motion and the need to move the mechanical modes from each suspension out of the measurement band.

4.2.3 Actuation

The auxiliary suspensions will have voice coil actuators on their intermediate stages for local alignment control. The 100 g suspensions will have voice coils on multiple stages in order to provide corrections for low frequency drifts, though the test mass stages will not have voice coils to prevent *Barkhausen* noise [139] coupling to the test masses via the actuators' magnets.

The filtering effect from the final pendulum stage means that the voice coil actuators will not be able to effectively correct test mass positions at high frequencies. Instead, *electrostatic drive* (ESD) actuators will be placed behind each ETM able to produce small corrective forces but without Barkhausen noise. ESD actuators have been demonstrated in GEO 600 [140] and Advanced LIGO [46] based on a metallic comb arrangement, though in the Sagnac speed meter experiment the intention is to use a new ESD design which will be discussed in chapter 6.

4.2.4 Sensing and control

While the main readout of the Sagnac speed meter experiment is to be the BHD, in order to control the interferometer a number of other signals from different ports will need to be extracted and fed back to actuators to keep the interferometer at the operating point. The control topics can be split into two broad categories representing the process used to bring the interferometer to the operating point, *lock acquisition*, and keeping it there, *low-noise control*.

The interferometer's test masses will drift from the operating point due to the noise imparted from quantum, seismic, thermal and other noise sources as discussed in section 2.1 if they are not actively controlled. This control takes the form of actuators at each test mass and feedback to the main laser's temperature and frequency. In the Sagnac speed meter experiment there are a few major control topics that must be solved in order to reach the required sensitivity:

1. the control of longitudinal drifts in the positions of the test masses, which leads to loss of cavity power and sensitivity;
2. the control of angular drifts in the optics, particularly from the triangular arm cavities and especially important for the small ITMs where radiation pressure effects will be significant;
3. and the control of intensity noise on the main laser.

The first two topics are tackled through the identification of the degrees of freedom of the interferometer and the selection of appropriate readout ports. The last involves the implementation of an appropriate frequency stabilisation control servo.

The process of bringing the interferometer to its operating point is challenging in interferometers with multiple degrees of freedom and typically requires modelling in the time domain to understand the effect that changes to actuator signals and mirror dynamics have on the system. In the case of the Sagnac speed meter experiment this is particularly challenging due to the coupling between the arm cavities from the counter-propagating modes. Some approaches to lock acquisition have been developed which should in principle be able to deterministically bring the interferometer to the operating point [141].

4.2.4.1 Data acquisition and software control

The control and data acquisition system developed for LIGO, CDS [142], is appropriate for the Sagnac speed meter experiment and it can benefit from the great deal of effort that has already gone into making this system useful and reliable for the control of complicated experiments. CDS takes the form of many “off-the-shelf” components and custom software to provide the ability to sense and feed back control signals at speeds of up to around 10 kHz. Extensive software is also available for offline data analysis.

The translation between the digital CDS domain and the analogue domain takes the form of analogue-to-digital and digital-to-analogue converters (ADCs and DACs) situated on *front-end* computers which run software control modules in real time. A *frame builder* computer communicates with the front ends over a fast network in order to build packets of measurements in GPS-synchronised intervals.

4.3 Topics of particular focus

There are many areas of research and development required in order to meet the challenging goals of the Sagnac speed meter experiment, but we will focus in particular on the

strategy for controlling the longitudinal degrees of freedom of the interferometer at its operating point. Chapter 5 will develop a longitudinal sensing and control scheme for the interferometer taking into account anticipated sensors, actuators and noise sources, and chapter 6 will present the design of the high voltage electronics and signalling for the ESD design to be used in the experiment.

Chapter 5

Concept for the longitudinal control of the Sagnac speed meter experiment

As shown in chapter 4, the Sagnac speed meter interferometer topology can potentially provide enhanced sensitivity to gravitational waves in the audio-band compared to equivalent Michelson interferometers. Using as an example the proof-of-concept Sagnac speed meter experiment in Glasgow, we discuss the issues surrounding the control of this type of interferometer and quantify the challenges using numerical simulations. We present a solution involving the extraction of multiple error signals that can be blended to produce corrective signals to be applied to the test mass actuators. Furthermore we show that this control scheme can be implemented without reducing significantly the quantum non-demolition character of this type of interferometer.

5.1 Introduction

The presence of arm cavities within the Sagnac speed meter experiment gives rise to challenges not previously encountered in the control of gravitational wave detectors and other experiments involving Michelson or Sagnac interferometers. In section 5.2 we describe the Sagnac speed meter experiment's sensors and actuators and its control requirements. We then describe in section 5.3 a control strategy for the Sagnac speed meter's differential degree of freedom based on that of Michelson designs, and demonstrate the challenges this approach introduces. In section 5.4 we present an alternative strategy which achieves adequate control of the interferometer to reach its design sensitivity over extended periods, including a comprehensive noise budget. Section 5.5 presents a calculation of a matched filter to optimally combine the two error signals to reduce the total quantum noise in post-processing. The parameters used for the control studies are listed in section 5.6 and a summary is provided in section 5.7.

5.2 Control of the proof-of-concept experiment

5.2.1 Degrees of freedom

The arm cavities of the Sagnac speed meter, like those of a Fabry-Perot Michelson interferometer, must be held resonant in order to maintain the light power required for the design sensitivity, and so these cavities represent a degree of freedom that must be controlled with active feedback. Meanwhile, the error signal is insensitive to the motion of the inter-cavity mode matching mirror, M_9 , since this is situated at half the total round trip distance and is sensed by the counter-propagating modes at almost the same time. Other mirrors are potentially significant: the beam splitter M_6 and steering mirror M_7 , as shown in figure 4.9. As these mirrors are situated near the start of one and the end of the other modes' round trips, a velocity dependent signal is created at the balanced homodyne detector (BHD, see section 4.2.1). We neglect all other auxiliary optics.

To assess the importance of the optics to the interferometer's sensitivity to differential arm cavity length $L_{(-)}$, transfer functions from individual mirrors to the BHD port, where the $L_{(-)}$ signal should by design couple most strongly, were calculated using Optickle (see appendix C.1.2). The results in figure 5.1 show that the cavity mirrors are the most important positions to control, with the arm cavity finesse enhancing the sensitivity of the BHD to the arm cavity mirrors such that they dominate the signals from M_6 and M_7 . These results have been confirmed both with Finesse and analytically [141].

The common mode motion of the arm cavities, $L_{(+)}$, will also need to be controlled by means of a photodetector placed at the input port to sense the light returning from the interferometer back towards the laser. This motion will be suppressed by applying strain and heat to the laser's crystal to change its geometry and therefore lasing frequency. This solution involves the creation of a wide bandwidth controller able to provide large corrections within the audio band. While the control of $L_{(+)}$ is crucial to maintain the light power within the arm cavities, we focus on $L_{(-)}$ given that it represents the main signal appearing at the output port and the one which will primarily contribute to the sensitivity of the interferometer in the context of gravitational wave detectors.

While the motion of M_9 can be suppressed at the main readout given suitable mirror positioning in order to cancel the signal from each of the counter-propagating modes, the effect of M_6 and M_7 is less clear cut. To assess the impact the motion from these mirrors has on $L_{(-)}$ sensitivity, a calculation of the effect of seismic noise from M_7 to the BHD can be made. M_6 need not be considered separately here: the transfer function is almost identical to that of M_7 and so we need only calculate one, and the suspension design—a work in progress at the time of writing—is intended to have better isolation than that of M_7 's auxiliary suspension.

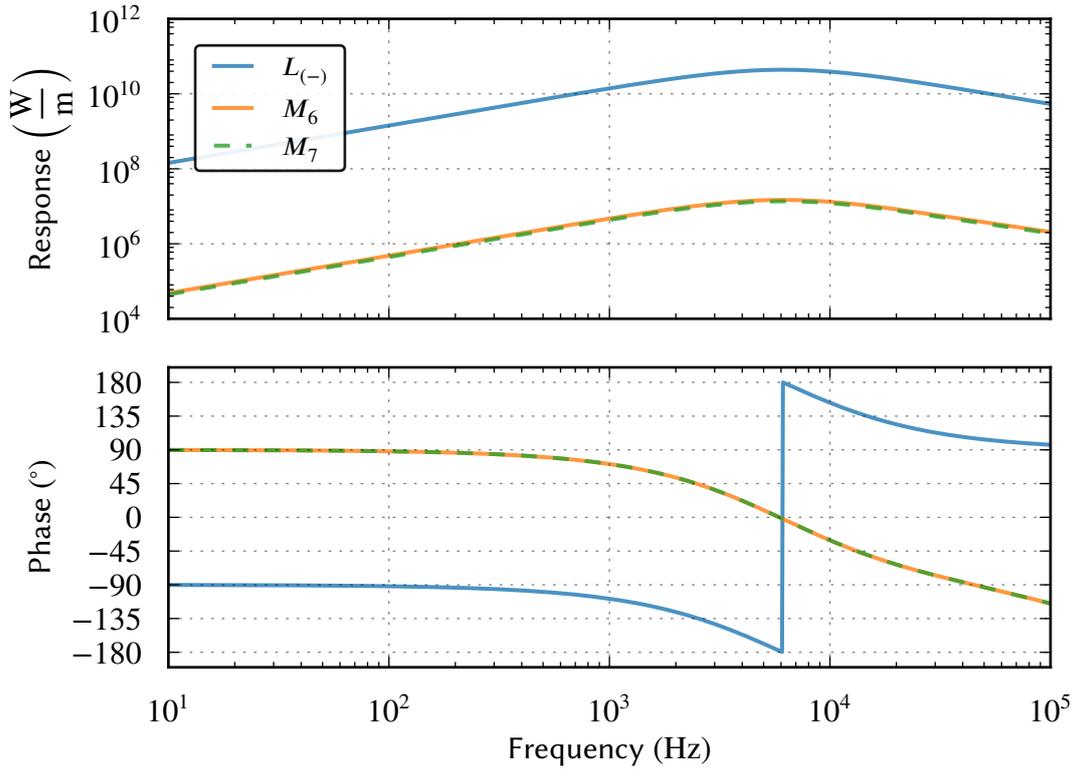


Figure 5.1: Transfer functions from important mirrors or combinations of mirrors in the Sagnac speed meter experiment to the balanced homodyne detector. The $L_{(-)}$ degree of freedom has the strongest response by design. The main beam splitter, M_6 , and the steering mirror for cavity A, M_7 , have response a factor of 10^{-3} that of $L_{(-)}$. Other mirrors have significantly lower coupling.

Measurements of the seismic motion present upon the ground outside the vacuum system can be propagated through a model of the passive seismic isolation within the vacuum system to obtain the effective seismic-induced motion of the tables upon which the suspensions sit. The seismic motion of M_7 can then be calculated by multiplying this spectrum with the transfer function of the auxiliary suspension from the table to the test mass, taken from a state-space model. This seismic noise can be projected into an effective differential arm cavity motion displacement spectral density by multiplying it by the ratio of the transfer functions of M_7 and $L_{(-)}$ to the BHD port¹, taken from figure 5.1. This can be compared with the requirement for sensitivity of the BHD to $L_{(-)}$. Figure 5.2 shows that M_7 's motion, projected into $L_{(-)}$, will meet the requirement above 100 Hz, and the result is similar for M_6 .

The results in figures 5.1 and 5.2 show that control of $L_{(-)}$ will be required to meet the sensitivity requirement at the BHD port above 100 Hz, where the measurement of reduced

¹This is the same as multiplying the motion of M_7 by its transfer function to the BHD port to yield a signal in $W \sqrt{\text{Hz}}^{-1}$, and dividing by the transfer function from $L_{(-)}$ to the BHD to yield an effective motion in terms of $L_{(-)}$.

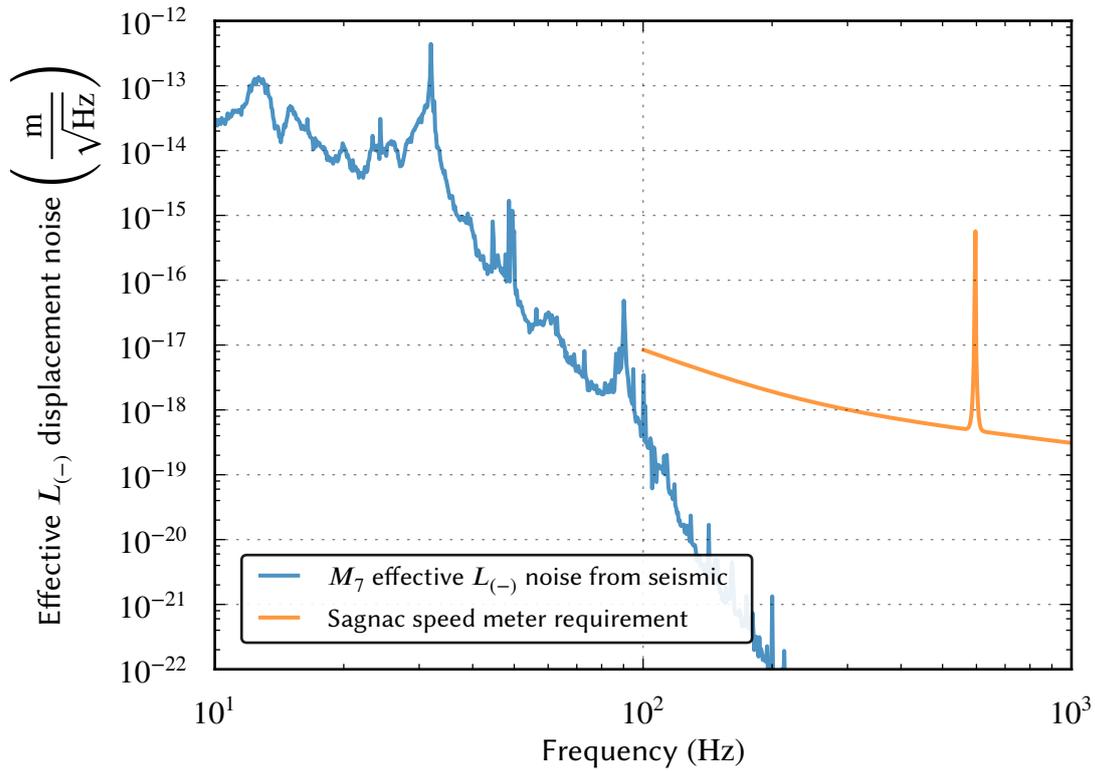


Figure 5.2: Effective $L_{(-)}$ seismic noise contribution from M_7 . This is calculated by first propagating a seismic noise spectral density for the laboratory near the vacuum system through damping and suspension models to obtain the motion of the M_7 test mass. With this figure, the response at the BHD can be calculated from the transfer function shown in figure 5.1, and this in turn can be expressed in units of differential arm cavity motion by dividing it by the response of $L_{(-)}$ to the BHD port. The requirement is given only for frequencies above 100 Hz where the measurement of reduced radiation pressure noise will be made, and this figure shows that seismic motion of M_7 will not represent a significant problem to the sensitivity of the experiment in the desired band. This conclusion applies also to the main beam splitter, M_6 , which is expected to have even greater isolation from seismic noise.

radiation pressure noise will be made. It should be noted, however, that the desired BHD homodyne angle depends on the relative path lengths of M_{11} to M_{16} and M_6 to M_{16} . This length will be controlled by an auxiliary loop not considered part of the longitudinal control strategy, and will be the subject of future work alongside a strategy for the control of $L_{(+)}$.

5.2.2 Sideband frequency

The eventual choice of sideband frequency, used to control cavity lengths using techniques such as Pound-Drever-Hall (PDH, see section 3.3.4.1), will depend on a number of factors both physical and technical. For the purpose of control simulations, however, the only requirement is that the sideband frequency is not resonant within the arm cavities, in order to act as a discriminant to allow for the control of the arm cavity lengths; this is described

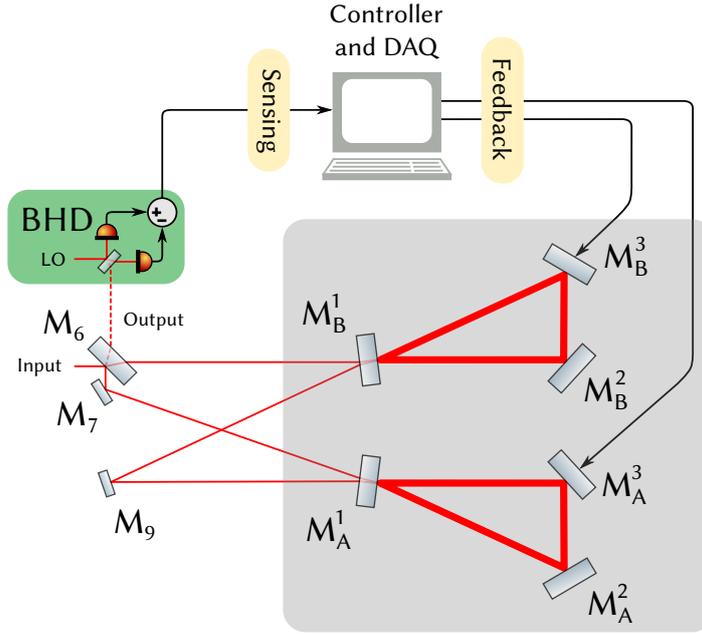


Figure 5.3: Simplified Sagnac speed meter layout including extraction of the BHD signal sensitive to the arm cavity differential mode, $L_{(-)}$, and the sensing and feedback paths. Light from the input optics (not shown) is incident upon the main beam splitter, M_6 . The triangular arm cavities are shown in the shaded grey area, and mirror M_9 couples light between them. The shaded green area shows the BHD extracting the signal from the main beam splitter’s output port (see section 5.2.5). The sensing and feedback signal paths are described in detail in section 5.2.8.

in more detail in section 7.4.2. In practice, this means the frequency offset from the carrier must be greater than the cavity’s FWHM (see appendix A.2.1). For control simulations the sideband frequency was chosen to be 15 MHz.

5.2.3 Control considerations

Figure 5.3 shows a simplified optical layout of the Sagnac speed meter experiment with the addition of a basic control loop. The main beam splitter (M_6) splits the input field towards the two triangular arm cavities where they form counter-propagating modes. One mode from each arm cavity is coupled into the other via the inter-cavity mirror M_9 , and the other modes recombine at the main beam splitter. Here, and for the rest of this chapter, we will consider only the cavity mirrors, the beam splitter and M_9 , defined as shown in figure 5.3.

As shown in section 5.2.1, frequency-dependent changes in $L_{(-)}$ lead to frequency-dependent signals at the BHD. Motion of an arm cavity mirror imparts signal sidebands upon the counter-propagating modes; these modes have different optical path lengths to the beam splitter and so the signal at the output port contains the superposition of signals representing the mirror’s displacement from different points in time, which is analogous to velocity. At dc the two modes at the output port contain the same displacement information and the

velocity signal is therefore zero².

The readout representing $L_{(-)}$ is sensed at the main beam splitter's output port by means of the BHD, as shown in the shaded green area in figure 5.3. The frequency dependence of the phase quadrature signal at the BHD s_{BHD} is given by the following relationship, ignoring the effect of losses³:

$$s_{\text{BHD}}(\Omega) \propto \frac{\Omega}{(\Omega^2 + \gamma_{\text{arm}}^2)} L_{(-)}, \quad (5.1)$$

for angular frequency Ω and with arm cavity half-bandwidth γ_{arm} defined to be:

$$\gamma_{\text{arm}} = \frac{c_0 T_{\text{ITM}}}{4L_{\text{RT}}}, \quad (5.2)$$

for speed of light c_0 , arm cavity input test mass (ITM) power transmissivity T_{ITM} and arm cavity round-trip length L_{RT} .

Other terms in the response function dependent upon mirror mass, laser power and mechanical modes are not frequency dependent. Note that for $\Omega \ll \gamma_{\text{arm}}$ the response is proportional to frequency, vanishing towards dc, as described above and shown in figure 5.4.

In order to maintain peak BHD sensitivity to $L_{(-)}$ and therefore gravitational waves, the positions of the cavity mirrors are controlled using *linear negative feedback*, where an error signal is extracted and applied through a control law to cavity mirror actuators. In the Sagnac speed meter experiment, voice coils and plate-capacitor electrostatic drives are used to actuate on the positions of the end test masses (ETMs) within each cavity. This feedback maintains the interferometer close to its operating point within the bandwidth of the controller.

5.2.4 Required controller precision

In order to achieve the required stability, the relative position of the cavity mirrors must be held at the dark fringe, as discussed in section 2.1.2. The noise present within the interferometer, however, produces an unintended *dark-fringe offset* at the output port. The dark fringe condition is strictly only met when there is no interferometer noise, though the slope of the fringe near the minimum is shallow within about 1% of the fringe's full width at half maximum (FWHM, see appendix A.2.1). A thorough analysis of the required control precision has been derived for the case of a dual-recycled Fabry-Perot Michelson interferometer with dc readout [143] but to the author's knowledge this has not been repeated for a Sagnac speed meter with BHD readout. As most technical noise sources in the Sagnac

²For a more complete description of the Sagnac speed meter's behaviour, see, for example, ref. [87, section IIb].

³A comprehensive treatment of the effect of loss is given in ref. [89].

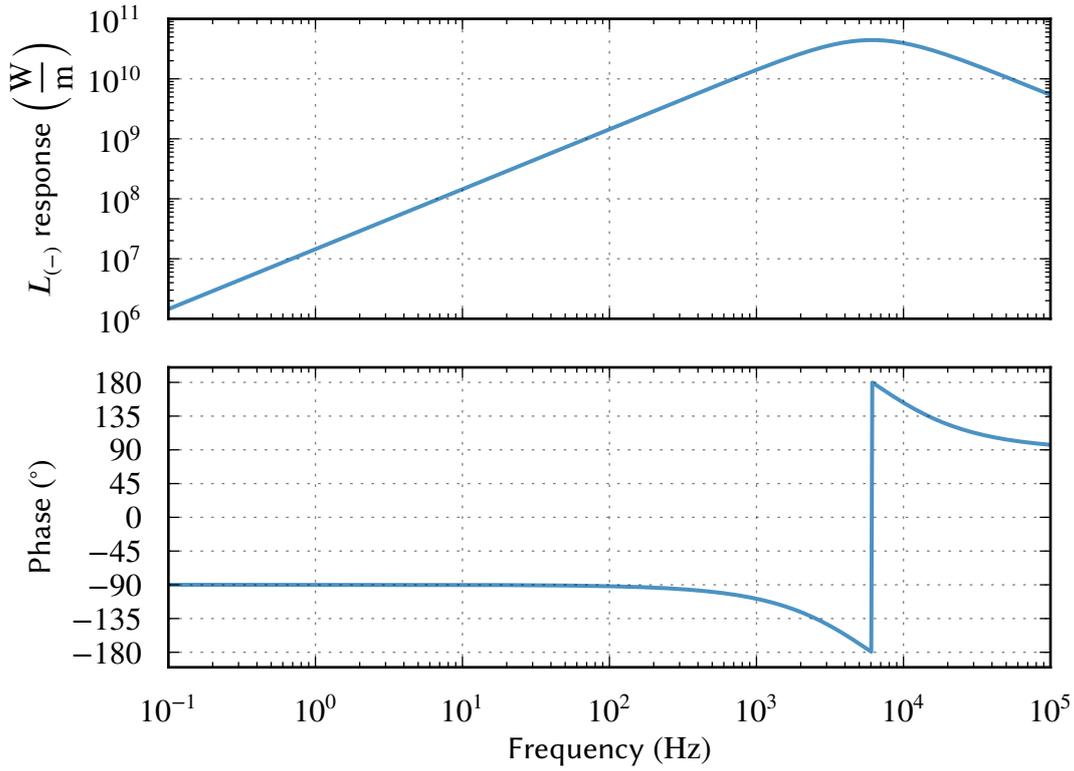


Figure 5.4: The frequency response of the Sagnac speed meter experiment's $L_{(-)}$ degree of freedom to the BHD, simulated numerically with Optickle. As the BHD is sensitive to the arm cavity mirrors' velocity, the signal is proportional to frequency below the cavity pole, and thus zero at dc.

speed meter have similar output port couplings to that of a Michelson interferometer, the requirement is expected to be similar. Assuming that the frequency equivalent fluctuations Δf fall within 1% of the arm cavity FWHM, we can derive a requirement to ensure that technical noise sources do not couple strongly to the gravitational wave channel.

Using the parameters listed in table 5.2 with the relation introduced in equation (1.6) linking laser carrier frequency fluctuations Δf and cavity length fluctuations $\Delta L_{(-)}$,

$$\frac{\Delta L_{(-)}}{L_{RT}} = \frac{\Delta f}{f_0}, \quad (5.3)$$

with f_0 representing the carrier frequency, the requirement for the Sagnac speed meter experiment is that the root-mean-square (rms) displacement of the mirrors due to noise must be less than 3.5×10^{-13} m.

As shown in ref. [89], asymmetries in the main beam splitter introduce common arm cavity mode coupling at the output port, which leads to further unintended dark fringe offset, and so the real requirement is likely to be more stringent. The controller should therefore have a reasonable factor of safety in terms of the gain it is able to apply to the system to hold it

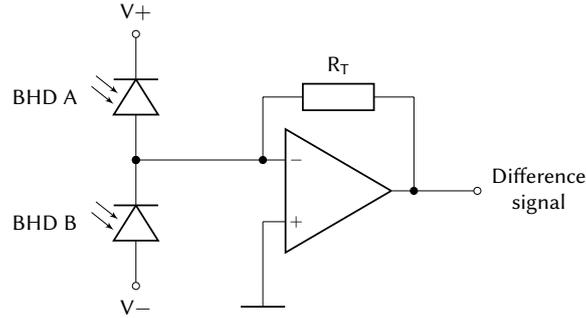


Figure 5.5: Simplified electronic schematic for the BHD readout within the Sagnac speed meter experiment. The difference current from two matched, high quantum efficiency photodiodes is amplified via a transimpedance op-amp stage, with this signal representing the differential motion of the arm cavity mirrors (see equation (5.1)).

at the operating point.

5.2.5 Balanced homodyne detection

The BHD at the Sagnac speed meter experiment's output port consists of two high quantum-efficiency photodiodes sensing the reflected and transmitted fields from the BHD's beam splitter, M_{16} . A local oscillator field is incident upon the BHD's beam splitter. The difference current is converted to a voltage by an op-amp with transimpedance resistor R_T before being sent to CDS.

An example circuit for the balanced homodyne detector is shown in figure 5.5. The op-amp introduces its own noise to the output, though a well-chosen op-amp will possess noise significantly lower than the signal representing $L_{(-)}$ in the intended measurement band. In order for an op-amp to contribute less than 1% of the uncorrelated noise in the measurement, its noise must be at least a factor of 10 below the dominating noise source in the measurement band⁴.

Op-amps used for control in audio-band interferometry typically possess a noise power spectrum inversely proportional to frequency (so-called *flicker noise* [144, section 11.2.3]) in the low audio band. As the BHD error signal is dependent upon the time derivative of the mirror positions, however, there will necessarily be frequencies at which the op-amp noise will dominate the BHD error signal. This makes control of slow drifts of the arm cavity mirror positions impossible with the velocity readout, despite the op-amp being well-chosen for a measurement band above 100 Hz. This control problem with relation to the Sagnac speed meter experiment will be described in more detail in section 5.3.

⁴As the uncorrelated noise sources are added in quadrature, a noise source a factor $\frac{1}{10}$ that of another will contribute less than $\left(\frac{1}{10}\right)^2 = 1\%$ to the overall noise.

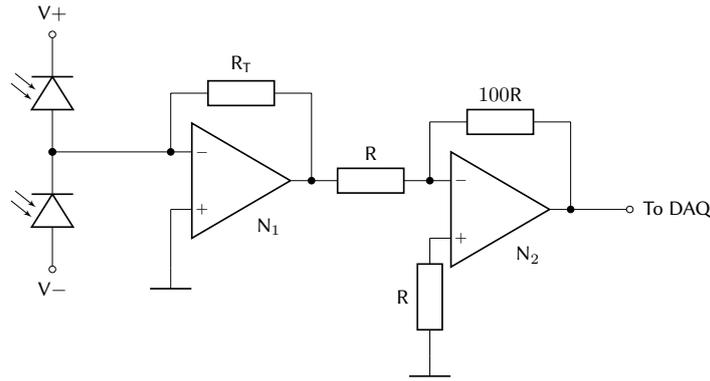


Figure 5.6: Electronic schematic for the measurement of noise from the BHD readout circuit shown in figure 5.5. The output from the transimpedance op-amp is multiplied by a factor of 100 by an identical op-amp. The level of multiplication was chosen to allow this noise to exceed the noise of the analogue-to-digital converters within CDS.

5.2.6 Op-amp noise

To measure the effect of a suitable op-amp's noise at low frequencies, the output from an applicable BHD circuit was investigated. The circuit shown in figure 5.6 was housed within a dark enclosure to minimise photocurrent, with one of the two op-amps within a Texas Instruments® OPA2227 integrated circuit being used to amplify the noise from the other by a factor of 100. This amplification step is necessary in order to allow the desired op-amp noise to be measured above the noise of the CDS data acquisition system's analogue-to-digital converters (ADCs) used to record the data. The OPA2227 op-amp is a low-noise precision amplifier designed for audio applications, and is thus suitable for the BHD readout in the Sagnac speed meter experiment depicted in figure 5.5 given the intended measurement band. The transimpedance resistor was set to $10\text{ k}\Omega$ to balance the first op-amp's contributions to its output from its input current and voltage noise.

The circuit's output noise was recorded for a period of 16 days alongside an open CDS input channel used to quantify some of the measurement noise. A time series of the data is shown in figure 5.7, where a drift over the course of the measurement period is apparent. An amplitude spectral density estimate of the measured op-amp noise time series (figure 5.8) shows a combination of flicker noise and an additional slope possibly due to resistor current noise below around 1 Hz [145]. ADC noise dominates above 4 Hz. The "Model (total)" spectral density in figure 5.8 shows the contributions to the measurements from the first op-amp N_1 's current and voltage noise and the Johnson-Nyquist noise of its transimpedance resistor R_T . This spectral density additionally contains the measured open channel noise summed in quadrature to show the agreement it has with the measurements down to around 1 Hz.

Since the signal measured at the BHD in the Sagnac speed meter experiment represents cavity mirror velocity, it must necessarily drop below the noise at low frequencies where the velocity tends to zero. The op-amp N_1 's noise drift produces an offset upon the BHD

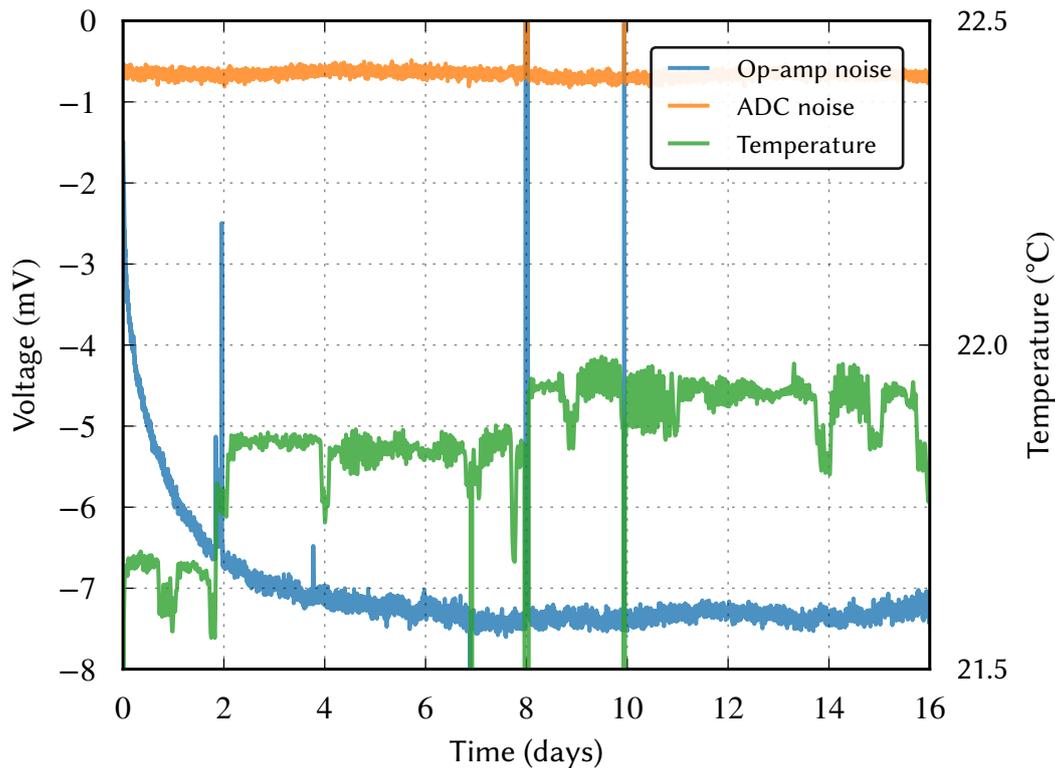


Figure 5.7: Time series of the output from the BHD noise measurement circuit shown in figure 5.6. The op-amp noise (blue) drifts from 0 to a level of approximately -7.3 mV over the 16 day measurement. Simultaneously, a CDS channel was measured without an input connected (orange) such that the input impedance was that of the channel’s line receiver, $34\text{ M}\Omega$, in order to act as a null stream. The temperature was also measured by a sensor within the same housing as the noise circuit (green), showing a drift of around $0.5\text{ }^{\circ}\text{C}$.

error signal which is to be fed back to the cavity mirror actuators, and thus op-amp noise contributes to cavity mirror displacement noise, affecting the experiment’s sensitivity.

5.2.7 Technical noise

5.2.7.1 Analogue to digital converters

As the CDS system runs its control system in software, the analogue signals sensed from the interferometer must be converted to digital form with ADCs suffering from quantisation noise as introduced in section 2.2.5.3. The feedback signals generated by the control system must similarly be converted from digital to analogue form with DACs.

In the case of the 16-bit ADCs in CDS, the effective number of bits⁵ (ENOB) is $b = 13.9$,

⁵One might decide to purchase an ADC based merely upon its number of bits, but this is not a good guide for determining its sensitivity. A 24-bit ADC is no better than a noise-free 16-bit ADC if the first 8 bits are noise. A more useful figure of merit is the ENOB.

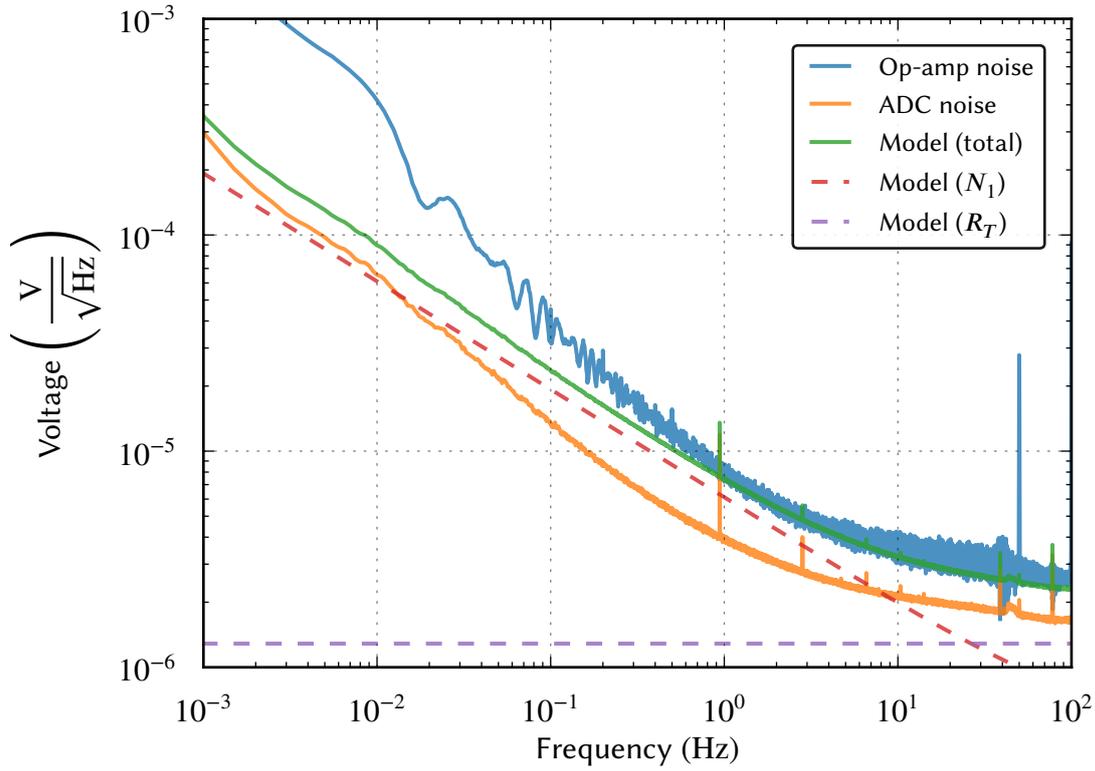


Figure 5.8: Amplitude spectral density of the noise measured from the BHD readout electronics. The op-amp noise spectrum (blue) shows the amplitude spectral density estimate (see appendix B.2.3) for the data in figure 5.7. The ADC noise spectral density is also given (orange) along with modelled op-amp and resistor noise sources projected into the same measurement point (red and purple, respectively). The most significant contribution to the output is from the first op-amp, as intended; though the noise model, which accounts for the op-amp’s input voltage and current noise and the Johnson-Nyquist noise of the resistors, departs from the measurements at lower frequencies.

corresponding to a noise level of $1.8 \times 10^{-6} \text{ V } \sqrt{\text{Hz}^{-1}}$ using the relation [50]

$$\epsilon_{\text{ADC}} = \frac{V_{\text{range}}}{2^b \sqrt{12 f_N}}, \quad (5.4)$$

where f_N is the Nyquist frequency, which, in the case of CDS is 32 768 Hz. This noise floor is flat across much of the bandwidth of CDS as it is determined by op-amps chosen for low noise in the audio band.

With the CDS system, the ADCs and DACs are well matched and possess the same noise floor.

5.2.8 Sensor and actuator dynamic range considerations

5.2.8.1 Whitening and dewatering filters

The signals sensed by the Sagnac speed meter experiment's photodetectors contain large components at low frequencies arising from seismic noise, and small components at higher frequencies where the measurement of radiation pressure can be made. The ADC's input range of ± 10 V and its noise $1.8 \times 10^{-6} \text{ V } \sqrt{\text{Hz}^{-1}}$ lead to a dynamic range D of

$$\begin{aligned} D &= 20 \log_{10} \left(\frac{V_{\max}}{\epsilon_{\text{ADC}}} \right) \\ &= 134.9 \text{ dB.} \end{aligned} \tag{5.5}$$

As the $L_{(-)}$ sensitivity of the BHD is shaped inversely to frequency, signals in the kHz band are typically many orders of magnitude smaller than those of seismic noise at a few Hz meaning that signals at the level of quantum noise are often smaller than the noise of the ADCs and DACs if the gain is chosen to avoid saturating the sensor at low frequencies. This makes it difficult to sense low and high frequency signals simultaneously. To avoid this problem, a technique called *whitening* can be used. This involves the application of a filter to the desired input or output signal in order to increase or decrease certain frequency components of a signal; the intention is to make the signal strength at all frequencies the same, i.e. *white*.

The effect of whitening is shown in figure 5.9 for a hypothetical signal and sensor. In the unwhitened case, the signal's dynamic range is greater than that of the sensor, and so the sensor is unable to measure the signal above around 100 Hz with fidelity. At higher frequencies, the signal is below the level of the sensor's noise and as such the sensor cannot distinguish the signal. By applying a whitening filter, the small signal content there can be enhanced to a level at which it can be detected by the sensor above its noise. This effectively reduces the dynamic range of the signal, and this is shown in the whitened case. The original signal can be recovered digitally by the application of an inverse whitening filter ("dewatering").

The use of whitening can lead to extra noise at higher frequencies, where the continually decaying signal becomes lower than the electronic noise of any real whitening filter arising from the op-amps and resistors. Care must be taken in the design of whitening filters, with component values chosen fit for the intended measurement band.

The whitening to be applied to the Sagnac speed meter experiment's sensor and actuator signals is shown in figure 5.10. These filters are sufficient to meet the requirement that the ADC and DAC noise contribute less than 1% of the noise power.

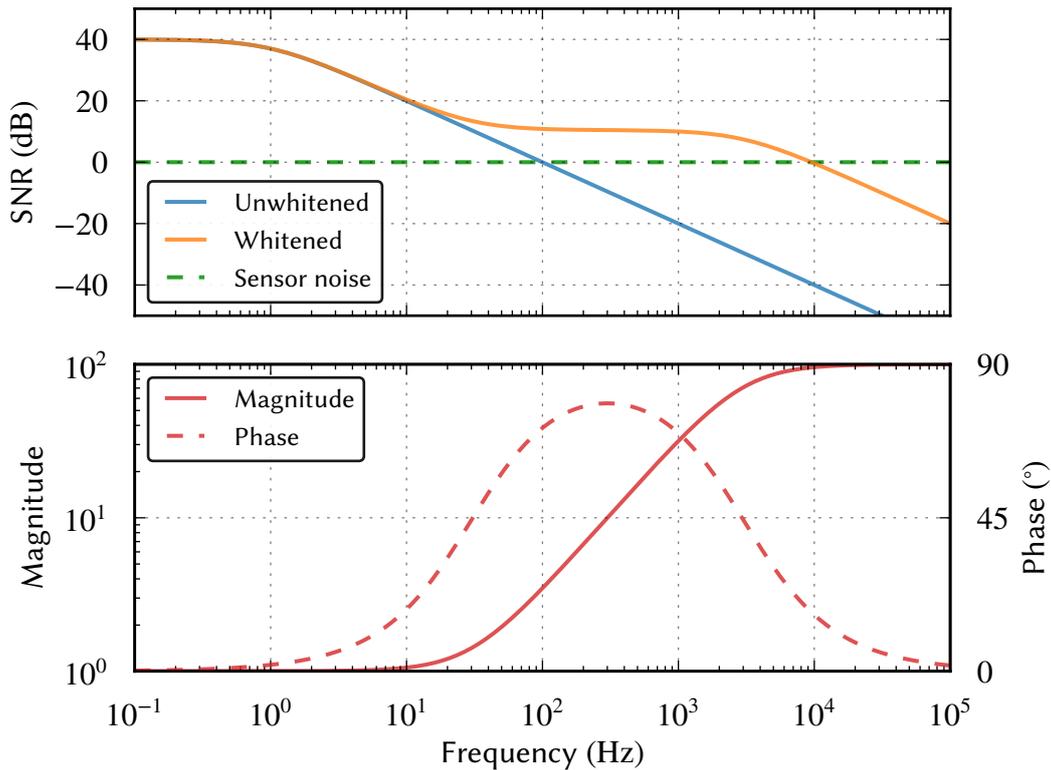


Figure 5.9: The effect of whitening on a signal. The unwhitened signal (blue, upper plot) drops below the sensor’s noise above 100 Hz, and so it will output only the noise of the sensor (at SNR = 0) at higher frequencies. A whitening filter, shown in the lower plot, has been applied to the same data to produce the whitened curve (orange, upper plot). In this case the whitening filter is a high pass filter, but in general its dynamics should be determined by the shape of the underlying signal to be whitened. The whitening filter increases the signal’s magnitude at higher frequencies and thus makes it detectable above the sensor’s noise up to a higher frequency, in this case 10 kHz. Once the whitened signal has been detected the underlying, unwhitened signal can be recovered in the digital domain through the application of the corresponding inverse whitening filter.

5.2.8.2 Aliasing and imaging

A consequence of the Nyquist-Shannon sampling theorem is that ac signal content can be exactly reconstructed by a sampler if and only if the signal power above the Nyquist frequency is zero [146]. Non-zero signal at frequencies $f > f_N$ will enter the band of the sampler every $\frac{f}{f_N}$ cycles and appear on top of the real signal content in that band. To prevent this occurrence, *anti-aliasing* filters can be utilised to aggressively suppress higher frequency content using analogue electronics before the signal is sampled by the ADC. Similarly, the output from a DAC can be propagated through an *anti-imaging* filter to prevent the DAC’s finite sample rate from creating higher frequency copies of in-band signal content.

With CDS, the sample frequency is 65 536 Hz, and so the Nyquist frequency is 32 768 Hz,

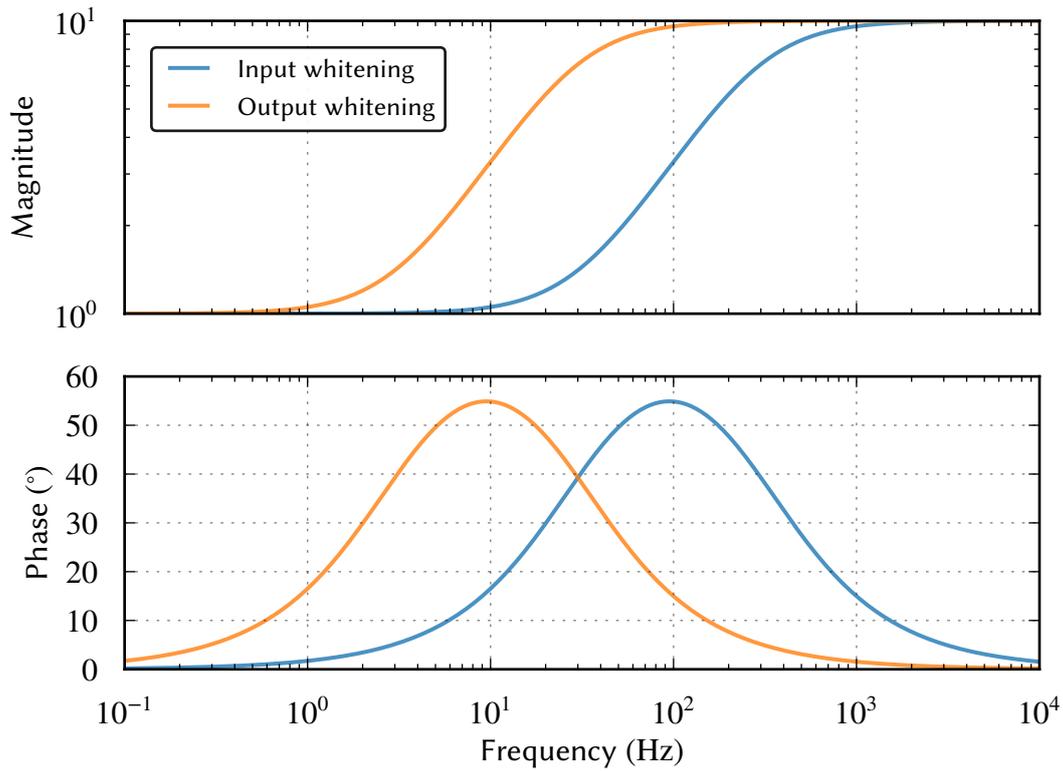


Figure 5.10: Input and output whitening filters in the Sagnac speed meter experiment. The input whitening will be implemented in the analogue domain while the output whitening will be implemented digitally in CDS. The equivalent dewhitening filters will be implemented in the digital and analogue domains, respectively.

and the anti-aliasing and anti-imaging filters have cut-off frequency 9 kHz to ensure that frequency content near the Nyquist frequency is practically zero. The filter is implemented as a 3rd order low-pass Butterworth, giving the flattest response in the band up to 9 kHz. In addition, a notch filter is implemented at the sample frequency to suppress pick-up from the laboratory: ultimately, the sample rate is generated by an oscillator which may produce electromagnetic radiation at nearby frequencies. Figure 5.11 shows the (identical) transfer functions of the anti-aliasing and anti-imaging filters implemented in CDS. At the Nyquist frequency, the signal is suppressed by approximately 10^2 and at the sample frequency it is suppressed by 10^5 .

5.2.8.3 Suspension gain hierarchy

As discussed in section 4.2.3 the suspension for each ETM will contain two actuator types: voice coils for control of the test mass motion at low frequencies where seismic noise is significant, and an ESD for control at high frequencies to correct small but fast perturbations. The feedback signal generated by the controller within the loop is a signal with frequency

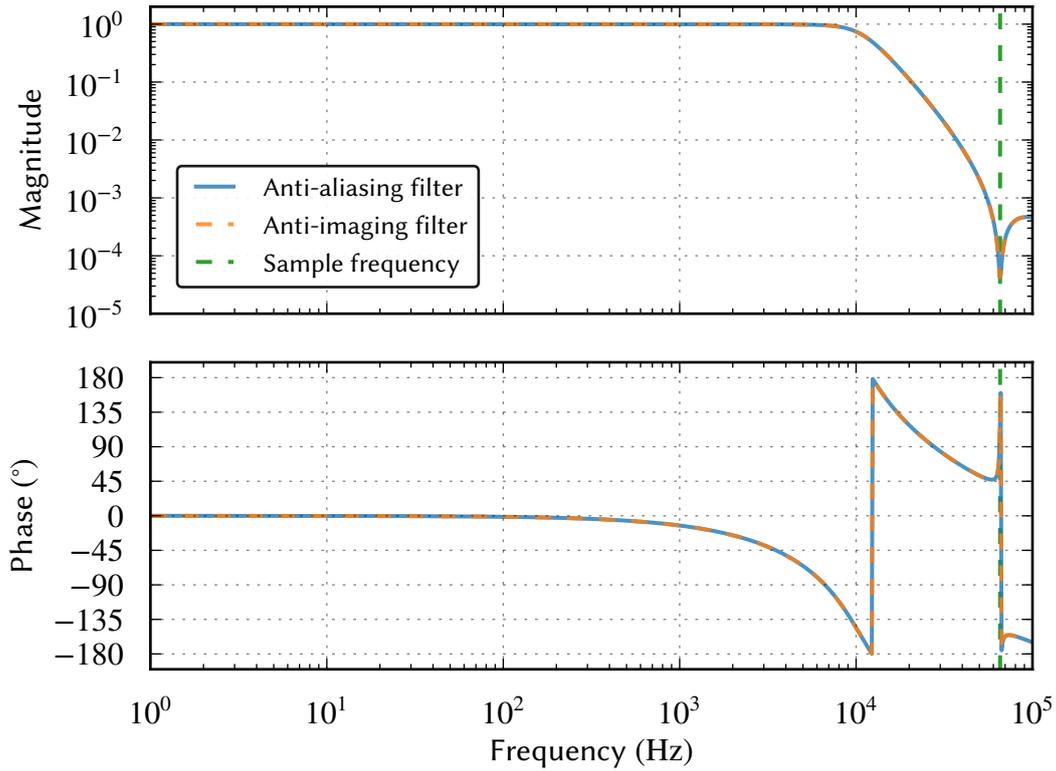


Figure 5.11: Transfer functions of the anti-aliasing and anti-imaging filters implemented in CDS. The gain at low frequencies is unity to allow signals to pass unperturbed. From 9 kHz, the filters suppress higher frequency signal content to prevent aliasing or imaging into lower frequencies. At the sample frequency, a notch filter removes all but a factor of 10^{-5} of the signal to prevent pick-up.

components at low and high frequencies, and a set of filters is required to split this feedback between the voice coils and the ESDs. This technique is termed *gain hierarchy* and it has been applied for example in the control of Advanced LIGO’s quadruple suspension systems [147].

Actuators are a potential source of noise in a control loop. Voice coils are susceptible to noise coupling via stray magnetic fields and Barkhausen noise. The ESD, however, is anticipated to have excellent noise performance (see chapter 6), so it would ideally be used for test mass positional corrections across the entire control bandwidth; however, its maximum voltage supply, and therefore force output, is very limited and is nowhere near capable of controlling the test masses due to seismic noise at low frequencies. The rms motion of the uncontrolled ETMs is expected to be of the order 10^{-6} m due to noise sources such as seismic, thermal, electronic and quantum, while the maximum force output of the ESD will be around $1.5 \mu\text{N}$, corresponding to a displacement of around 690 nm at 1 Hz, or just over half a wavelength. The response of the ESD at high frequencies is what would be expected of a force to displacement coupling on a free mass, proportional to $\frac{1}{f^2}$. The voice coils’ response at high frequencies, however, contains both a $\frac{1}{f^2}$ slope from force to displacement

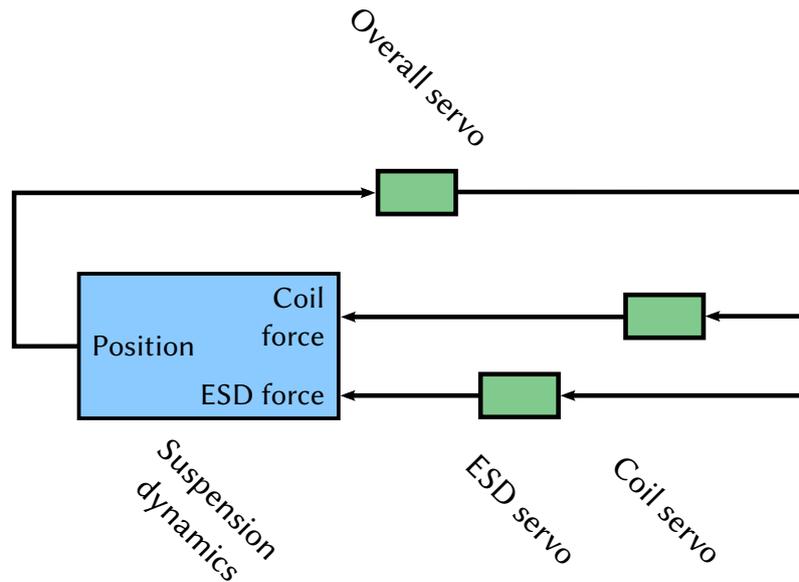


Figure 5.12: ETM suspension actuator control loop block diagram. The output of the state-space model representing the test mass motion is fed back to the actuators via an “overall” servo. This servo can in principle represent the interferometer’s response, since the test mass motion from the suspension will be altered by the interferometer before being fed back to the actuators, though in the development of the gain hierarchy the interferometer’s response is assumed to be unity. The real effects can be compensated for in the controller (see section 5.2.8.4).

of the intermediate stage, as well as another $\frac{1}{f^2}$ term from the pendulum stage between the intermediate mass and the test mass, giving an overall filtering effect at high frequencies proportional to $\frac{1}{f^4}$. Given these constraints the effort of the ESD is focused at high frequencies where its actuation is stronger than that of the voice coils, while the voice coils are utilised at low frequencies where their vastly increased range is available to correct for the larger expected noise disturbances.

In order to create the gain hierarchy a series of filters were implemented around a state-space model of the ETM suspension. This model includes the response of the actuators from force to displacement for each degree of freedom of the suspension. The use of filters on the input path to each actuator allows us to split the feedback signal into low- and high-frequency corrections, while an overall filter allows common corrections to be applied to both actuators. The control loop built to configure the suspension gain hierarchy is shown in figure 5.12.

From the control precision requirement presented in section 5.2.4 it can be shown that the rms motion as a function of frequency drops below this requirement around 100 Hz, meaning that the interferometer’s controller must at least feed back signals up to this frequency. To give the controller some extra headroom to control particularly large noise transients—expected due to the stationary random noise present within the system—the unity gain frequency should instead be set at a higher frequency.

While the shaping of the hierarchical gain is an iterative process that requires some trial and error, the following paragraphs attempt to explain the methodology behind the features of the Sagnac speed meter experiment's implementation for the ETMs.

Stable unity gain frequency Primarily due to the dynamic range of sensor and actuators, the controller has finite bandwidth and cannot feed back signals at an infinite number of frequencies. The control bandwidth was set to 350 Hz to give a reasonable factor of safety over the 100 Hz requirement. This means the unity gain frequency is 350 Hz. As the ESD will be providing most of the feedback at this frequency, we must ensure that its slope is proportional to $\frac{1}{f}$ at 350 Hz to facilitate a stable unity gain crossing (see appendix B.4.3). As the force-to-displacement response is proportional to $\frac{1}{f^2}$ above the pendulum resonance, we want to add an f response to make this $\frac{1}{f}$, and so we use a transitional differentiator between 30 Hz and 1 kHz.

Stable actuator crossover frequency A frequency at which the magnitude of feedback from the voice coils and ESD is equal is called a *crossover* frequency. This point has the same requirement as the unity gain frequency, in that the feedback must not be 180° out of phase with the input. To facilitate a stable crossover at around 18 Hz—chosen to prevent the ESD's range from being consumed by corrections to low frequency noise—a transitional differentiator between 2 Hz and 50 Hz was added to the voice coil servo. As the voice coil's high frequency response is proportional to $\frac{1}{f^4}$, this results in a response of $\frac{1}{f^3}$ to compliment the ESD's response of $\frac{1}{f^2}$ —a difference of 90°.

Increasing the voice coil feedback at the pendulum resonances At 1.8 Hz the final pendulum stage is resonant and so the ground motion is amplified on the test mass. To avoid saturating the control electronics with signal at this frequency, an additional boost was provided to the voice coil feedback through a 2nd order resonant gain filter with pole and zero at the resonant frequency and a quality factor of around 3 to allow for slight changes in the resonant frequency due to temperature drift.

Control of pitch coupling At 10.2 Hz, a coupling between pitch and longitudinal modes of the suspension's final stage pendulum leads to a suspension resonance. To prevent this mode from ringing, a 2nd order resonant gain filter was applied at 10.2 Hz with a quality factor of 4 to allow for manufacturing tolerance. As the voice coil and ESD feedback is of similar magnitude at this frequency, this filter was placed within the common feedback path.

Notching of violin modes Violin modes (see section 2.2.2.2) are present on the ETM suspensions starting at 800 Hz. Although this frequency is well above the control bandwidth, the modes have high enough quality factor and amplitude and are resonant peaks with a 180° phase change such that they can potentially lead to positive (unstable) feedback. Instead of damping these modes with resonant gain, we avoid feedback at this frequency by applying a 2nd order notch filter. For illustration we've added damping for the first violin mode. If the higher order violin modes become a problem for control we can add additional filters as necessary.

Implementation considerations To provide an equal number of poles and zeros in the ESD servo, we add an additional transitional differentiator between 0.05 Hz and 5 Hz. This has little effect on the response as the ESD's gain is much smaller than that of the voice coils in this range, but it simplifies the digital implementation of the servo. Ideally, the differentiator's corner frequency would be at dc, but this is not possible in a digital implementation.

If low frequency seismic noise must be suppressed further, an extra boost can be applied to the voice coil servo with the inclusion of a transitional integrator between for example 0.01 Hz and 1 Hz, though this has not been considered in the analysis.

The open-loop transfer functions for the voice coils and ESD, normalised to the voice coil's dc response, are shown in figure 5.13. This shows the gain of the ESD with respect to the voice coils as a function of frequency, resembling the performance of the system in operation. The coil transfer function was calculated by breaking the loop between the coil and ESD servos and the suspension dynamics block, injecting a signal at the suspension dynamics block input corresponding to the coil servo and then reading it out at the coil servo output. The transfer function is the ratio of the readout and injected signals. The ESD transfer function was calculated in a similar way. This process is depicted in figure 5.14.

5.2.8.4 Interferometer compensation

The suspension gain hierarchy in section 5.2.8.3 was developed by feeding the test mass motion directly back to the suspension actuators, as shown in figure 5.12. In reality, the test mass motion affects the interferometer which changes the signal on the BHD, and so the frequency components of the signal are altered. In order for the suspension gain hierarchy to operate as designed, a filter must be placed in the controller to compensate for the interferometer's response. This was implemented in CDS as transitional integrators: the first between 0.1 Hz and 2 kHz and the second between 1 Hz and 100 Hz. While these filters do not reproduce the exact transfer of $L_{(-)}$ displacement to the BHD readout, it approximates it closely enough to allow the suspension gain hierarchy to operate as per its design.

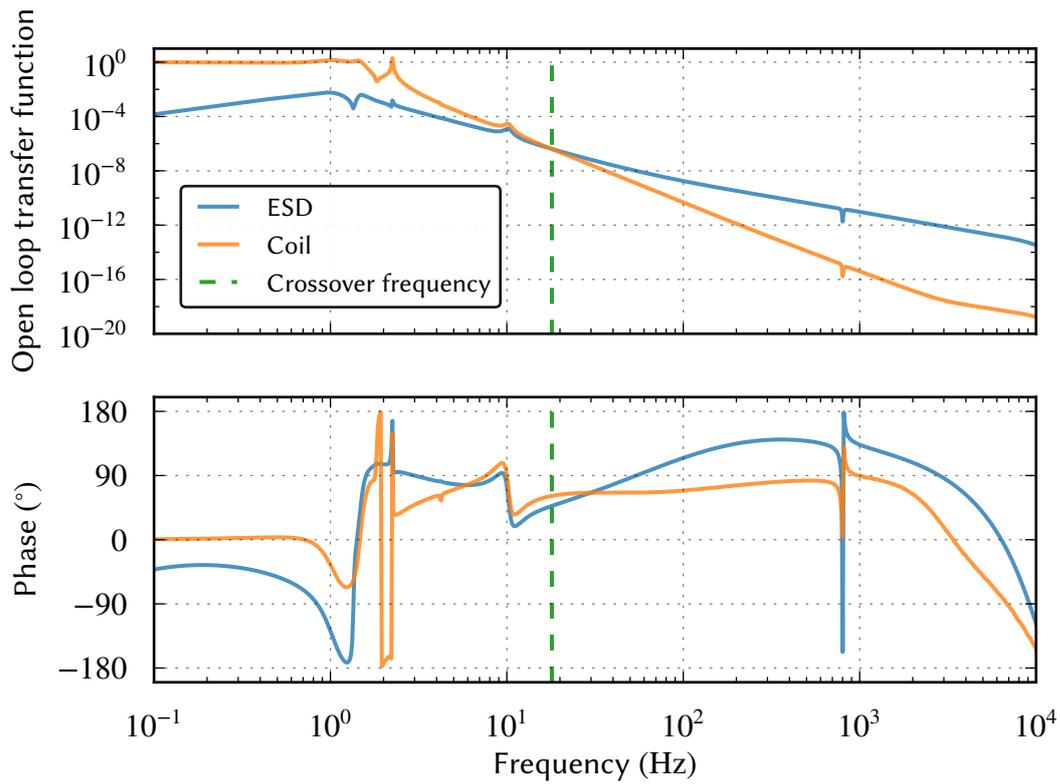


Figure 5.13: Simulated Sagnac speed meter experiment ETM suspension actuator open-loop transfer functions showing the difference in gain between the voice coils and ESD. The voice coils provide extensive actuation range but are suppressed at high frequencies by the final stage pendulum. The ESD actuates directly upon the test mass and is therefore capable of providing stronger correction than the voice coils at higher frequencies. A 2nd order notch filter is present on both actuators at 800 Hz to prevent excitation of the first suspension violin mode.

During commissioning it may be useful to adjust the shape of this compensation to better fit the real interferometer.

5.2.8.5 Photodiode quantum efficiency

A photodiode's *quantum efficiency* relates to how well it converts incoming photons into electrons. A real photodiode cannot fully convert incoming light power into photocurrent without some loss.

The number of photons corresponding to a given light power is governed by the wavelength, and a photodiode's efficiency changes with the wavelength. To calculate the photocurrent from a photodiode for a given light power at a given wavelength, we use its *responsivity*. For 1064 nm laser light there exists some high quantum efficiency models providing responsivity upwards of 0.9 A W^{-1} , which is the figure we assume for the photodiodes of the BHD.

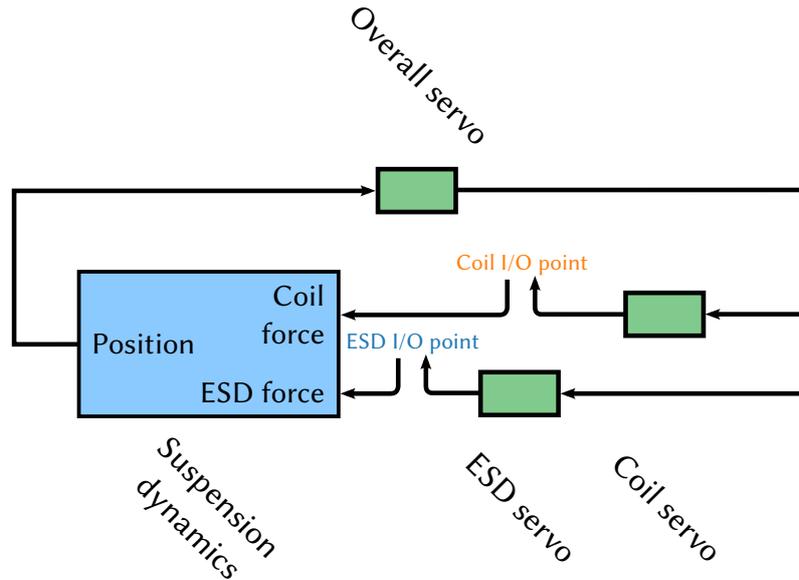


Figure 5.14: ETM suspension actuator control loop block diagram showing injection points. To calculate the open loop transfer functions, both the coil and ESD feedback loops are severed between the servos and the suspension dynamics block. The coil transfer function is calculated by injecting a signal at the suspension dynamics block input corresponding to the coil and reading out the resulting signal at the output of the coil servo, having propagated through the loop. The ESD transfer function is calculated in the same way but using the ESD servo and corresponding suspension dynamics block input as readout and injection points. The transfer functions are shown in figure 5.13.

5.2.8.6 Loop gain

With the response of $L_{(-)}$ to the BHD readout calculated by Optickle and implemented in the control loop alongside the suspension feedback and interferometer compensation filters, the strength of the control loop's suppression of displacement noise is determined by the *loop gain*. This is a dimensionless number determined by the response of each component within the loop into its connected components. It can be increased manually through the use of a dc gain stage placed anywhere within the loop, or for instance by utilising a photodetector with higher quantum efficiency or by using stronger actuators. It is in practice easiest to place a manual, "overall" gain stage within the controller, in this case the CDS system where gain is "free" within the limits set by numerical precision. An increase in the loop's dc gain leads to a higher open loop unity gain frequency and due to the shape of the suspensions' hierarchical gain this means the feedback at lower frequencies is stronger, with seismic noise being more aggressively suppressed. The residual displacement of the test masses decreases with higher loop gain, leading to better control, until the rms range of an actuator or sensor is reached. A rule of thumb with the operation of control loops within the field of gravitational wave interferometry is to design the overall gain servo shape such that the phase margin at the unity gain frequency is at least 35° [148].

Sensing noise (unsuppressed)	Displacement noise (suppressed)
Quantum shot noise	Seismic noise
ADC noise	Quantum radiation pressure noise
Op-amp input voltage noise	Coating Brownian thermal noise
Op-amp input current noise	Coating thermooptic noise
Photodiode quantum efficiency	Suspension thermal noise
	DAC noise

Table 5.1: Noise categories within the Sagnac speed meter experiment. Sources of sensing noise arise from the detection of electronic signals from the interferometer for the purposes of data acquisition and control. Sources of displacement noise arise from the path between the controller’s feedback signal and the test masses being controlled. Displacement noise is suppressed by the controller’s loop gain, but, as discussed in appendix B.4, the controller can only usefully suppress noise to the level to which it can sense the noise, i.e. the level governed by sensing noise.

5.3 Velocity control

In this section we approach the control of the Sagnac speed meter in a similar fashion to the control of a Michelson interferometer by feeding back the signal measured at the output port to the arm cavity actuators.

5.3.1 Control loop

A control loop schematic using the calculations, filters and servos presented in section 5.2 is shown in figure 5.15. The items contained within the grey box are implemented in software and hardware as part of CDS. The lower section contains the blocks which exist in the analogue domain.

5.3.2 Low frequency noise projection with velocity feedback

To reach the desired sensitivity of the interferometer it is crucial to understand the noise characteristics associated with the sensing and control apparatus employed in the experiment. Individual noise sources, arising for example from the BHD op-amp electronics, can be projected into units of differential displacement-equivalent noise using the linear projection technique [149]. The sources of noise can be logically separated into two categories: *sensing noise* and *displacement noise*, as shown in table 5.1. Both sources of noise are fed back to the test masses because in practice it is not possible for the controller to distinguish them. Note that we neglect force noise from the actuators, for instance arising from stray magnetic fields, due to the absence of good models. These noise sources are not, however, expected to dominate the noise sources under consideration in the frequency band of interest.

Sources of sensing noise are associated with the readout of the variable of interest—in the case of the Sagnac speed meter the positions of the test masses’ surfaces—but do not directly influence the variable of interest in an open loop measurement. Sources of sensing noise include quantum shot noise, electronic noise including op-amp noise as measured in section 5.2.6 and quantisation noise due to the ADCs as described in section 5.2.7.1.

Displacement noise sources directly influence the positions of the test mass surfaces being measured by the interferometer and are therefore transformed by the dynamics of the test masses [89]. As the readout variable in the BHD is the time derivative of position, the control system measures and actively suppresses these noise sources. Significant sources of displacement noise in the Sagnac speed meter experiment are quantum radiation pressure noise, seismic noise, suspension thermal noise (see section 2.2.2.2) and coating Brownian noise arising from the dielectric coatings present upon the cavity mirrors (see section 2.2.2.1).

The noise projection for $L_{(-)}$, calculated using Optickle and the control noise modelling tool *SimulinkNb*⁶, is shown in figure 5.16. The rms $L_{(-)}$ displacement this creates is shown in figure 5.17 as a function of time. It shows that, as the interferometer is held at its operating point, over a period of several hours the expected drift is large enough for the cavities to become uncontrollable (see section 5.2.4).

Although for sensing noise we only consider the sources listed in table 5.1, in the real experiment there will be other contributing forms of time-varying offset present upon the BHD error signal:

- residual local oscillator light due to temperature-driven imbalances in the BHD beam splitting ratio and photodiode quantum efficiencies,
- signal from common mode arm cavity motion due to imbalanced beam splitting at the main beam splitter [89],
- changing thermoelectric potentials and op-amp drift in electronics,
- and any other time-varying effects.

As such, the estimated rms displacement shown in figure 5.17 represents a “best case” scenario where the op-amp’s electronic noise is the dominant effect at low frequencies, and this drift becomes unacceptably large after a few hours. To allow for long term cavity stability it is essential for the error signal to contain a signal significantly above the electronic noise at low frequencies. In the next section we present a strategy for obtaining an error signal of suitable magnitude across the entire control bandwidth.

⁶Available as of the time of writing at <https://github.com/cwipf/SimulinkNb/>.

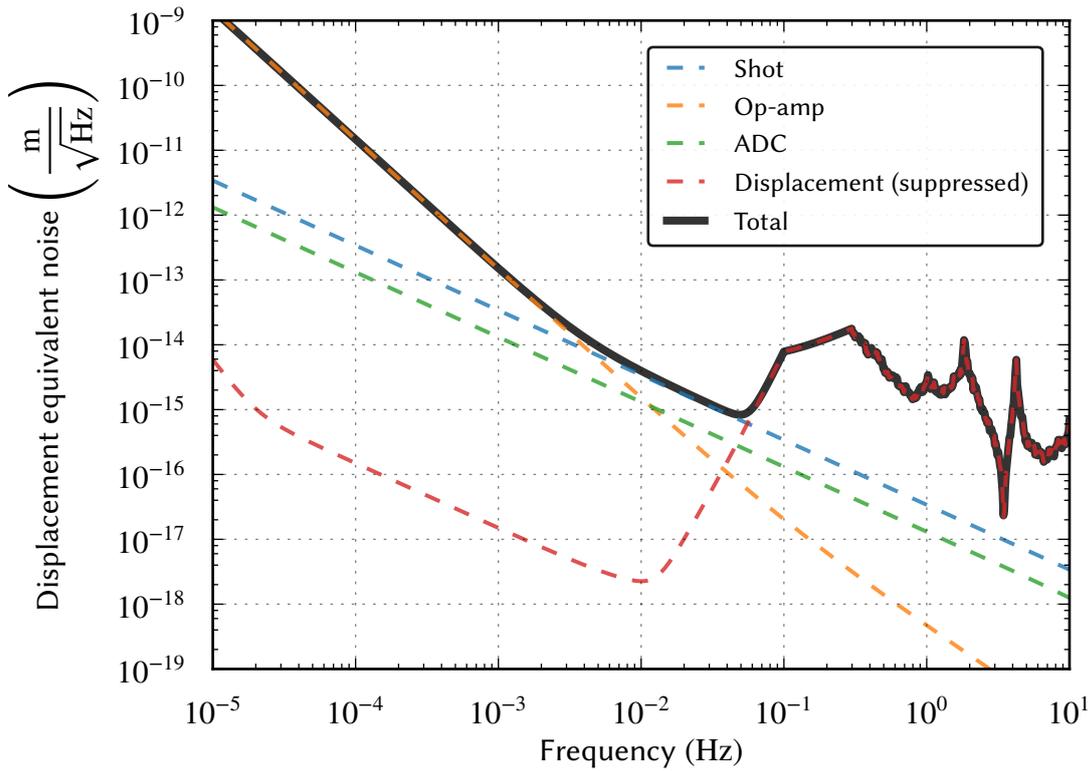


Figure 5.16: Spectral density showing the noise associated with the readout of $L_{(-)}$ at the BHD. The significant noise sources associated with sensing (shot, op-amp and ADC noise) are shown alongside the contribution from suppressed displacement noise sources. Below around 5 mHz the dominating sensing noise is due to the op-amp electronics. Lab measurements of seismic noise have been made down to 0.3 Hz, and this dominates the displacement noise between 10 mHz and 10 Hz. The assumption has been made that the noise is sharply suppressed below the microseism at 0.1 Hz. Dominant displacement noise below 10 mHz is due to the DACs.

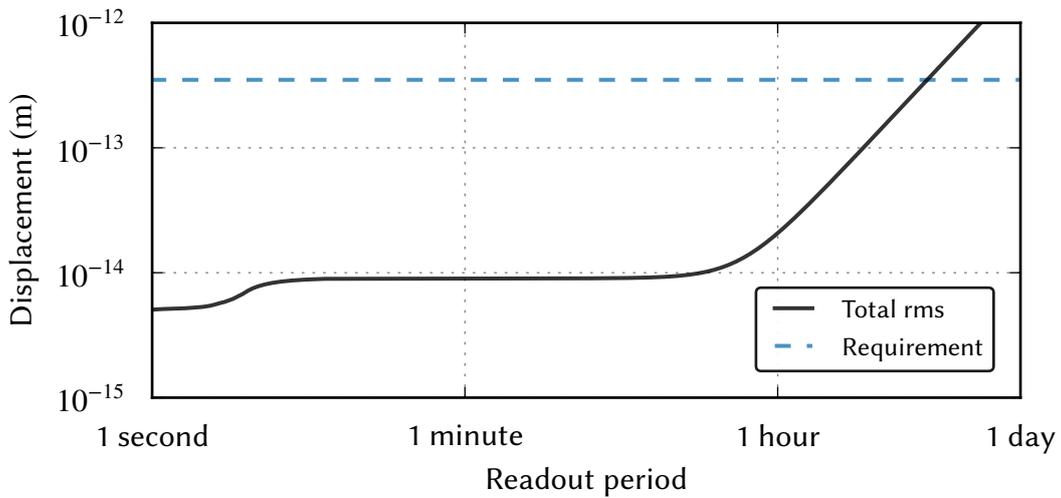


Figure 5.17: Root-mean-square noise projection for $L_{(-)}$ using velocity feedback. The requirement is exceeded beyond a few hours, after which the noise due to the BHD readout is enough for the cavities to drift beyond the displacement requirement and lose sensitivity.

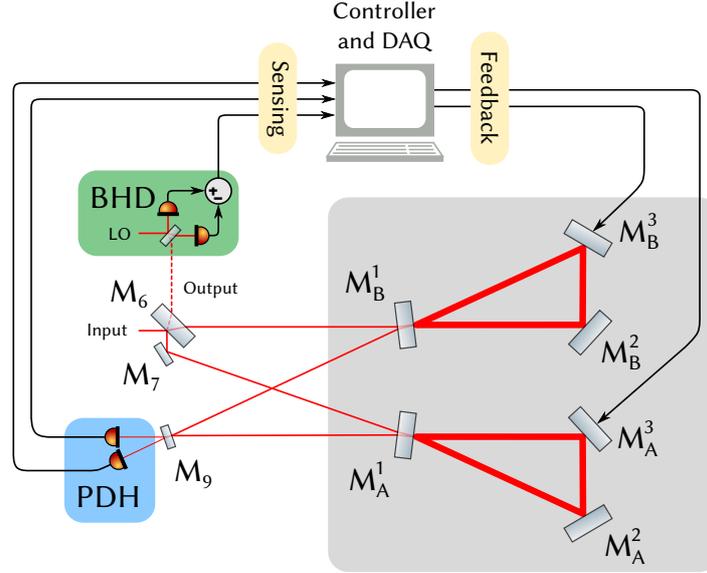


Figure 5.18: Simplified layout of the Sagnac speed meter experiment including both displacement and velocity feedback paths. Apart from the components shown already in figure 5.3, this diagram includes the PDH readout used to provide a displacement error signal at low frequencies.

5.4 Velocity-displacement control

Light from each counter-propagating mode is incident upon M_9 , and as such this is a natural port in which to separate the modes and sense the motion of each arm cavity (see the shaded blue region of figure 5.18). Using RF modulation, for instance via the Pound-Drever-Hall (PDH) technique [121], it is possible to obtain a displacement error signal for each cavity that, unlike the velocity signal from the BHD, has flat response at dc, with a similar cavity pole frequency (see figure 5.19). The individual cavity PDH signals can be mixed to obtain a measurement of $L_{(-)}$, and the frequency dependence of the signal s_{PDH} is, following ref. [43], given by

$$s_{\text{PDH}}(\Omega) \propto \sqrt{\frac{\gamma_{\text{arm}}}{(\Omega^2 + \gamma_{\text{arm}}^2)}} L_{(-)}, \quad (5.6)$$

ignoring again the effect of losses and constant terms as with equation (5.1). Note that for $\Omega \ll \gamma_{\text{arm}}$, the response is flat as expected for a displacement measurement and as such the PDH readout offers a suitable signal to sense $L_{(-)}$ at low frequencies.

5.4.1 Combined filter

The separate velocity and displacement readouts contain the same fundamental information about $L_{(-)}$, albeit with different response functions. We can express the signal at output

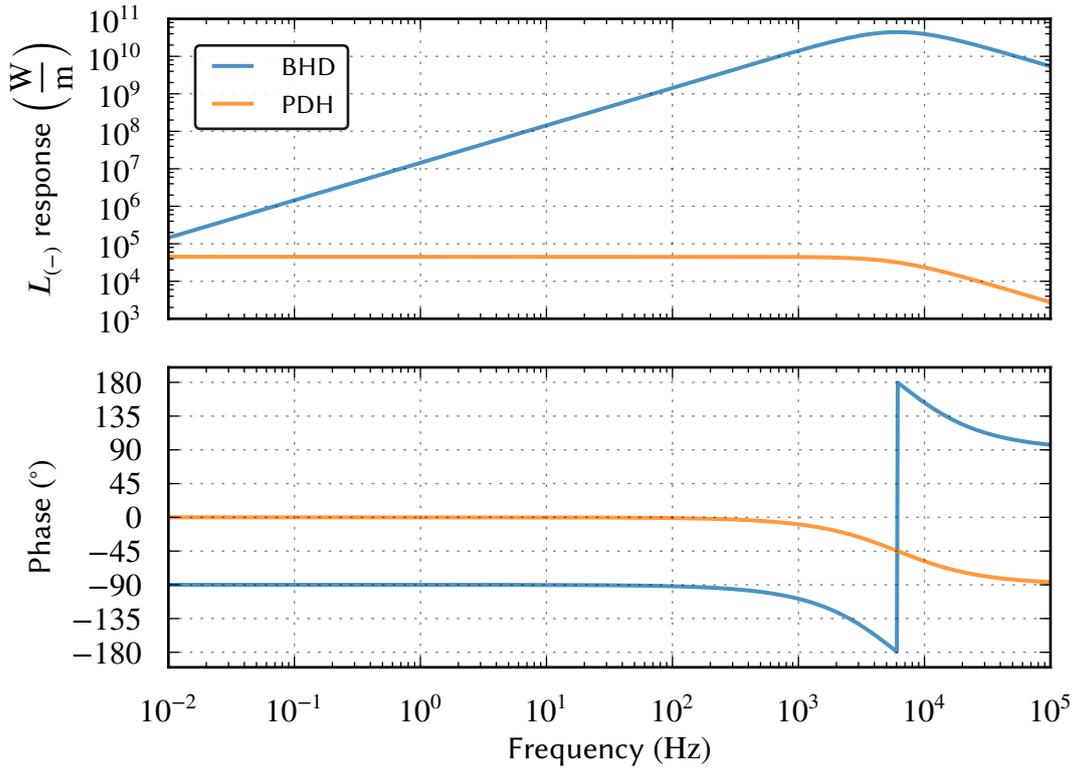


Figure 5.19: The frequency response of the differential arm cavity degree of freedom to the PDH readout alongside that of the BHD readout, simulated numerically with Optickle. At low frequencies, the PDH readout is flat whereas the BHD signal decays towards zero. The flat PDH error signal assists with the long term stability of the Sagnac speed meter experiment.

field i as a function of the k^{th} mode of motion, $\hat{o}_{k,i}(\Omega)$, as [43]

$$\hat{o}_{k,i}(\Omega) = L_k(\Omega) + \frac{\hat{n}_i(\Omega)}{R_{k,i}(\Omega)} \quad (5.7)$$

where L_k is the position of mode k , $\hat{n}_i(\Omega)$ is the noise at field i and $R_{k,i}(\Omega)$ is the optomechanical transfer function of mode k to field i . The definition of a field in this case refers to that of a single signal sideband at frequency Ω . The total time domain signal on a sensor due to the k^{th} mode at the location of the output field will see a combination of the upper and lower signal sidebands:

$$\hat{o}_{k,i}(t) = \int_0^\infty \frac{d\Omega}{2\pi} \left(\hat{o}_{k,i}(\omega_0 + \Omega) + \hat{o}_{k,i}^\dagger(\omega_0 - \Omega) \right) e^{-i\Omega t}, \quad (5.8)$$

where ω_0 is the angular frequency of the carrier.

Displacement noise sources are implicit in L , and we assume the sensing noise other than quantum noise associated with both the BHD and PDH readouts is the same. The excess noise at each readout port is therefore due to \hat{n}_i , the quantum vacuum entering at open ports

within the interferometer. The presence of such vacuum noise limits the sensitivity of the readout in the measurement band. For this reason the reflectivity of M_9 must be chosen to be close to unity, therefore only a small amount of light is available to the displacement readout for use as a low frequency error signal.

The blending of the two readouts must be causal in order to preserve the stability of the control loop. Due to the different frequency response of the PDH and BHD signals, a practical approach is simply to combine the two readouts together with different gain. In this case, the blending filter has a 90° phase difference at the crossover frequency determined by the relative gains for frequencies below the cavity pole, providing adequate phase margin. As we are concerned with greatest sensitivity above 100 Hz, yet must use the displacement information provided by the PDH signal to control drifts at low frequencies, the blending filter's crossover frequency can in this case be placed in the region of 1 Hz. This results in BHD feedback with magnitude at least a factor 100 greater than that of the PDH feedback above 100 Hz, preventing the PDH's displacement-proportional noise in the measurement band from adversely affecting the overall sensitivity there.

5.4.2 Control loop

The intended control loop schematic for the experiment with combined feedback is shown in figure 5.20. This is similar to the loop shown in figure 5.15, but with the addition of signal paths from the interferometer's PDH readout to the controller and two gain blocks to control the way in which the velocity and displacement readouts are combined. An additional step is modelled with the PDH readout: the demodulation gain. This is the gain the signal receives as a result of the mixing of the sideband frequency as part of the PDH technique. The responsivity of the photodiodes used for the PDH readout has also been assumed smaller, at 0.8 A W^{-1} , due to this channel's relaxed loss requirements. For the same reason, the photocurrent from the PDH sensors is not directly subtracted, instead being combined into a signal representing $L_{(-)}$ within CDS.

5.4.3 Low frequency noise projection with combined feedback

The $L_{(-)}$ noise projection for a simple combined filter as discussed in section 5.4.1 is shown in figure 5.21, showing the sensitivity the control system has to the interferometer's differential arm cavity motion. At the expense of a very small increase in noise at 1 Hz over the velocity-only feedback, the rms curve in figure 5.22 shows a clear reduction in residual displacement over longer periods.

The open loop gain of the system is shown in figure 5.23. This shows the unity gain frequency to be 350 Hz, with a phase margin of 44° . With this controller the system's differ-

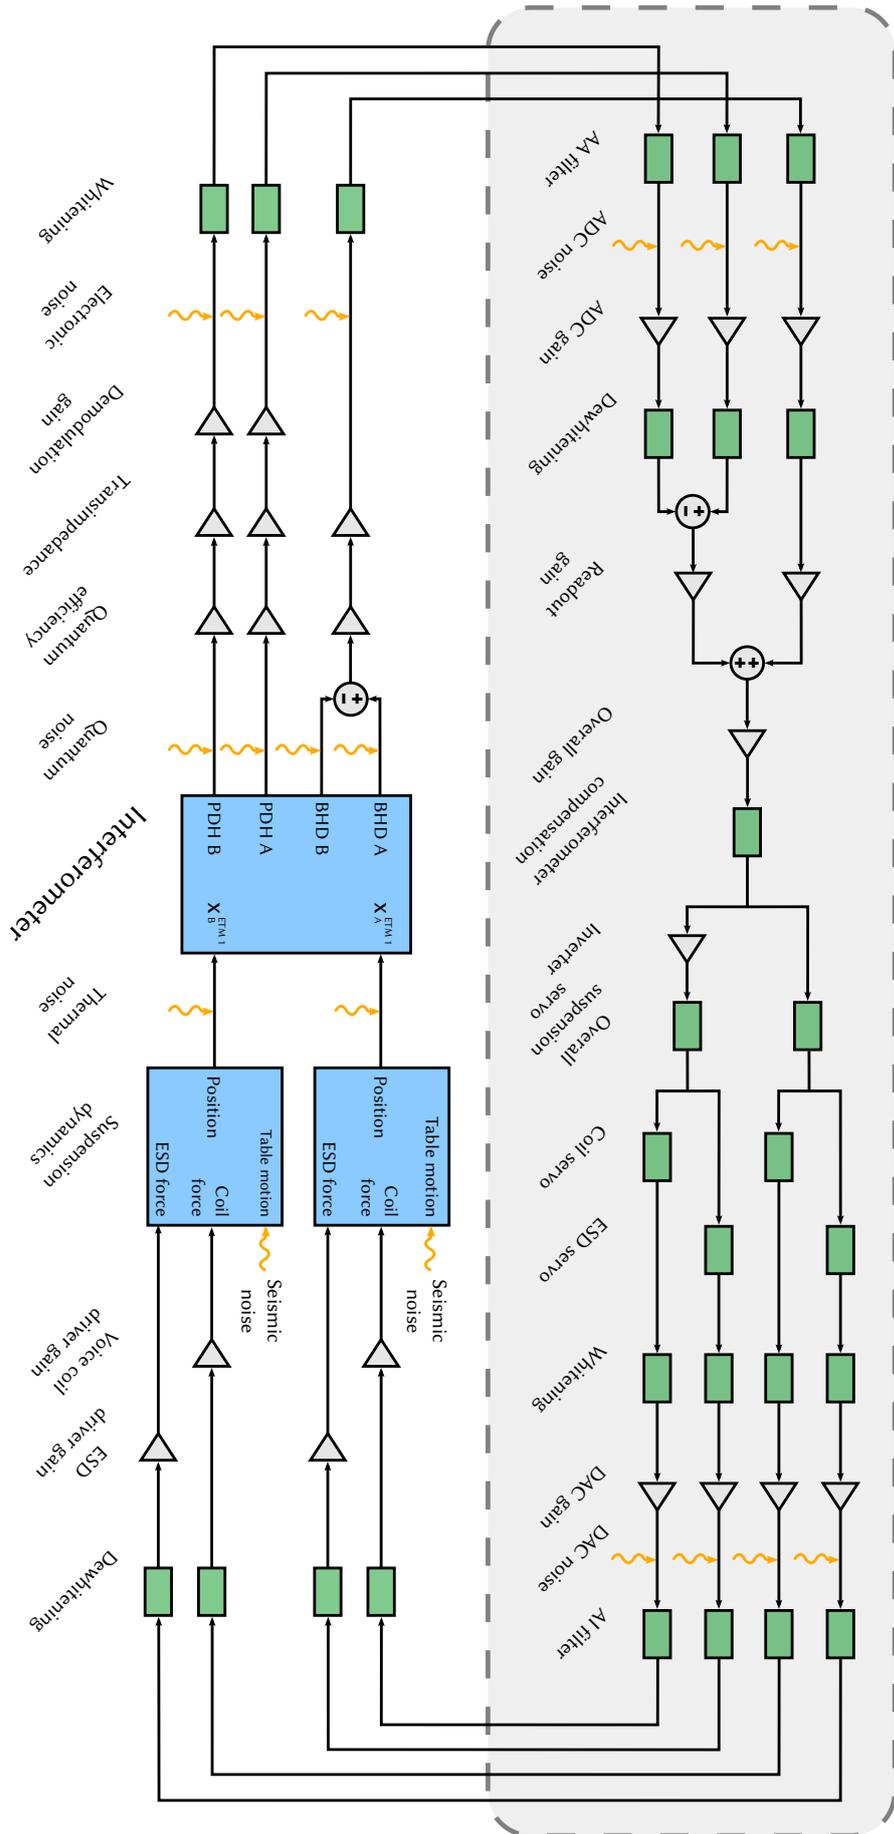


Figure 5.20: Sagnac speed meter experiment control loop model. This control loop is similar to that shown in figure 5.15, but with the addition of components used to send the displacement-sensitive PDH readout to CDS. Within CDS, additional gain blocks allow for control over the way in which the velocity and displacement readouts are combined into one feedback signal.

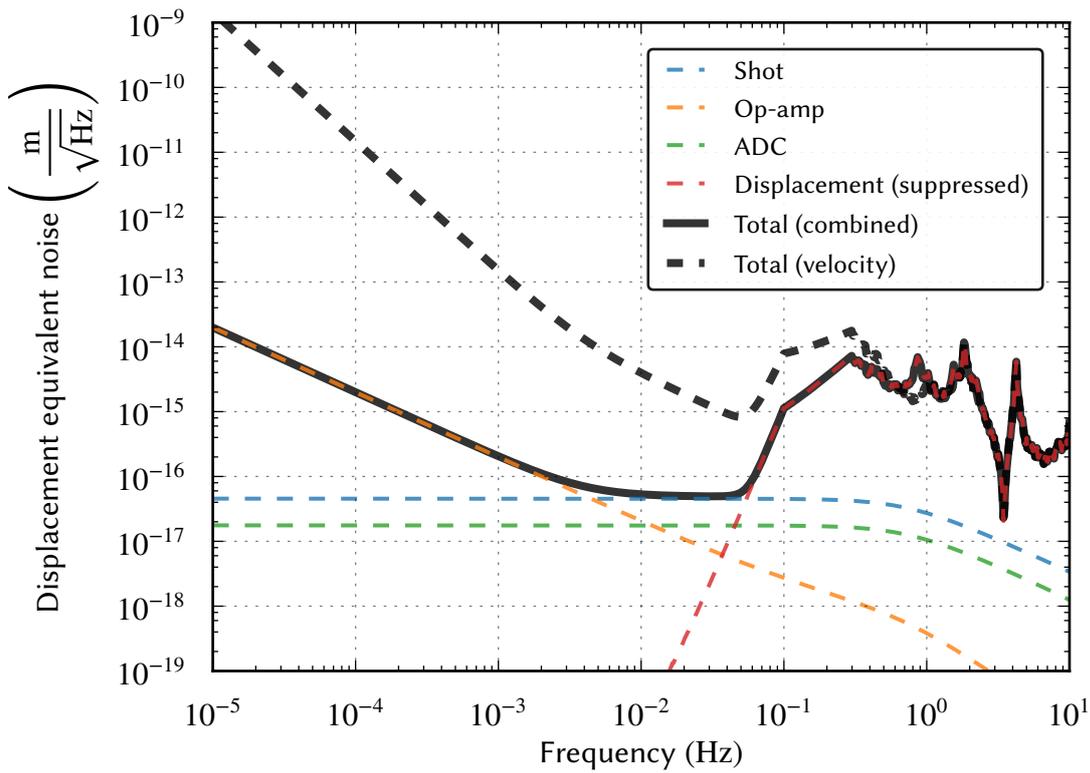


Figure 5.21: Spectral density noise projection for $L_{(-)}$ using both displacement and velocity feedback. The mixing of displacement information into the feedback signal at low frequencies leads to greatly reduced equivalent $L_{(-)}$ noise, as the noise from the velocity readout electronics is suppressed by the strong displacement response. As with figure 5.16 some important individual contributions to the overall noise are shown.

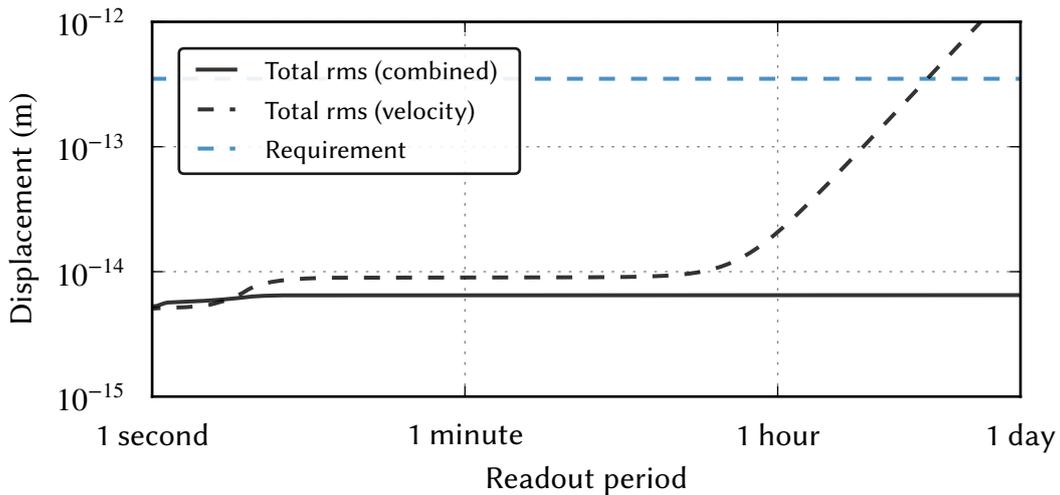


Figure 5.22: Root-mean-square noise projection for $L_{(-)}$ using both displacement and velocity feedback. Unlike the velocity-only feedback, the combination of velocity and displacement feedback prevents the rms cavity mirror displacement from exceeding the required control precision after a few hours, instead allowing stability over much greater periods.

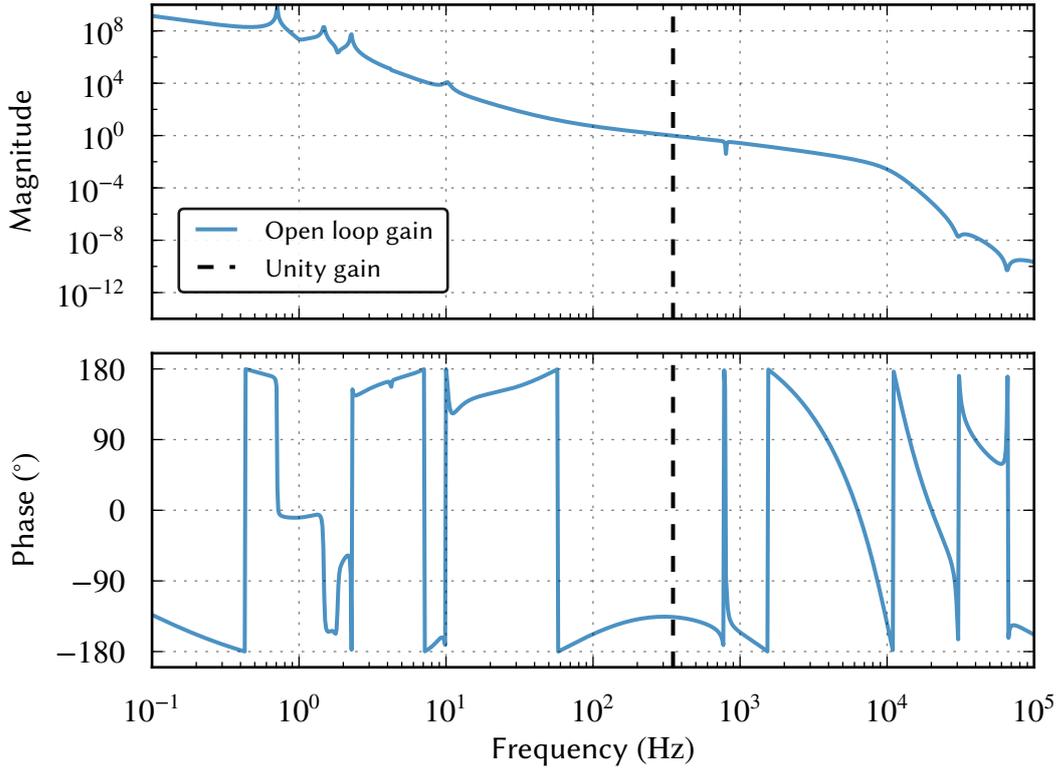


Figure 5.23: Simulated Sagnac speed meter controller open loop gain. The majority of the gain is applied to correct displacements due to seismic noise below 10 Hz. The unity gain frequency is 350 Hz and the phase margin is 44°.

ential arm cavity displacement is able to be controlled to within the requirement shown in section 5.2.4.

5.5 Matched filter

By considering cross-correlations in the quantum noise at the BHD and PDH readouts, it is possible to calculate a matched filter with which to combine the two outputs in post-processing so as to minimise the total noise spectral density. This can be applied to the offline data in order to calculate the noise spectral density of the interferometer during the data acquisition period.

The noise at each readout is the sum of the quantum noise inputs at open ports propagated through the interferometer with appropriate transfer functions, so we can rewrite \hat{n}_i in equation (5.7) in terms of the quantum noise amplitudes \hat{q}_m entering at N_p open ports:

$$\hat{n}_i(\Omega) = \sum_{m=1}^{N_p} M_{m,i}^{\text{ff}}(\Omega) \hat{q}_m(\Omega), \quad (5.9)$$

where $M_{m,i}^{\text{ff}}(\Omega)$ represents the transfer function between input field m and output field i for signal sideband Ω . The cross-correlation spectral density for unity noise at the i^{th} and j^{th} output channels, for the k^{th} mode, is given by [40]

$$S_{k,ij}(\Omega) = \sum_{m=1}^{N_p} \frac{\left[M_{m,i}^{\text{ff}*}(\Omega) M_{m,j}^{\text{ff}}(\Omega) + M_{m,j}^{\text{ff}*}(-\Omega) M_{m,i}^{\text{ff}}(-\Omega) \right]}{\left[R_{k,i}^*(\Omega) + R_{k,i}(-\Omega) \right] \left[R_{k,j}(\Omega) + R_{k,j}^*(-\Omega) \right]}. \quad (5.10)$$

This reduces to the following form for noise entering the same port in which it exits:

$$S_{i,i} = \frac{1}{2} \frac{\left| M_{i,i}^{\text{ff}}(\Omega) \right|^2 + \left| M_{i,i}^{\text{ff}*}(-\Omega) \right|^2}{\left(\left| R_{k,i}(\Omega) \right| + \left| R_{k,i}^*(-\Omega) \right| \right)^2}. \quad (5.11)$$

Assuming a filter $\alpha(\Omega)$ combines the BHD ($i = 1$) and PDH ($i = 2$) fields, its output for L_k would be:

$$\begin{aligned} \hat{o}_{k,\text{combined}}(\Omega) &= \alpha(\Omega) \hat{o}_{k,1}(\Omega) + (1 - \alpha(\Omega)) \hat{o}_{k,2}(\Omega) \\ &= \left(\alpha(\Omega) L_k(\Omega) + (1 - \alpha(\Omega)) L_k(\Omega) \right) \\ &\quad + \frac{\alpha(\Omega) \hat{n}_1}{R_{k,1}(\Omega)} + \frac{(1 - \alpha(\Omega)) \hat{n}_2}{R_{k,2}(\Omega)}. \end{aligned} \quad (5.12)$$

The corresponding total noise power spectral density of the combined readout is then:

$$\begin{aligned} S_{\text{readout}} &= |\alpha|^2 S_{n_1,n_1} + |1 - \alpha|^2 S_{n_2,n_2} \\ &\quad + \Re \left[\alpha^* (1 - \alpha) S_{n_1,n_2} \right] \\ &\quad + \Re \left[\alpha^* (1 - \alpha) S_{n_2,n_1} \right], \end{aligned} \quad (5.13)$$

where S_{n_1,n_1} is the noise power spectral density at the BHD port due to vacuum entering at the BHD port, S_{n_2,n_2} is the noise power spectral density at the PDH port due to vacuum entering at the PDH port, and S_{n_1,n_2} and S_{n_2,n_1} are the noise power spectral densities for noise entering at one port and exiting at the other. The optimal matching filter α_{matched} can be determined by minimising equation (5.13) over α :

$$\alpha_{\text{matched}} = \frac{S_{n_1,n_2} - S_{n_1,n_2}^*}{S_{n_1,n_1} + S_{n_2,n_2} - \Re \left[S_{n_1,n_2} \right] - \Re \left[S_{n_2,n_1} \right]}. \quad (5.14)$$

The reflectivity of M_g is implicit in both the field-to-field and mode-to-field transfer matrices for each signal sideband, \mathbf{M}^{ff} and \mathbf{R} , respectively, and as such α_{matched} depends on the value of M_g .

The matrices \mathbf{M}^{ff} and \mathbf{R} are not calculated in Optickle by default, and so some modifications to the code were necessary (see appendix C.3). The effect of M_g 's reflectivity on α_{matched} is

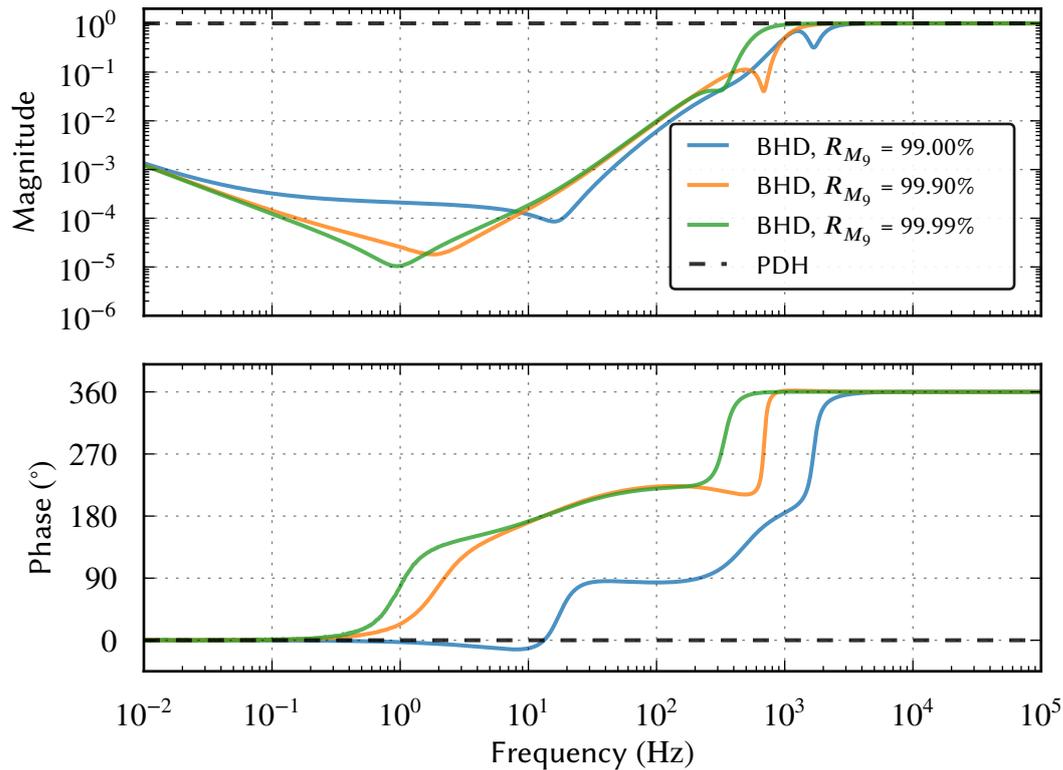


Figure 5.24: Matched filters to combine the BHD and PDH signals for different values of M_0 reflectivity. The red, yellow and green curves are the coefficients to be applied to the BHD signal with respect to the PDH signal before the two are combined, for different M_0 (power) reflectivities. The black, dashed curve is the (unity) coefficient to be applied to the PDH signal. For all values of M_0 shown, the optimal combination involves suppressing the BHD signal with respect to the PDH at frequencies below around 1 kHz.

shown in figure 5.24. Note that, because it is calculated with offline spectral densities and not tested for stability, the filter predicted by equation (5.14) is not necessarily realisable and therefore cannot be implemented in the real-time control system. A *causal* Wiener filter—one which is implementable in a physical experiment—has previously been calculated for *single-readout* interferometers [150, 151], but a similar calculation for more than one readout has to the author’s knowledge not been found as of the time of writing.

While the matched filter is presented in the context of the Sagnac speed meter experiment, it is a general solution for an interferometer with multiple readouts for a single variable and may prove useful for future gravitational wave detectors utilising QND techniques regardless of topology.

5.5.1 Noise budget

In order to show that quantum noise is reduced with respect to an equivalent Fabry-Perot Michelson interferometer, the design of the Sagnac speed meter experiment intends for it to be the limiting noise source in a frequency band in the region of a few hundreds of Hz [137]. Using the linear projection technique outlined in section 5.3.2, each anticipated significant source of noise has been estimated and projected into $L_{(-)}$ noise to discover the limiting sources across the control bandwidth, and verify that the experiment will be limited by quantum noise in the intended band. The noise budget was created in steps. First, the individual noise sources contributing to the displacement of the test masses and the sensing of the interferometer signals were estimated individually, as described in section 5.2. Each individual noise source was then projected to the point in the loop where the data is recorded—CDS—with all other noise sources switched off. Here, the open loop gain of the controller was applied to simulate the effect the loop has in suppressing displacement noise sources. Finally, to understand the sensitivity in terms of $L_{(-)}$, the noise spectral density was divided by the transfer function from differential arm cavity mirror motion to CDS. The noise budget for each significant noise source, calculated with SimulinkNb, is presented in figure 5.25. This budget is similar to the one presented in ref. [137], with the difference that this noise budget is the product of a comprehensive control noise study. The noise contribution from the PDH feedback is shown to be vastly lower than the limiting noise in the intended measurement band, justifying its inclusion.

The sensitivity between 100 Hz and 700 Hz, shaded in blue, is the quantum noise limited measurement band. This band is constrained from below by expected test mass suspension mechanical mode cross-couplings (not shown) and from above by the first violin mode of the ETM suspensions. Suspension thermal noise is the second highest noise source present in this band and is at most a factor of 2.3 below quantum noise, allowing a careful direct measurement of quantum radiation pressure noise to be made in this region. The contribution to the quantum noise from the PDH feedback is far below the total quantum noise, showing that the use of the displacement readout as part of the combined filter presented in section 5.4.1 does not significantly affect the sensitivity of the Sagnac speed meter in the desired band.

5.6 Experimental parameters

The parameters, including those developed over the course of this work, are shown in table 5.2. Unless otherwise stated, the mirrors specified in the figures and simulations are assumed to have unity reflectivity (apart from beam splitters, which are assumed to have equal transmissivity and reflectivity). All listed transmissivities represent power, no sub-

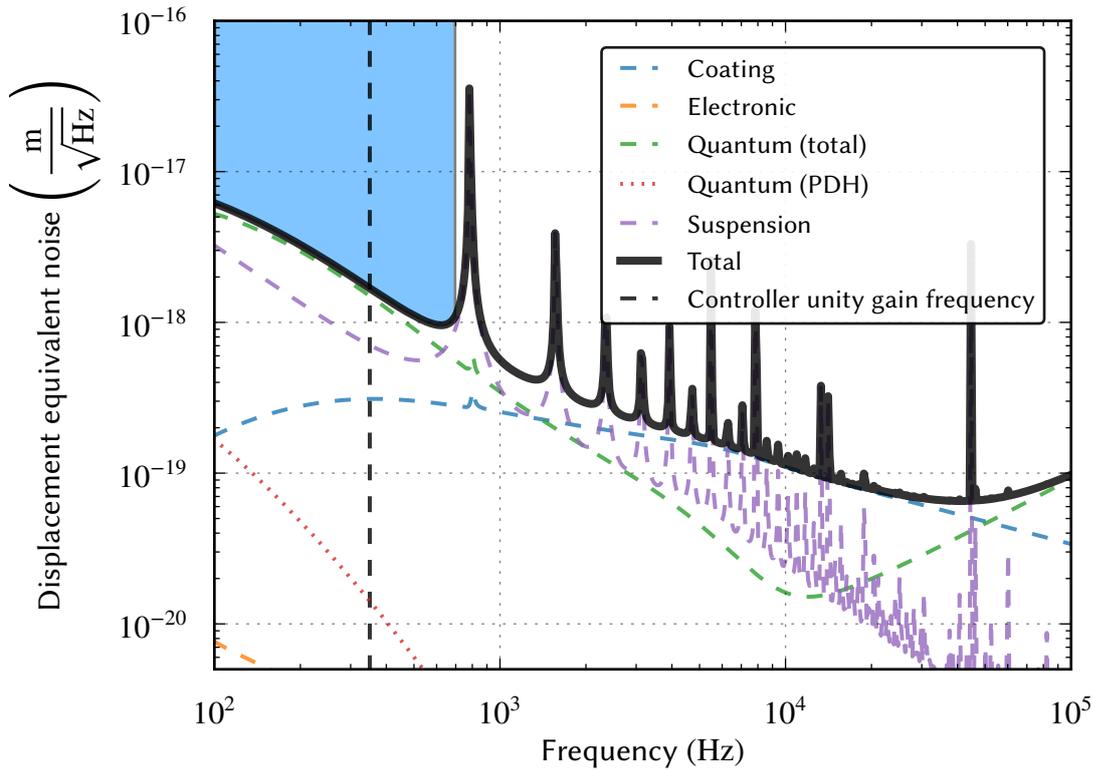


Figure 5.25: Sagnac speed meter differential mode noise budget for the combined filter scheme with sensing and control noise taken into account. The shaded blue region represents the frequency band at which the intended direct measurement of reduced quantum radiation pressure noise is to be made in the experiment. The quantum noise contribution from the PDH readout is more than an order of magnitude smaller than the total quantum noise, showing that its inclusion in the combined filter is not harmful to the overall sensitivity in this band. Displacement noise sources, such as coating noise, are suppressed by the loop gain below the unity gain frequency.

strate loss is assumed for any optic and all simulations have been performed using the plane-wave approximation.

5.7 Summary and Outlook

We have outlined a realistic control strategy for the Sagnac speed meter experiment taking into account the sensors and actuators to be used, and demonstrated that positional drifts of the cavity mirrors at low frequencies due to sensing noise lead to an inability to control the cavity mirrors over time scales longer than a few hours. We have shown that this drift can be suppressed by taking a small amount of light from the path between the arm cavities to provide a displacement readout, and that this does not significantly affect the sensitivity of the main, velocity readout. A combination of the displacement and velocity readouts provides a suitable error signal for the control of the arm cavity differential mode

Parameter	Fiducial value
Laser wavelength λ_0	1064 nm
Input power	1.8 W
L_{RT}	2.83 m
M_0 transmissivity	10^4 ppm
T_{ITM}	700 ppm
Arm cavity FWHM	12.2 kHz
Arm cavity finesse	8663
BHD quantum efficiency	0.95 A W^{-1}
PDH quantum efficiency	0.80 A W^{-1}
R_{T}	10 k Ω
PDH demodulation gain	21 dB
ADC/DAC quantisation noise	$1.8 \mu\text{V} \sqrt{\text{Hz}^{-1}}$
ETM mass	113 g
ETM fibres	4
ETM fibre diameter	40 μm
ETM fibre length	200 mm
ITM mass	0.86 g
ITM fibres	2
ITM fibre diameter	10 μm
ITM fibre length	100 mm
Suspension vertical-to-horizontal coupling	0.01

Table 5.2: Experimental parameters. The properties for the suspensions and test masses are given should the reader wish to reproduce the suspension thermal noise spectral density presented in figure 5.25.

at all relevant frequencies without spoiling the quantum non-demolition effect at higher frequencies, facilitating measurements with arbitrary integration time and allowing the Sagnac speed meter to reach its design sensitivity.

The effect that the mixing of displacement and velocity information should have on the interferometer's stability should be quite straightforward to measure. As shown in figure 5.22, the drift from the lack of an error signal at low frequencies in the velocity-only case will cause the interferometer to lose its sensitivity after a short period of time. To enable measurement integration times longer than a few hours—which will probably be necessary to obtain well defined measurements of radiation pressure noise to meet the experiment's goals and indeed for future gravitational wave observatories to conduct science runs of months in duration—the displacement-proportional signals will have to be fed back at frequencies below the measurement band. A simple test of the two control laws should highlight a noticeable contrast.

Since the main readout of any interferometer primarily sensitive to velocity will encounter the problem of vanishing signal in the presence of flat or increasing sensing noise at low frequencies, we believe the solution presented in this work is applicable to any audio-band

speed-meter.

Chapter 6

Infrastructure for the control of a plate capacitor electrostatic drive with reduced seismic coupling

6.1 Electrostatic drives as actuators in suspended interferometer experiments

Suspended test masses in interferometers require positional corrections in order for the interferometer to be kept at its operating point. This is typically provided via actuators on the suspension system and predominantly involves voice coil actuators composing magnets and wound wire. Force noise can be introduced to the test masses by their actuators due to various effects such as stray magnetic field coupling, electronic noise in the driver circuitry, Barkhausen noise [139] and seismic coupling via the actuator attachment point. The first two effects can usually be mitigated with appropriate design and shielding, for example by choosing appropriate electronic components and by making the magnets small and the electromagnetic environment quiet. The third effect is often mitigated by suspending the actuators from a separate suspension behind the test mass called a *reaction* suspension. This provides seismic filtering to the actuators such that the ground motion coupling introduced to the test masses from the actuators is of similar magnitude to the ground motion the test masses would in any case receive with no actuation.

As introduced in chapter 5, electrostatic drives (ESDs) are a type of actuator employed in GEO 600 and Advanced LIGO for fast (high frequency) corrections to the interferometer. This actuator creates a force on a dielectric test mass by creating a potential difference between anodes and cathodes upon the face of the dielectric reaction mass. Electromagnetic field gradients are then formed in such a way that a force can be applied to the test mass in

a particular direction.

6.1.1 Comb electrostatic drive design

The ESD design used in current generation detectors involves a comb of interlocking anodes and cathodes across which a voltage is applied to create the desired force. Alignment control is achieved through the use of multiple sets of combs on the face of the reaction mass, and the sign of the voltage applied to each set can be controlled to induce torque.

There are a number of problems with this approach actuation. There are obviously cost and technical implications for the use of a reaction suspension system behind each main suspension. The alignment of this second suspension must be controlled and damped in the presence of displacement noise just like the first suspension. With the use of ESDs, the gap created between the reaction and test masses can also lead to noise from *squeezed film damping* due to residual gas in the vacuum system [152]. One of the most important issues in the use of this type of ESD and the more typical voice coil actuators, however, is the limited clear aperture behind each test mass. In the case of this ESD design, the metal comb pattern on its corresponding reaction mass can clip the transmitted beam. The beam size on the ETMs of Advanced LIGO is around 6 cm and so in this case if the transmitted light were to be measured for the purposes of sensing and control the choice would have to be made between allowing the beam to be clipped with the associated technical problems this introduces and a reduction in the space available on the reaction mass for the electrostatic comb structure.

6.1.2 Plate capacitor electrostatic drive design

An alternative design to the comb arrangement is to use parallel metal plates with faces perpendicular to the beam axis. Applying a potential difference between these plates effectively creates a capacitor, and the fringe electric field between the edge of the plates and a dielectric test mass produces a small component of force along the beam axis.

The electrostatic energy in the capacitor is given by the volume integral of the electric field created by the potential difference multiplied by the permittivity of the volume enclosed by the plates. For the case when the dielectric mirror is partially inside the volume enclosed by the plates, it can be shown that the energy is given by [153]

$$E = \frac{1}{2}wd (\Delta\phi/d)^2 (\epsilon z + \epsilon_0 (l - z)), \quad (6.1)$$

where w and l are the plate width and height, respectively, d is the dielectric slab's thickness, $\Delta\phi$ is the potential difference, ϵ and ϵ_0 are the dielectric and vacuum permittivities

and z is the offset of the slab along the beam axis. The electrostatic force the capacitor applies to the dielectric slab in the direction of the beam axis is given by the gradient of the electrostatic energy:

$$\begin{aligned} F(z) &= \nabla E(z) \\ &= (\epsilon - \epsilon_0) \frac{w\Delta\phi^2}{2d}. \end{aligned} \tag{6.2}$$

This shows that the force produced by the parallel plates depends on the voltage applied across the space separating the plates alongside the plate geometry and separation. The force is greater for higher voltage and wider plates with smaller separation.

Originally suggested by Wittel et al. [154], this parallel plate capacitor ESD has potential applications as an actuator for test masses in suspended interferometers where it was shown that the dc force provided by such an actuator with dimensions applicable to the AEI 10 m prototype is around $1.5 \mu\text{N}$ at 1 kV, corresponding to a displacement of around $0.3 \mu\text{m}$. This result taken alongside inspection of equation (6.2) for silica mirrors, feasible plate geometries and voltages shows that this ESD design is only suitable for small corrections. As the pendulum systems used in suspended interferometers filter ground motion to a greater extent at higher frequencies, this type of actuation tends to lend itself more to the control of radiation pressure and other displacement effects at frequencies above 50 Hz where seismic motion is typically insignificant.

Apart from the availability of a clear aperture behind the test mass for sensing and control, a further advantage of this type of actuator is that it is highly unsusceptible to displacement noise induced by seismic motion; this is due to the shallow gradient of the fringe field at the edge of the plates. Small perturbations with respect to the position of the plates do not strongly couple to the direction or magnitude of the force produced by this type of ESD; this contrasts to the seismic noise coupling produced not only by voice coils but also by the metal comb type ESD, as shown in ref. [154]. The remaining contributions to displacement noise with this type of ESD come from misalignments in the plates with respect to one another and electronic noise in the creation of the potential difference across the plates. This chapter will address the design of electronics capable of providing low noise actuation upon the ETMs within the Sagnac speed meter experiment.

6.2 Electrostatic drives for the Sagnac speed meter experiment

The plan for the Sagnac speed meter experiment is to adopt a plate capacitor design for the actuation of the ETMs so that the transmitted beam is available for the purposes of sensing

Parameter	Value
ETM diameter	48.6 mm
ETM thickness	24.5 mm
Single plate width	48.6 mm
Single plate length	50 mm
Nominal plate separation	58.6 mm

Table 6.1: Plate capacitor and optic parameters for the Sagnac speed meter experiment.

and control and to reduce the number of suspensions required in the limited space within the vacuum enclosure.

Due to the dimensions of the 100 g ETMs for which the ESDs will eventually be used, the plate capacitor parameters shown in table 6.1 were deemed appropriate. As equation (6.2) assumes that the dielectric test mass completely fills the region between the plates it is only an approximation for this situation where each *round* test mass is offset from the plates to avoid ground motion coupling and friction. In order to gain a more complete understanding of the force produced on the test mass by the plates finite element simulations were conducted. A basic model of the plates and test mass described in table 6.1 was built with ANSYS in order to model the effect of force coupling in each direction from an applied voltage, given perfectly aligned plates. Figure 6.1 shows the results of this simulation, with the force along the beam axis given by the z-direction and the forces in the transverse directions given by x and y as a function of potential difference. For a voltage of 750 V, which is what we think our vacuum feedthrough connectors should be able to handle, the ESD is able to provide a force of around -1.48×10^{-6} N. The force-voltage behaviour is approximately linear in this region giving a gradient of -3.68 nN V⁻¹.

6.2.1 Maximum actuation requirements

At the operating point the ESD's feedback force will be significantly lower than its maximum. The requirement for the ESD's range is set by the lock acquisition sequence performed to reach the low noise state, where the ESDs have to be able to respond to changes in classical radiation pressure forces as the light power is increased in the cavities. The suggested lock acquisition schemes for the Sagnac speed meter experiment require of the order μ N actuation at high frequencies to bring the test masses to the operating point [141], and this therefore requires an HV amplifier capable of providing close to 750 V limit set by the vacuum system. The maximum displacement each ESD can apply to each ETM as a function of frequency given this limit is shown in figure 6.2, calculated using the state-space model for the ETMs discussed in chapter 5.

Given the experiment's constraints in terms of noise, output voltage and signalling we

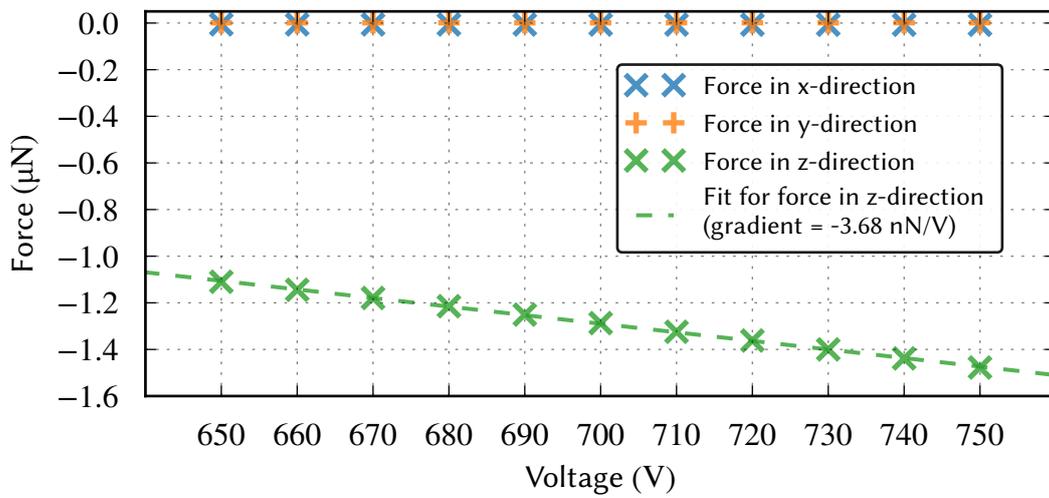


Figure 6.1: Simulations of the actuation force produced by the proposed ESD design upon a 100 g cylindrical test mass of diameter 48.6 mm and depth 24.5 mm resembling that of the Sagnac speed meter experiment's ETMs. The plate separation and the position of the mirror with respect to the plates influence the level of force produced. Greatest force coupling is produced when the mirror centre of mass is aligned to the edge of the plates and the plates are as close as possible to the mirror without touching. Although the force follows a quadratic relationship with voltage given by equation (6.2), the range shown is in a highly linear region.

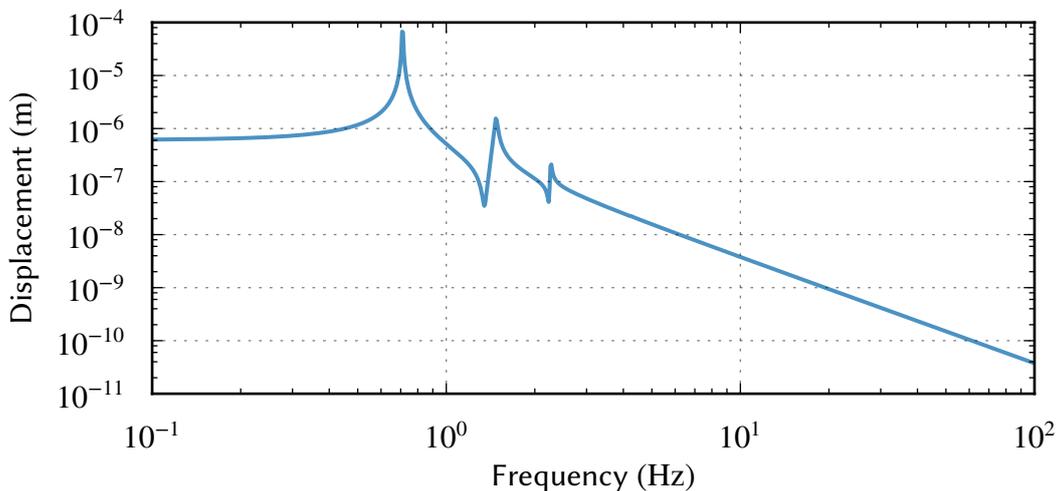


Figure 6.2: Maximum ETM displacement the ESD can create as a function of frequency with a potential difference of 750 V. This is produced by multiplying the force to displacement transfer function calculated for the ETMs using a state-space model by the force produced by the ESD at maximum output. This assumes that the output is entirely concentrated in a sine wave at each particular frequency, but in reality the feedback signal will contain many frequency components and so the displacement applied at any one frequency will be reduced.

consider a bespoke HV amplifier design. The next sections discuss the design, construction and testing of such a device.

6.3 High voltage amplifier design and implementation

Given that there are four ETMs in the arm cavities of the Sagnac speed meter experiment there will be four ESDs and so there will need to be four electronic feedback channels. The feedback signal to be applied to each ESD via an amplifier will be supplied by the CDS system as discussed in section 4.2.4.1. Given that a high voltage power supply, safety interlocking and monitoring infrastructure can be shared across the four channels, it is sensible to develop a single electronic amplifier circuit supporting four channels.

6.3.1 Differential sending and receiving

For the control of the amplifier and its high voltage output it is beneficial to utilise *differential* sending and receiving to allow for signals to be transmitted across the laboratory with minimal noise pick-up.

6.3.1.1 Control input

In CDS, the signal S is sent from the digital to the analogue domain via ADCs where it is split into two channels, A and B , containing the same signal but with opposite sign. These signals are sent to the amplifier in a two-core cable, and during transmission the signals are susceptible to noise pick up n_A and n_B , respectively, arising from electromagnetic interference; each channel then contains contributions from the signal and noise:

$$A = S + n_A, \quad (6.3)$$

$$B = -S + n_B. \quad (6.4)$$

We can represent these noise sources in terms of common and differential modes at the amplifier input, $n_{(+)}$ and $n_{(-)}$, respectively:

$$n_{(+)} = n_A + n_B, \quad (6.5)$$

$$n_{(-)} = n_A - n_B. \quad (6.6)$$

Standard op-amps contain *inverting* and *non-inverting* inputs and the output current is proportional to the difference between the two inputs. This subtraction results in the cancellation of common mode noise at the two inputs, so it makes sense to use op-amps to receive

the signals sent from CDS. An op-amp's ability to remove common mode noise is expressed as its *common mode rejection ratio* (CMRR), defined as the logarithm of the ratio of the amplifier's differential and common mode gains $G_{(-)}$ and $G_{(+)}$, respectively:

$$\text{CMRR} = 20 \log_{10} \left(\frac{G_{(-)}}{G_{(+)}} \right), \quad (6.7)$$

with the resulting number expressed in decibels.

The output of an op-amp in terms of its differential and common mode gain is:

$$S_{\text{out}} = G_{(-)}(A - B) + G_{(+)} \frac{(A + B)}{2}, \quad (6.8)$$

and given the definitions of A and B above, this becomes:

$$S_{\text{out}} = G_{(-)}(2S + n_{(-)}) + G_{(+)} \frac{n_{(+)}}{2}. \quad (6.9)$$

Given equation (6.7), an op-amp with high CMRR implies $G_{(-)} \gg G_{(+)}$, and so equation (6.9) shows that an op-amp with high CMRR will result in only differential noise having any significance at the output. The unsuppressed differential noise is kept to a minimum by ensuring that A and B are transmitted through the same shielded cable.

A standard low-noise op-amp such as the OPA277 provides CMRR of 120 dB at 100 Hz; this would result in $n_{(+)}$ being suppressed by a factor 10^6 which should be sufficient for a standard laboratory environment.

6.3.1.2 High voltage output

Only the potential difference between the plates of each ESD determines its force actuation, not the magnitude of the voltage with respect to ground at any one plate. To suppress the effect of noise from the op-amps used to produce the high voltage signals, however, it is best to create the potential difference V by producing balanced contributions of $\frac{V}{2}$ at the cathode and $\frac{-V}{2}$ at the anode. This is to reduce the effect of current noise, which depends on the output voltage, from the HV op-amp used to generate the signal. The noise from identical HV op-amps is then suppressed by a factor $\frac{\sqrt{2}}{2}$ between the plates. This requires two HV op-amps per output channel.

6.3.2 Switchable dewhitening

During lock acquisition the majority of the ESD's signal will contain corrections for low frequency oscillations created by seismic noise and radiation pressure effects in order to

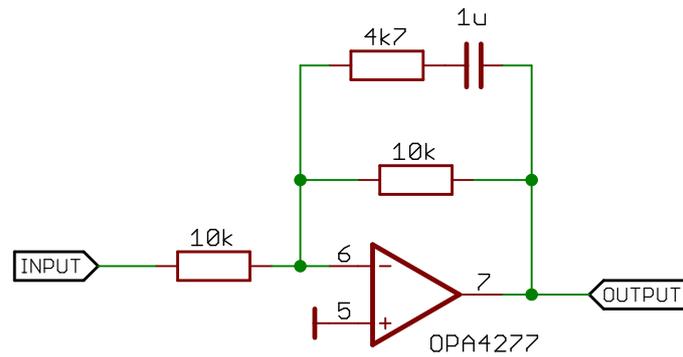


Figure 6.3: Active inverting dewhitening circuit. This filter uses an inverting op-amp with parallel feedback resistors to achieve the desired frequency response. At low frequencies, the capacitor's impedance is high and so the feedback path is dominated by the impedance of the 10 k Ω resistor and the gain is 1. At high frequencies the capacitor produces very low impedance and so the feedback path's total impedance is the equivalent resistance of the parallel 4.7 k Ω and 10 k Ω resistors and the gain is then $\frac{10\text{ k}\Omega}{3.2\text{ k}\Omega} \approx -10\text{ dB}$. The simulated transfer functions for one and two of these filters is shown in figure 6.4.

damp the residual test mass motion to the required level. As the operating point is neared the signal will have to contain feedback at higher frequencies to reduce the displacement noise sources shown in table 5.1 to reach the target sensitivity.

Whitening and dewhitening was introduced in section 5.2.8.1. To achieve strong actuation at low frequencies during lock acquisition the input to the HV amplifier from CDS will be unwhitened to achieve maximum low frequency signals at the expense of higher frequency noise from the DACs. When low frequency noise is suppressed and greater locking precision is desired in order to reach the low noise operating point, a whitening filter on CDS will be engaged to enhance the higher frequency feedback above the DAC noise. To compensate for this digital filter an equivalent dewhitening filter will be required on each analogue HV channel as shown in figure 5.20. For maximum flexibility we intend to include two dewhitening filters for each HV channel, and due to the speed at which the control system will have to respond to high frequency perturbations these filters will have to be switchable by an electronic means.

The simulations conducted in chapter 5 suggest that the whitening at the output from CDS should provide amplification by a factor of 10 above 100 Hz, which means the dewhitening should reduce the magnitude by the same amount. Simulating one or two 10 dB active dewhitening circuits as shown in figure 6.3 with the circuit simulation tool *LISO* the resulting transfer functions are shown in figure 6.4. The difference in phase between the two filters can be compensated with a sign flip applied to the feedback signal during filter transitions.

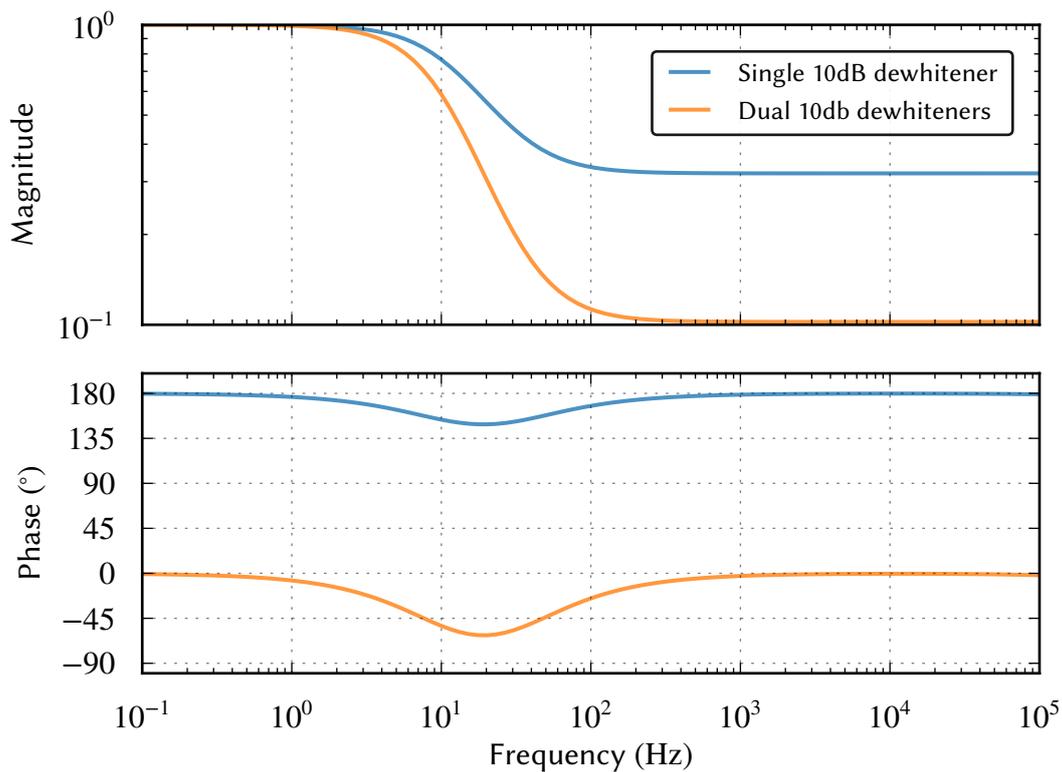


Figure 6.4: Frequency response of the dewhitening filters simulated with LISO. Each dewhitening filter provides -10 dB gain at high frequencies and so the combined pair produces an overall high frequency gain of -20 dB $\approx 10^{-1}$.

6.3.2.1 Digital switching electronics

A series of digital outputs from CDS can be used to control the HV amplifier's dewhitening filters. The Contec DO-32L-PE output card provides a 32 channel binary switch with an equivalent schematic shown in the left side of figure 6.5. The control signal from CDS for a particular output is inverted and attaches to the negative input of an optocoupler which acts as a relay without an electrical connection between the input and output. The positive input is attached to a voltage supply such that a digital output of 1 results in a closed circuit once inverted. An output of 0 is inverted to 1 and so there is no potential difference to close the optocoupler's circuit. The optocoupler's output in turn connects to the base of a transistor which controls current flow between the collector and emitter. Once the potential difference between the transistor's base and emitter exceeds around 0.7 V, current is allowed to flow from the collector to the emitter, forcing the potential difference between collector and emitter to be zero.

The analogue input to the HV amplifier for a particular channel is split into two with one passed through the dewhitening filter and the other passed through unfiltered. These two signals, the "dewhitened" and "normal" inputs, respectively, as shown in figure 6.5, form the

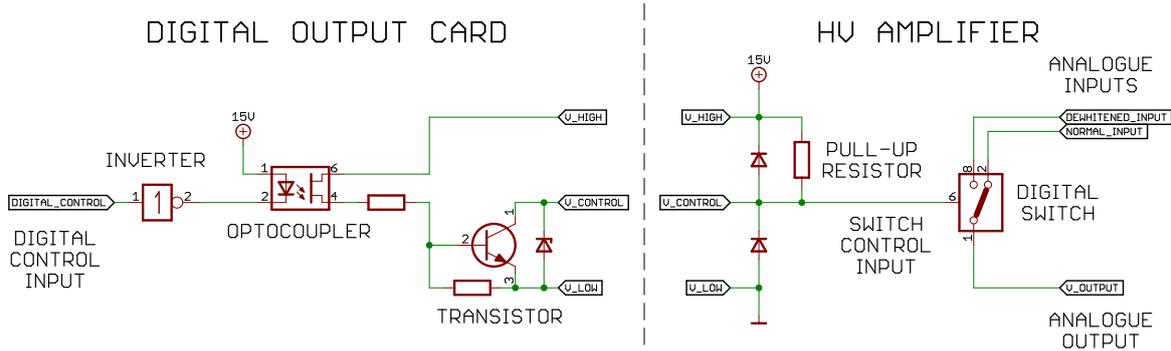


Figure 6.5: Digital signalling between CDS and a channel within the HV amplifier. The digital control signal operates an optocoupler which allows current to flow from the V_{HIGH} voltage reference to a transistor which operates the analogue control signal V_{CONTROL} . This signal is used by the CMOS switch in the HV amplifier to select either the dewhitened or unfiltered inputs.

CDS software logic level	V_{CONTROL}	Dewhitening status
0	V_{HIGH}	On
1	V_{LOW}	Off

Table 6.2: Truth table for digital switching of dewhitening filters in the HV amplifier, showing the effect that a software logic level in CDS has on the dewhitening of the input signal for a particular channel in the HV amplifier.

poles of a *complementary metal-oxide semiconductor* (CMOS) switch, selected for its rapid switching speed. The analogue output from this switch is either the dewhitened or normal input, as determined by a threshold sensor at the switch's control input. If the control voltage is near to 15 V the output is the dewhitened input, and if it is near to ground the output is the normal input. The voltage drop across the CMOS is negligible. The switch control input is connected to one of the digital outputs of the CDS card via a shielded transmission line, represented by the dashed vertical line. Also transmitted between the circuits are the reference voltages V_{HIGH} and V_{LOW} , set to 15 V and ground, respectively. The software on CDS determines whether the control voltage is set to V_{HIGH} or V_{LOW} . Due to the *active low* operation of the digital output card, where the application of a current to the optocoupler's LED (corresponding to a digital control input of 1) results in the control voltage becoming V_{LOW} , the dewhitening filter threshold signal V_{CONTROL} must be referenced to V_{HIGH} . This is achieved in the circuit through the use of a pull-up resistor. When dewhitening is desired, an output of 1 at the optocoupler results in a low-resistance path between the control signal and ground and so the control signal within the amplifier becomes V_{LOW} . A truth table for the logic between CDS and the HV amplifier is shown in table 6.2.

The electronics shown in figure 6.5 are for a single dewhitener. This configuration must be repeated twice for each of the channels in order to provide available dewhitening of 0 dB, 10 dB or 20 dB.

6.3.3 Choice of high voltage op-amp

Due to the nature of the load the amplifier does not need to drive a significant current, but the voltage noise it produces must be significantly lower than the displacement requirement for the experiment in order for it not to limit its sensitivity. The lowest noise high voltage amplifier integrated circuits available tend to be MOSFET-type op-amps. The choice of device tends to be motivated by the bandwidth, maximum output voltage and noise of each particular model.

The *gain-bandwidth product* specifies an op-amp's open-loop gain as a function of the bandwidth it is able to provide it over, and this figure is derived from the speed at which the op-amp's output is able to react to a change in its input (its *slew rate*). The full output voltage is not provided at the unity gain frequency, however, and so a more useful figure of merit is the bandwidth over which the maximum output can be provided. For the Sagnac speed meter experiment it is expected that radiation pressure and thermal noise will require fast corrections in the kHz range, and to avoid becoming limited by the device's slew rate at higher frequencies (which appears as phase lag on a plot of the frequency response) it is reasonable to require a bandwidth of at least 20 kHz.

The required dc op-amp gain should be known ahead of time in order to fully estimate the effect an op-amp's noise will have on the experiment. The maximum CDS input voltage is ± 10 V (differential) and so to achieve the maximum output voltage requirement of ± 350 V with the greatest dynamic range the amplifier circuit's gain should be at least $\frac{\pm 10 \text{ V}}{\pm 350 \text{ V}} = \pm 35$.

Finally, the quiescent current drawn by each op-amp determines the amount of heat it will produce, and since the HV amplifier will be used continuously during interferometer operation this parameter should be as low as practical in the chosen model.

Table 6.3 shows the aforementioned parameters for some popular op-amp models produced by *Apex*. The models shown have sufficient output voltage and identical input noise. The PA89's bandwidth is limited and the full output voltage is not available beyond 7 kHz, while the quiescent power in the PA94 and PA98 is high enough to warrant challenging heat sink requirements. In this case the optimal choice of op-amp is the PA95, which provides enough bandwidth and for which only a passive heat sink will be required.

6.3.4 Amplifier signal path

The constructed amplifier signal path is shown for a single channel in figure 6.6. The inputs IN+ and IN- are differentially received and converted to a single-ended signal which is then passed through the two dewhitening stages controlled by digital switches. The signal is then split into two parts, with one part being inverted, and these become the inputs

	PA89	PA94	PA95	PA98
Maximum output	1140 V	900 V		450 V
Bandwidth @ 400 V output	7 kHz	90 kHz	30 kHz	60 kHz
Input noise @ 100 Hz	6 nV $\sqrt{\text{Hz}^{-1}}$			
Quiescent power @ ± 400 V	3.8 W	14.1 W	1.3 W	17.5 W

Table 6.3: Performance specifications for various HV op-amps. All of the models listed are manufactured by Apex and the values have been obtained from their respective data sheets. The PA95 was selected due to its low quiescent power for the desired voltage range, which in turn eases the requirement for heat sinking, and the bandwidth is sufficient for the Sagnac speed meter experiment.

to the PA95 HV op-amps. The inverting and non-inverting inputs on each op-amp are bridged with diodes to ensure that the voltage difference is less than ± 0.7 V; this is to prevent accidental overvoltage at the inputs. Although the differential input range of the PA95s is ± 20 V, the ± 0.7 V limit is sufficient to reach the required bandwidth in this application, and avoids potential damage to the costly op-amps from incorrect input. A 1% voltage divider at each of the differential outputs provides a signal that is within the input range of CDS for the purpose of output monitoring. This so-called *monitor* provides a reference channel for the controller in order to assist with the calibration of the device as well as a possible error signal extraction point for a feedback loop intended to suppress amplifier noise. The resistors near pins 7 and 8 of each op-amp provide current limiting, discussed in section 6.3.5.1. The capacitor placed between pins 4 and 6 on each amplifier provides compensation for the phase of the op-amp at high frequencies, and its value is determined from a look-up table provided by the manufacturer.

6.3.5 Practical and safety features

As the HV amplifier handles potentially lethal current and voltage, a number of additional features beyond the signal circuitry are present.

6.3.5.1 Current limiting

At the output of each HV rail there are 47 k Ω series resistors which passively limit the current on each rail to around 10 mA. The current limit does not impact the driving of capacitive loads, but ensures that the current produced by the HV amplifier is not lethal. Given this series resistance the HV transmission lines to the vacuum system must be kept short (< 3 m) to prevent capacitance in the cable from creating a low-pass response within the bandwidth of the experiment.

The resistor placed between pins 7 and 8 also acts as an *active* current limit. The voltage drop across this resistor is used by each PA95 to limit its output current. The combination

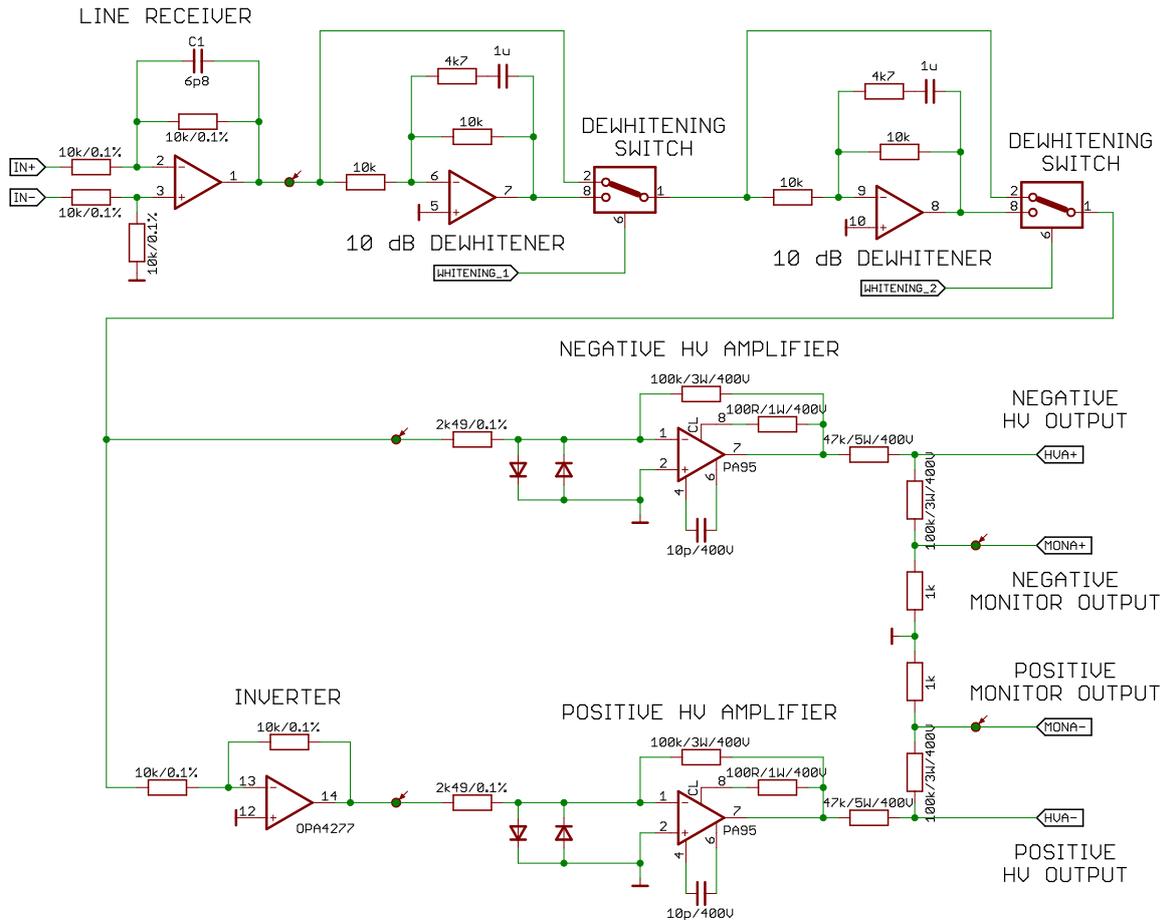


Figure 6.6: Schematic of a single high voltage amplifier channel. The input from CDS is differentially received, dewhitened and then split into differential outputs that are amplified by PA95 op-amps. A pick-off reads 1% of the output voltage for calibration and noise projection purposes.

of the two current limiting features ensures that a single fault will not lead to lethal output current, and as such this feature of the circuit meets European standard EN 61010 [155].

6.3.5.2 Soft-start

Capacitors are present upon the HV supply lines to filter ac noise, and these must be charged when the device is switched on. Normally these capacitors would present very little impedance to the power supplies and so a large initial current would be drawn potentially damaging components in its path. The simplest technique to prevent this from happening would be to put resistors in the path of the power supplies, but resistors that would charge the capacitors at a safe rate in a short time would also dissipate a lot of power and require additional heat sinking. Instead, we use a *soft-start* mechanism which controls the current flow during the charging of the capacitors. The HV amplifier's on-off switch operates optocouplers which allow current to flow into the circuit. Initially, when the circuit's HV capacitors are discharged, the current on each HV rail is limited by the parallel

5 k Ω resistors. A 2.5 % pick-off from each HV rail is compared to a reference 5 V potential at an op-amp, the output of which operates a second optocoupler on each rail. When the voltage surpasses 200 V the power the 5 k Ω resistors dissipate is 8 W, which is near their limit. At this point, the capacitors are almost fully charged and the pick-off voltage surpasses the 5 V reference and so the op-amp's output operates the optocoupler to open up a low-resistance path that bypasses the 5 k Ω resistors and prevents them from overheating. The soft-start circuit is shown in figure 6.7. The capacitors to be charged by this circuit would be attached between the +HV and -HV rails and ground.

6.3.5.3 Pressure and temperature interlock

The breakdown voltage of the plate capacitors as a function of pressure, given by Paschen's Law, has a minimum in the region of 10^{-1} mbar to 10^1 mbar depending on the separation and geometry of the anode and cathode. If the gas pressure and plate separation are favourable, it is possible for arcing to occur between the anode and cathode of each ESD. Estimates of the Paschen curves for parallel plates with varying separations in nitrogen are shown as a function of pressure in figure 6.8. Apart from arcing, related effects such as creep-*age*—tracking of charge across insulative surfaces—can lead to arcing at voltages above 50 V in low vacuum, though the use of highly insulating material such as ceramics can help to mitigate this risk almost entirely [155].

Although the use of high voltage plate capacitors is in general safe at both atmospheric pressure and high vacuum, the act of pumping gas out of the vacuum system passes through pressures at which arcing and surface tracking can occur. To prevent this possibility a cut-off function was implemented within the circuit to prevent HV output unless a control signal is supplied (see figure 6.9). The switching mechanism uses the same CMOS switches as the digital dewhiting filters shown in figure 6.5, with the threshold signal being generated by CDS based on a signal from a separate pressure monitor input. Additionally, as a temperature fail-safe for the amplifier components, temperature sensors are present within the enclosure which operate threshold switches able to remove the supply to the HV op-amps using the same mechanism as the pressure interlock. The outputs from each interlock are sent to an AND gate, and only if both the pressure and temperature switches are in a safe state will the HV op-amp's supplies switch on. The interlock circuit is shown in figure 6.9.

6.3.6 Transfer functions and noise measurements

The monitor output provides a means of measuring 1 % of the full HV output with a signal that is within the input range of CDS. This allows transfer functions from the input to the

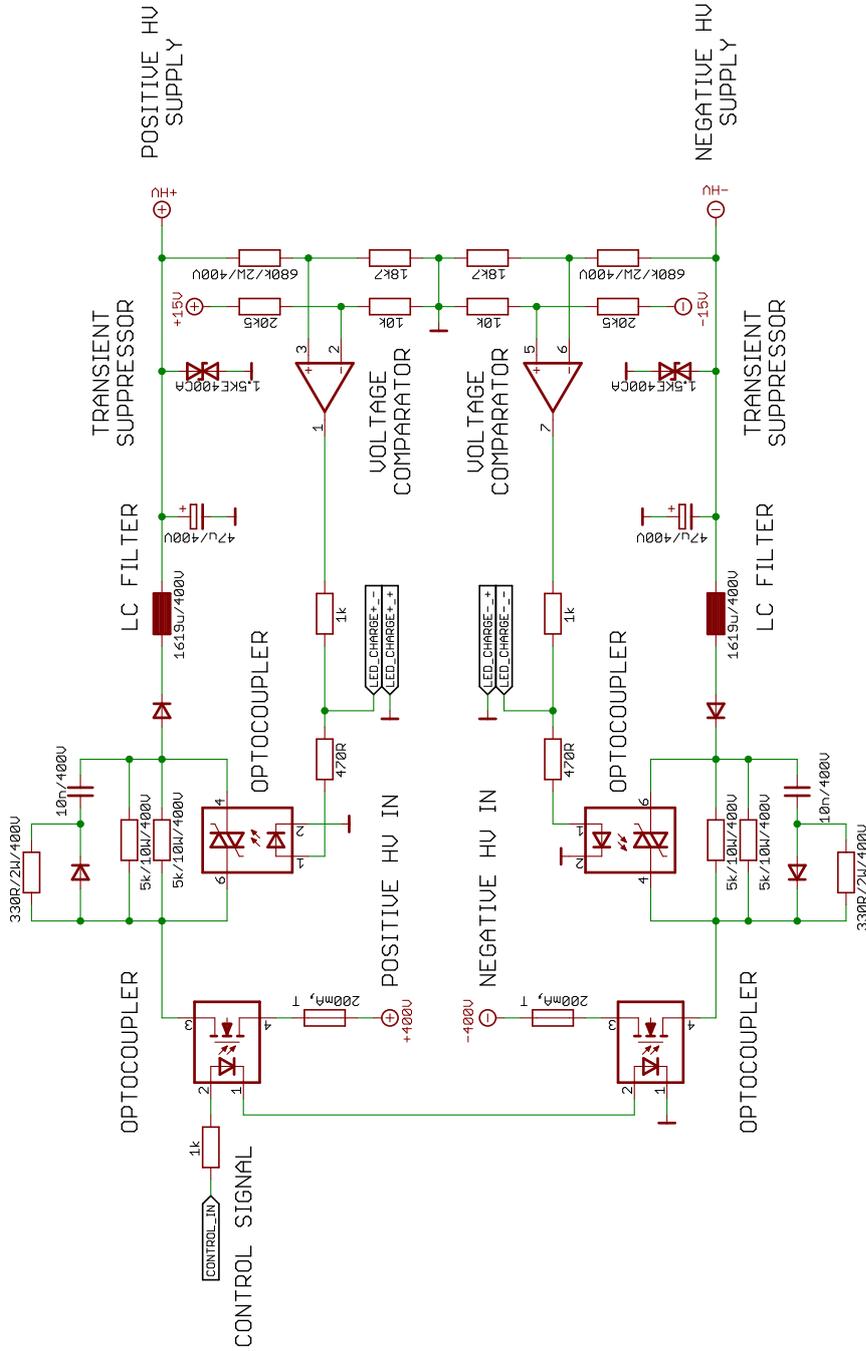


Figure 6.7: Soft-start mechanism to prevent excessive power supply current draw. This is based on a design for the AFI 10 m prototype by Andreas Weidner. The leftmost optocouplers control whether current is allowed to flow from the power supplies given a control signal from an on-off switch on the enclosure. After the electronics are switched on, the current flow is initially limited by two 5 kΩ resistors. Any ac signal content above around 500 Hz is filtered by the presence of a low-pass LC filter formed of the 1619 µH choke and 47 µF capacitor (the low-pass corner frequency decreases further when additional capacitive load is attached). A 2.5 % pick-off from each HV rail is compared to a ±5 V reference, with the resulting difference fed back to a second pair of optocouplers (different models from the first due to their lower duty cycle) which gradually open up a low-resistance path on each supply rail. Without these optocouplers, the power the 5 kΩ resistors would need to dissipate would be considerable at maximum voltage. As the supply voltage increases beyond around 200 V, the resistance through the optocoupler is reduced far below the 5 kΩ resistors and so the power dissipated in the resistors is negligible.

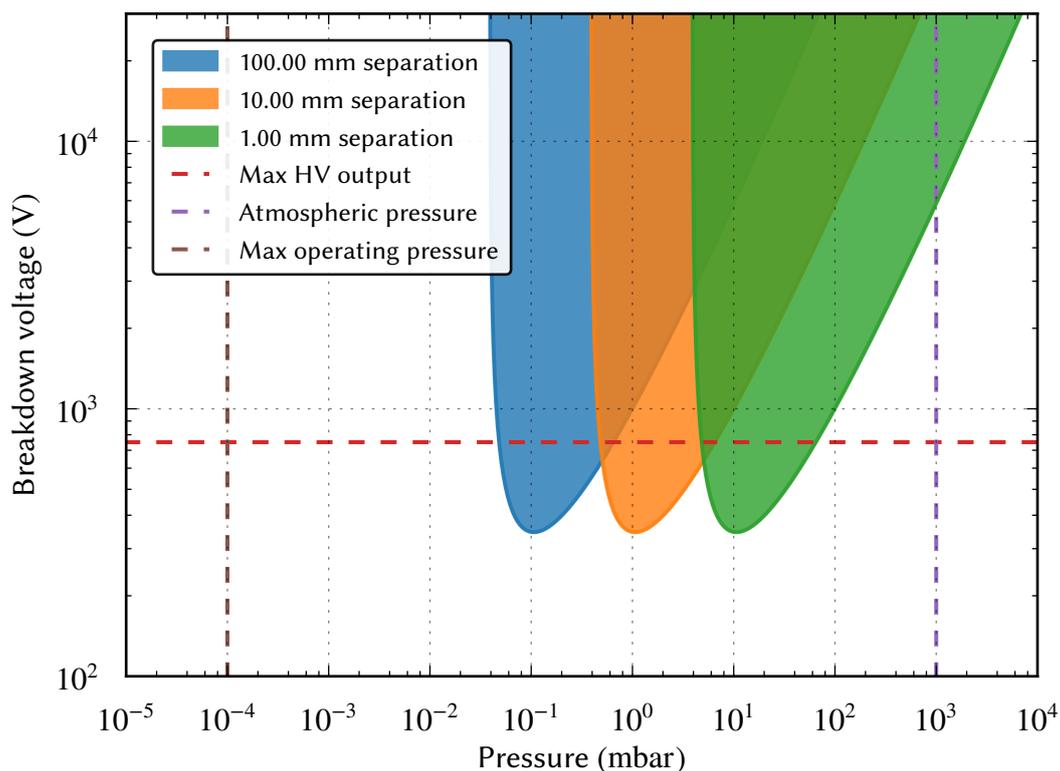


Figure 6.8: The minimum breakdown voltage between the two plates of the ESD for different separations. This is calculated using Paschen’s Law, assuming nitrogen gas and a flat plate geometry. The real effect is a lot more complicated than this model, but the steep slope at lower pressures shown here indicates that the voltage created by the HV amplifier for the ESDs will avoid problems associated with arcing as long as the vacuum system is operated at atmosphere or high vacuum.

output of each channel to be calculated in software to assist in the eventual calibration of the experiment. The HV amplifier should provide gain to the input signal without applying ac filtering beyond the desired dewhitening within the intended bandwidth of the experiment. Similarly, the noise measured at the output should not be significant enough to affect the sensitivity of the interferometer by applying force noise to the test masses. The following subsections describe measurements to assert these characteristics.

6.3.6.1 Swept sine response of each channel

Transfer functions for the HV amplifier can be measured by injecting a known signal into one channel and measuring the corresponding output. Figure 6.10 shows the *swept sine* response, calculated by injecting a sine wave at a given frequency, measuring the output signal and dividing it by the injection, and repeating this process over a given frequency range. These measurements were made without the digital dewhitening switches disengaged such that they provide 20 dB of low-pass filtering. The figure shows the response expected from the predictions made by LISO shown in figure 6.4, and the measurements

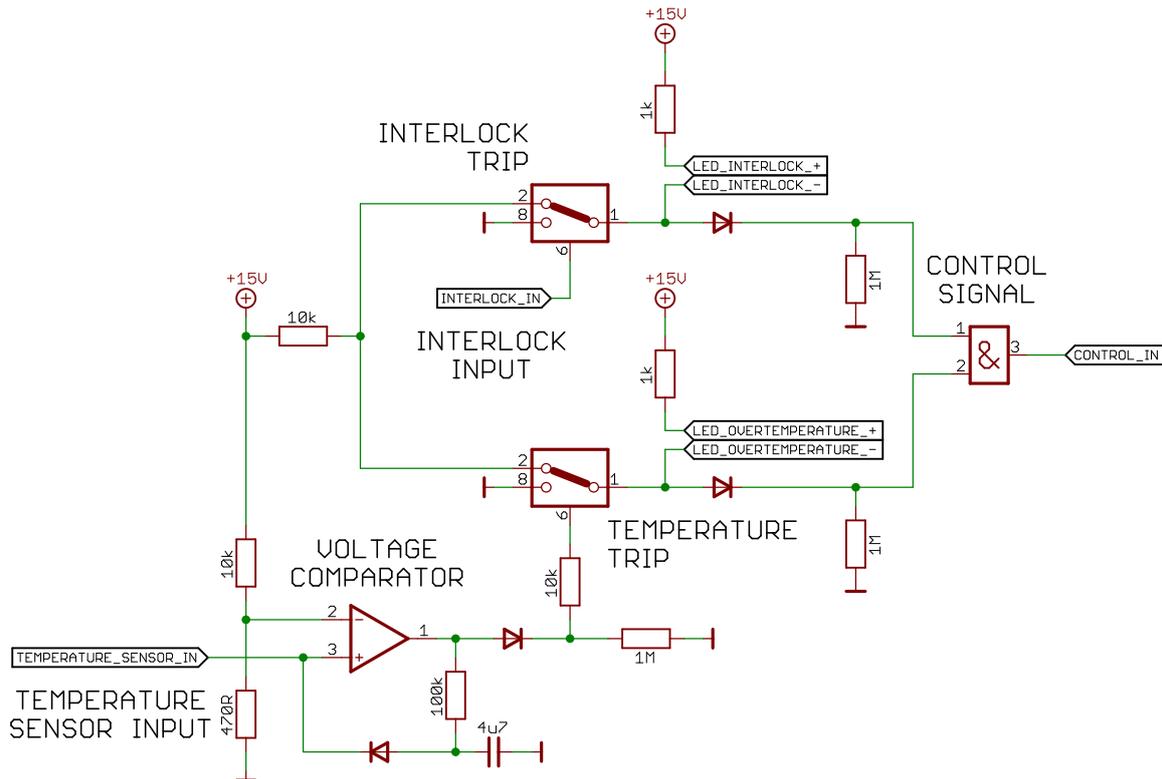


Figure 6.9: Pressure and temperature interlock circuit. The digital interlock signal from CDS operates one switch and the output from a temperature sensor threshold switch operates another. Only if both switches output 15 V will the control signal used to operate the HV supplies switch on.

agree once the filtering effect of the CDS anti-aliasing filters are taken into account.

6.3.6.2 Response with and without dewhitening

Figure 6.11 shows swept sine measurements of the first channel with the dewhitening filters in various states: both on, the first on and the second off, the first off and the second on, and both off. The measurements match predictions and the response with both dewhiteners off is flat as intended across the measurement band.

6.3.6.3 Coherence between channels

The channels should be isolated from one another such that a signal injected at one input does not appear at the output of another. Op-amp power supply filtering is implemented using decoupling capacitors, inductors and diodes such that there should be minimal cross-coupling between the channels. Figure 6.12 shows the coherence for each channel to each other channel, measuring whether the output signal has the same phase angle as the input signal. Coherence of 1 indicates that there is causal coupling between the two channels,

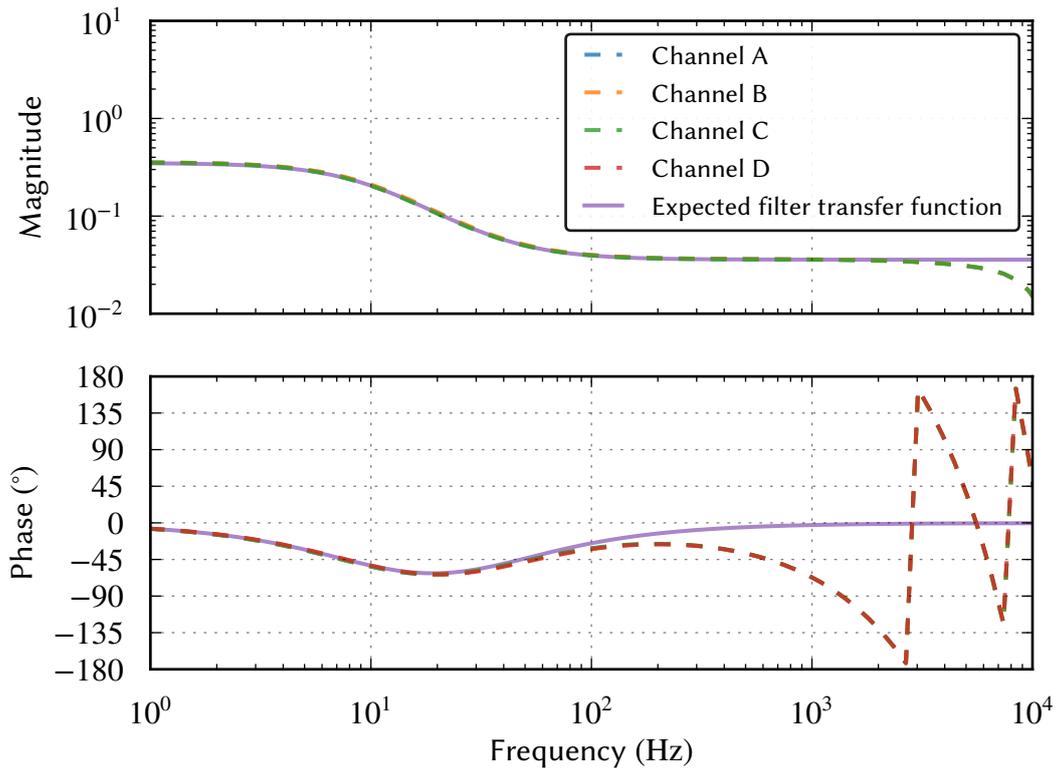


Figure 6.10: Second amplifier transfer functions with dewhitening enabled. The expected performance of the dewhitening filter from theory is shown in purple alongside the transfer functions of each channel. The curves agree closely, showing that the implemented filter operates as expected. The mismatch at high frequency is caused by the anti-aliasing filters implemented in CDS, which aggressively filter signals above a few kHz.

whereas coherence less than around 0.5 is expected from statistical random noise processes. The results confirm that a high level of isolation between channels has been achieved.

6.3.6.4 Output voltage noise

The noise at the output of the first channel of the HV amplifier is shown in figure 6.13. This was calculated via the HV amplifier's monitor with the measured noise being projected into effective HV noise by multiplying the signal by the inverse of the monitor pick-off fraction, 100. To measure the noise at frequencies comparable to the Sagnac speed meter experiment's cavity bandwidth a Stanford Research SR785 spectrum analyser was used; as shown in figure 5.11 the anti-aliasing and anti-imaging filters significantly reduce the input to and output from CDS above 9 kHz and so they are not practical for this measurement.

The following measurements were made:

- the HV amplifier's output noise when it is disconnected from CDS and has no input signal;

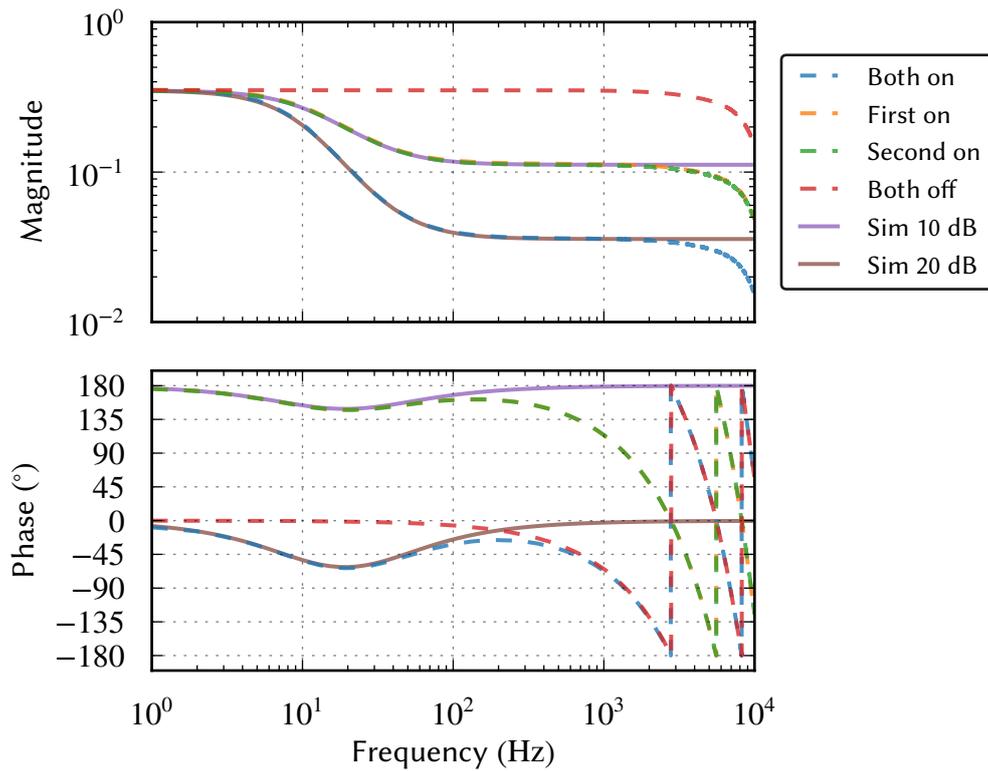


Figure 6.11: Transfer functions of the high voltage amplifier input to monitor output with the de-whitening filters on and off. The monitor output is a 1% pick-off from the main HV output, and so the gain is 0.4 instead of 40. The curves simulated with LISO agree exactly with the measurements within the bandwidth of CDS, and at higher frequencies the transfer functions are suppressed by anti-aliasing filters.

- the HV amplifier's output noise when it is connected to the CDS DACs with an input signal equivalent to maximum HV output;
- the spectrum analyser noise floor with a 50 Ω load.

Each of the spectral densities shown were produced using amplitude spectral density estimates of the time domain signal (see appendix B.2.1), and to avoid windowing effects as discussed in appendix B.2.3 the measurements were made with averages across three bands: 250 mHz to 200 Hz (16 averages), 4 Hz to 3.2 kHz (64 averages) and 128 Hz to 102.4 kHz (1024 averages). The measured monitor noise is above the input noise of the spectrum analyser in all cases.

The left y-axis shows the measured monitor noise. The right y-axis shows the noise projected into equivalent noise on the HV output. The noise is different in the case of zero (orange) or full (green) dc output not because of an input voltage dependency but rather due to the inclusion of the CDS DACs in the signal path in the latter case which contribute noise as discussed in section 5.2.7.1. Some large amplitude noise is present around 50 Hz and its first harmonic due to pick-up from the electricity mains on the power supply.

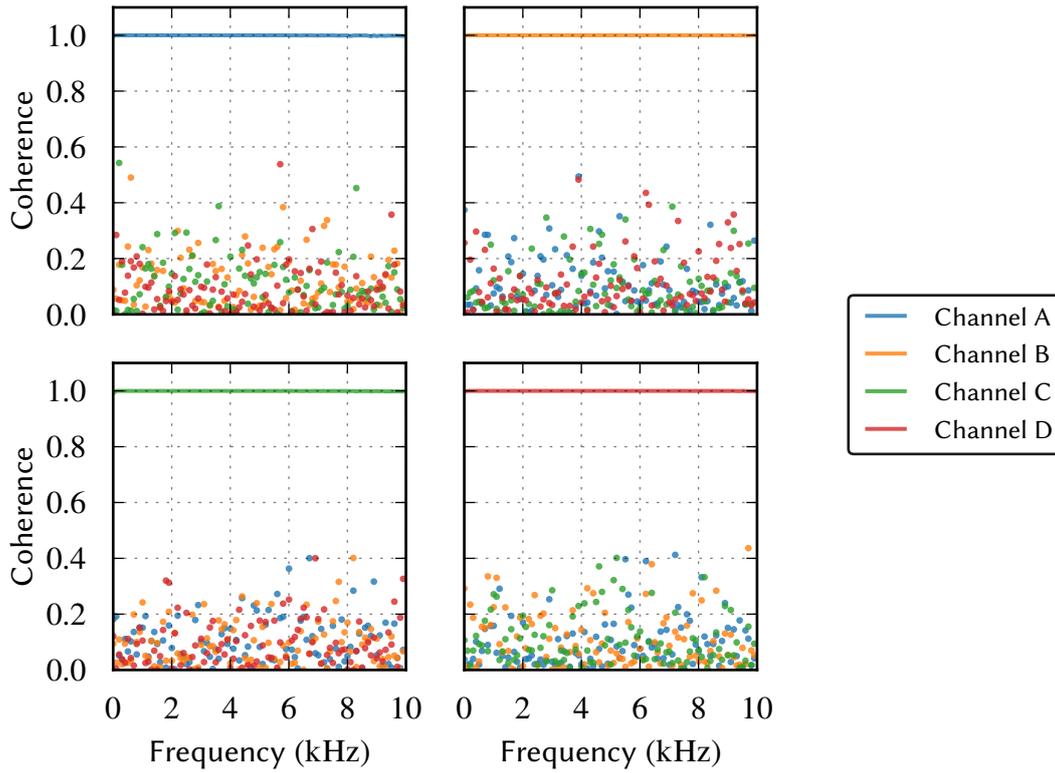


Figure 6.12: HV amplifier cross-channel normalised coherence. A swept sine was injected into each channel in turn while measurements of the output phase were made on all four channels. In each case, only the channel with the injection has coherence of 1, while other channels only show the effect from noise.

6.3.6.5 Effective displacement noise

As shown in equation (6.2) the ESD force depends on the square of the potential difference. We can rearrange it to show:

$$\frac{F(V_{\text{signal}})}{V_{\text{signal}}^2} = \frac{F(V_{\text{signal}} + V_{\text{noise}})}{(V_{\text{signal}} + V_{\text{noise}})^2}, \quad (6.10)$$

where V_{noise} is the HV amplifier's output noise. This shows that the force noise created by the HV amplifier will be greatest at maximum output signal, $V_{\text{signal}} = 750$ V. The effective force noise is simply the force produced due to the signal subtracted from the total force produced in the presence of signal and noise:

$$F(V_{\text{noise}}) = F(V_{\text{signal}} + V_{\text{noise}}) - F(V_{\text{signal}}). \quad (6.11)$$

The ETM displacement noise this would create is then simply the product of the force noise and the suspension transfer function, and this is shown in figure 6.14. The effective displacement noise is below the noise budget presented in chapter 5 within the measurement

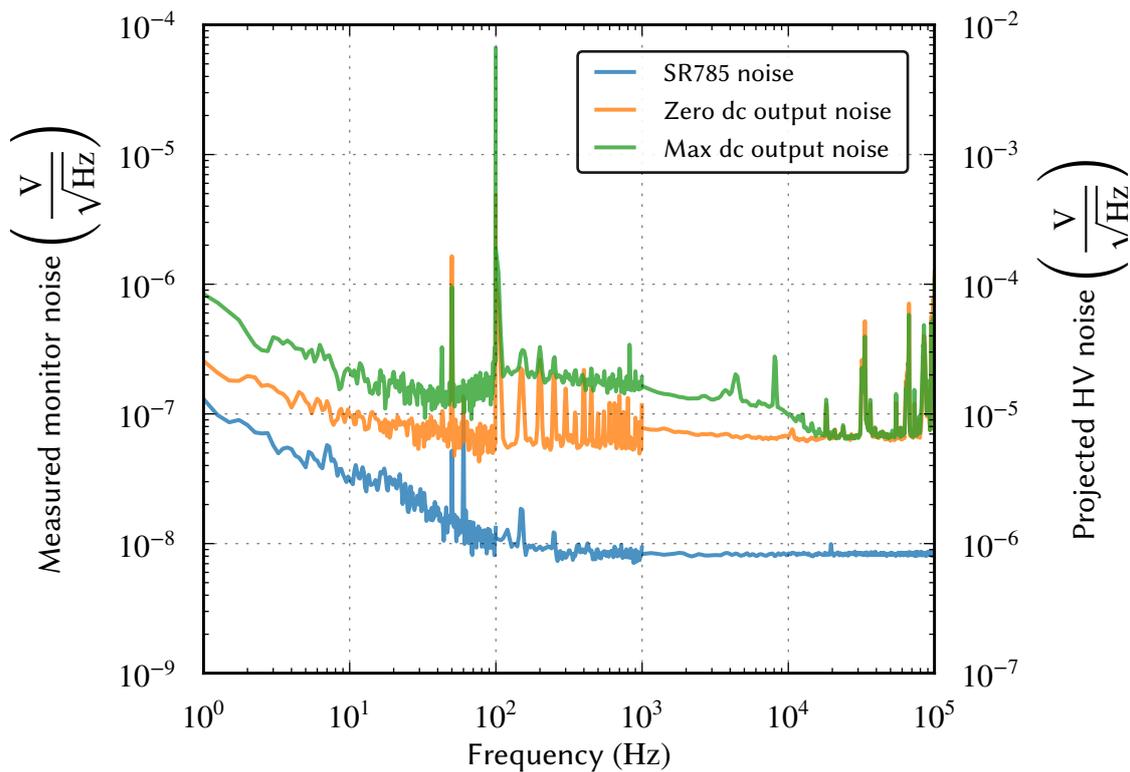


Figure 6.13: HV amplifier output noise. The noise of the amplifier was measured using the monitor output of the first channel, and this is shown in the green and orange traces using the left y-axis. The output noise projected from the monitor to the HV rails is shown on the right y-axis. The green trace represents the noise from the amplifier when the input is at maximum and includes the noise contribution from the DACs on CDS. The orange trace shows the output noise when CDS is not connected. The noise floor of the SR785 spectrum analyser used to make these measurements is shown in blue.

band, and given that this assumes the ESD is at maximum output the displacement noise when the interferometer is at the operating point, requiring only a fraction of its maximum range, will be lower. If the noise from the HV amplifier is found to be higher in practice, it should be possible to feed the monitor outputs back to their respective inputs in order to stabilise each channel.

6.4 Outlook

For the control of suspended test masses at high frequencies, parallel plate capacitor ESDs provide a low noise alternative to voice coils and their geometry can prevent clipping losses suffered in the use of metal comb ESDs. It is intended for parallel plate capacitor ESDs to be used as the high frequency actuators in the Sagnac speed meter experiment presented in Chapter 4, and in order to provide actuation at the required level an HV signal with suitably

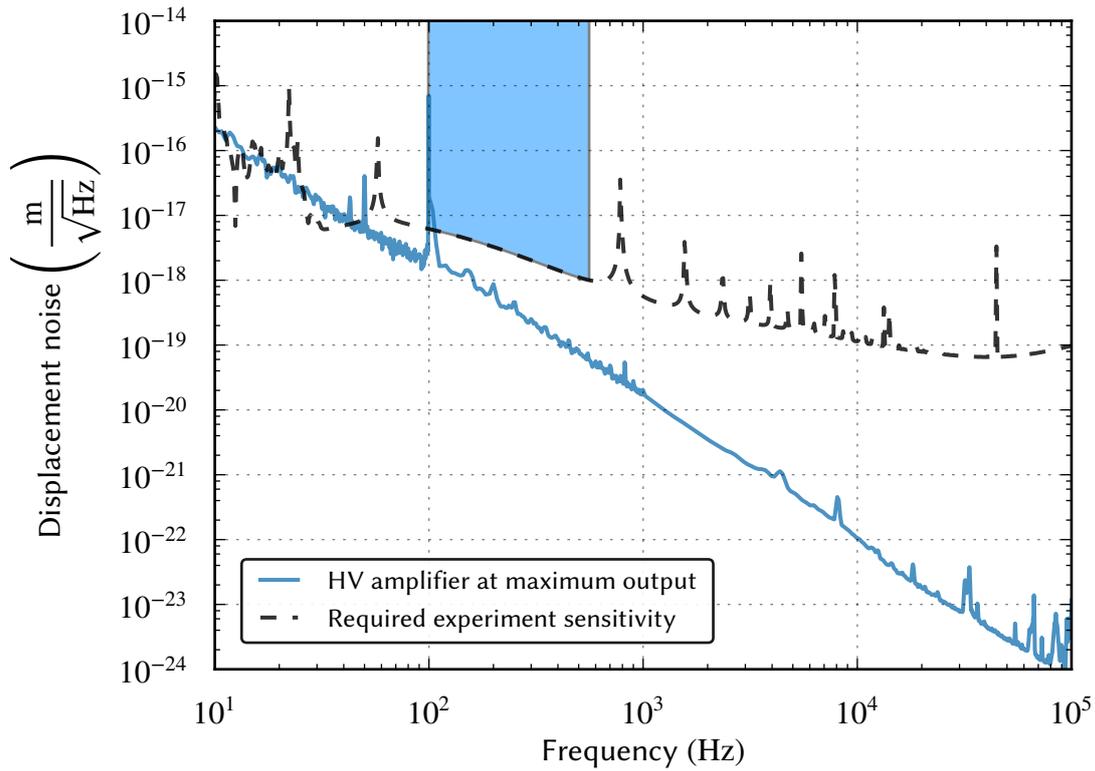


Figure 6.14: HV amplifier output voltage noise (without dewhitened) projected into effective ETM displacement noise. The force produced by the ESD is a function of the square of the voltage across its plates, and so the noise produced by the HV amplifier has greatest effect when the signal is at a maximum. This plot shows the displacement noise for maximum signal, and within the measurement band (blue shaded region) it is below the requirement shown in black taken from the analysis conducted in chapter 5. If necessary the HV amplifier’s noise could be reduced further by implementing a control loop between its monitor outputs and signal inputs.

high magnitude and low noise must be available. The technical design of an HV amplifier was presented which meets the experiment’s range and noise requirements, providing output up to 750 V with noise of around $20 \mu\text{V} \sqrt{\text{Hz}^{-1}}$ in the frequency band of interest.

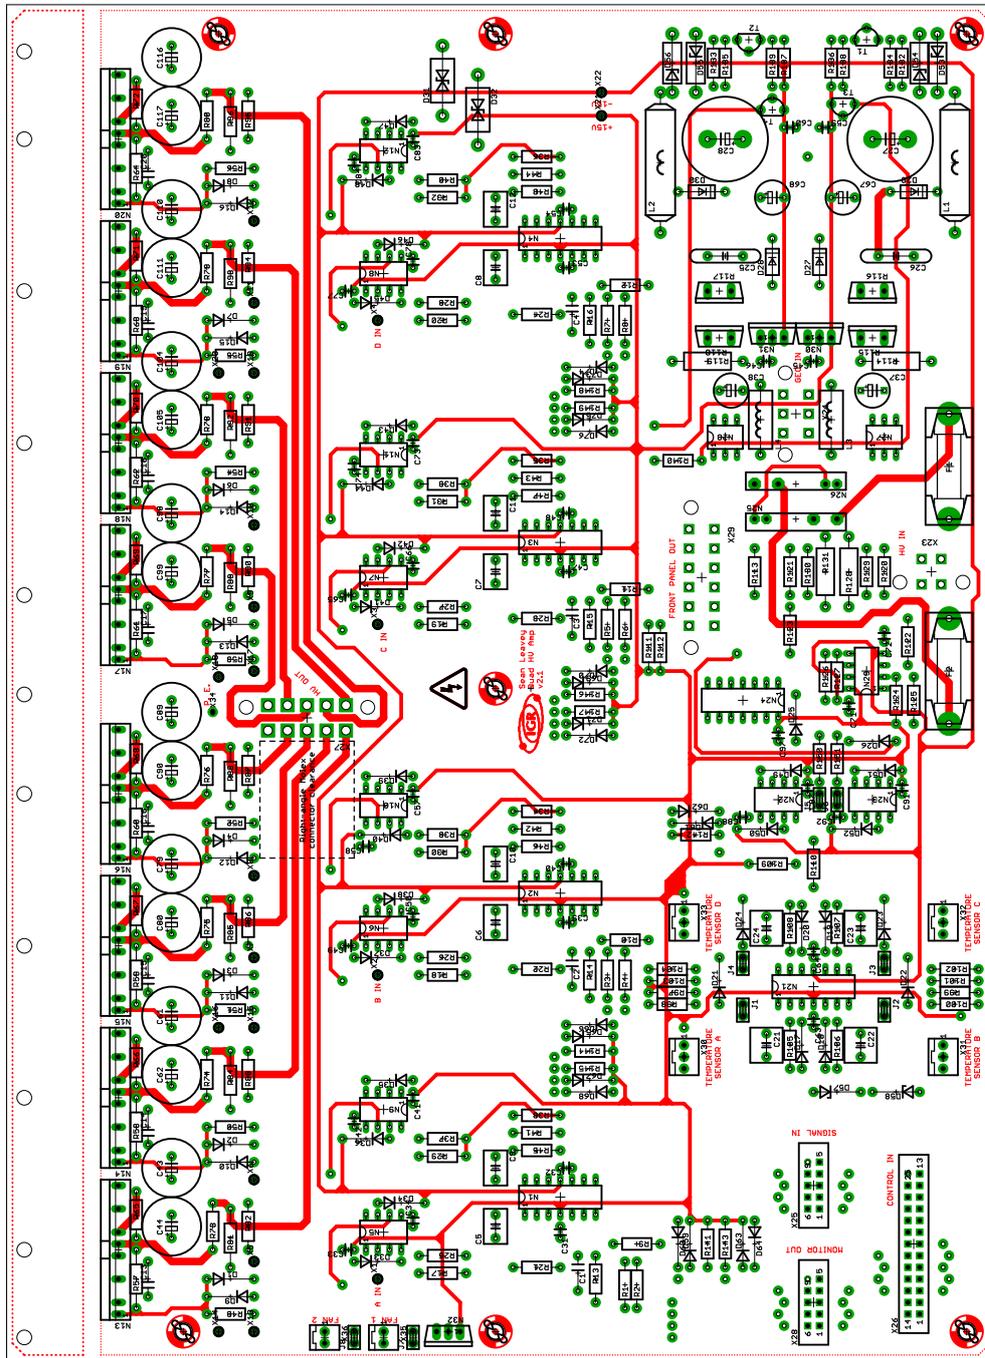


Figure 6.15: Top view of the HV amplifier printed circuit board showing the component and copper track layout. This side contains most of the components and all of the connectors that go to and from the enclosure. Areas not carrying signals are part of the ground plane, except for a relief at one side where the heat sinks for the PA95s are attached, isolated from the ground plane to minimise the risk of arcing.

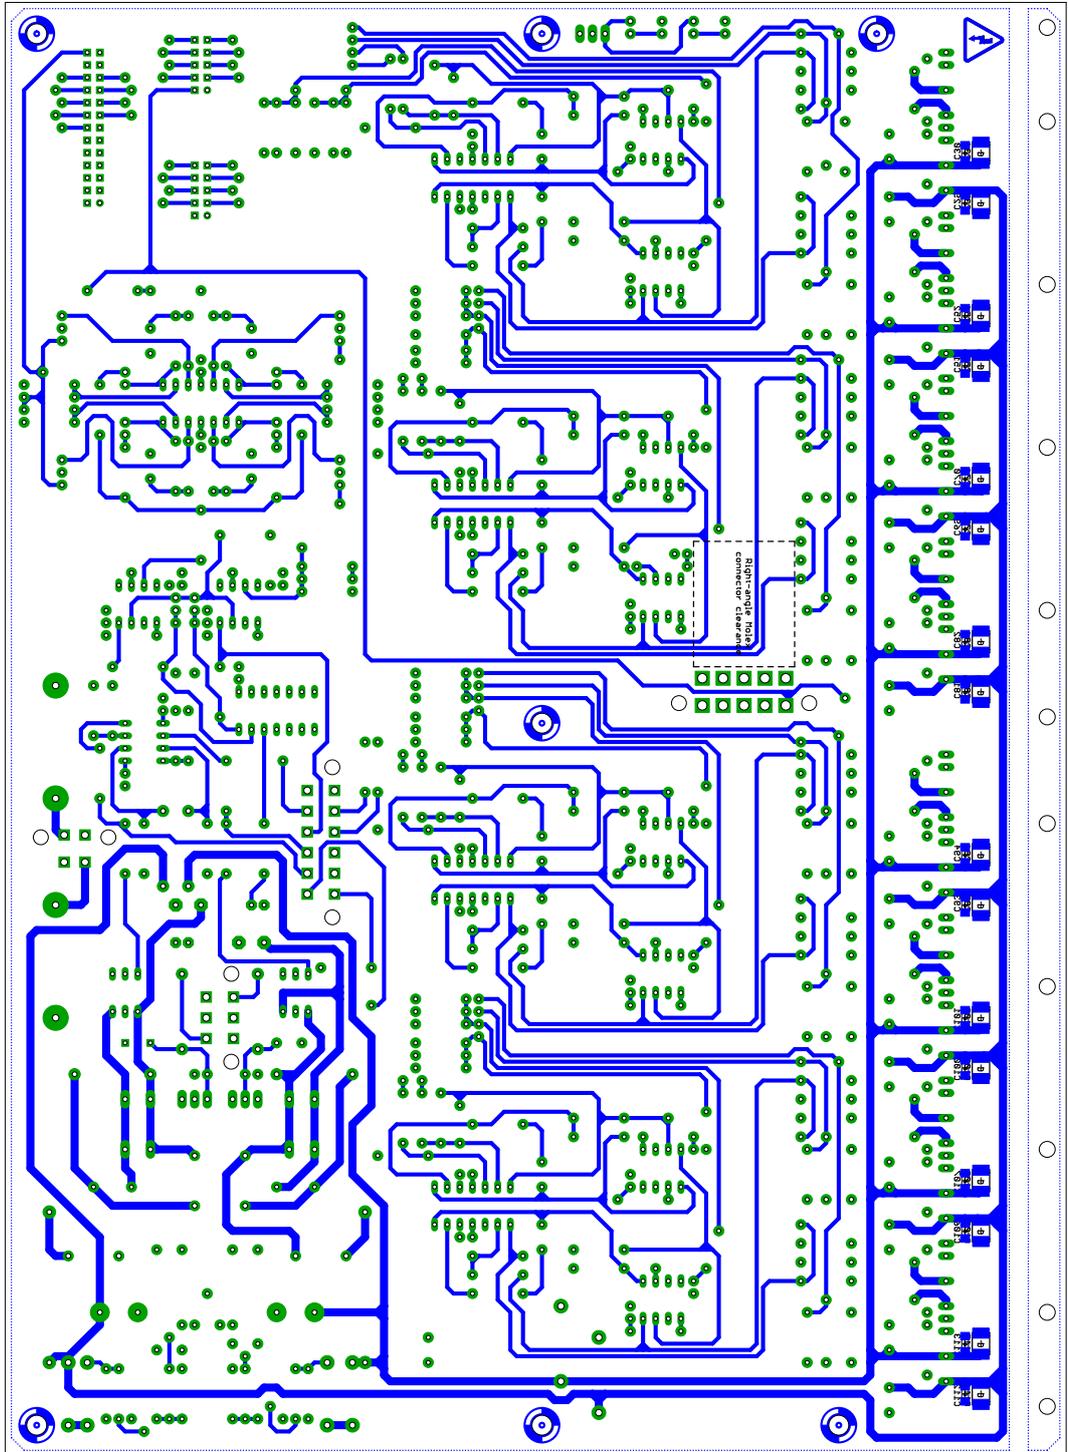


Figure 6.16: Bottom view of the HV amplifier printed circuit board showing the component and copper track layout. This side is where the top components are soldered to the copper tracks. Some surface mount decoupling capacitors for the supply rails of the PA95s are present on the bottom as they could not be placed as closely to the op-amps on the top side.

Chapter 7

Conceptual longitudinal sensing scheme for the low frequency Einstein Telescope detector

7.1 The Einstein Telescope facility

In 2011 a group of scientists primarily based in Europe completed a design study [16] to examine the infrastructure requirements for a gravitational wave observatory that pushes the Michelson interferometer topology to its limits, whilst making any newly built facility generic enough to allow for the implementation of new topologies as the state of the art evolves. In this study they laid out the expected improvements in technologies to mitigate fundamental and technical noise sources currently limiting the sensitivity of the second generation of detectors, with a number of practical differences to existing facilities.

The design for the Einstein Telescope targets an increase in sensitivity over Advanced LIGO and Advanced Virgo by a factor of 10 over a wide bandwidth. In order to surpass the sensitivity of the current detectors, which already expect to improve upon that of the first generation by an order of magnitude, a number of differences and improvements have been envisaged. The proposed Einstein Telescope facility composes six dual-recycled Fabry-Perot Michelson interferometers split between the three corners of a triangle with 10 km edges. The design exploits the geometry to implement interferometers with 10 km arms meeting at the three vertices to benefit from the colocation of multiple interferometers.

Seismic noise limits the sensitivity of current generation detectors below 10 Hz, and there are astrophysical advantages to being able to achieve good sensitivity at these frequencies [58], particularly in the ability to see the inspiral and merger of high mass black hole binary coalescences. The majority of spinning neutron stars discovered via optical tech-

niques have also had orbital frequencies below 10 Hz where fundamental noise limits the ability for the current generation of detectors to see such signals. Signals evolving with frequency, such as the binary black hole merger witnessed as GW150914 [1], which was seen for a few tens of ms, are present at frequencies around 2 Hz for hours. Having this extra observation time provides the possibility not only to better estimate the source parameters but also to track the signal evolution with corresponding changes in the signal recycling cavity tuning to provide optimal sensitivity [56, 156]. The plan for the Einstein Telescope is to have sufficient sensitivity in this band to provide access to this new science; low frequency sensitivity is difficult to achieve in Earth-based detectors, however.

As discussed in Chapter 5, the noise influencing gravitational wave interferometers can be split into groups arising from the sensing of signals and from sources that directly influence the test mass displacement. Sources of sensing noise are in general independent of the arm length and test mass parameters and instead rely on the amount of light power and the readout scheme being used. The effect on the strain sensitivity of displacement noise sources, however, typically scales inversely to arm length¹; because other detector parameters must be re-optimised for each new arm length, this scaling is non-trivial [17].

As an example of both the benefits and challenges that longer arms can create, consider the arm cavities within a Fabry-Perot Michelson interferometer. For a given readout technique, longer arms provide better strain sensitivity, as shown by equation (1.7). Some sources of noise are also reduced with respect to the increased signal [18]: quantum shot noise scales, for a fixed detector bandwidth, as the square root of the ratio of the lengths, such that an arm twice as long as another has a factor $\sqrt{\frac{1}{2}} = 0.7$ of the shot noise with respect to the signal; quantum radiation pressure noise reduces even further, proportionally to the power $\frac{3}{2}$. Coating thermal noise scales, for fixed beam size, inversely to the arm length. The width of a Gaussian beam also scales inversely to the arm length, however, so the beam spreads over a wider surface area leading to clipping loss unless the mirror surface area is also scaled. While larger beams provide a further improvement in coating thermal noise, it also requires significantly larger test masses given that an increase in surface area corresponds to a significantly larger increase in mass, especially considering that the aspect ratio of the test masses should be kept close to that of the test masses employed in current detectors due to thermal noise [157]. With larger test masses and longer arms come more complex suspension design requirements, particularly in the alignment of the cavity test masses. Significant research and development is ongoing to determine designs to mitigate these issues for the next generation of detectors.

¹This is not because the displacement noise decreases with arm length, but because the response of the interferometer is improved.

7.1.1 New facility

As discussed in section 2.2.4.1, seismic noise in current detectors limits the sensitivity at low frequencies and creates challenging control requirements due to the mirror motion created by ground vibrations. The Einstein Telescope interferometers will be 100 to 200 m underground to mitigate seismic noise. This location also helps to limit the impact of gravity gradient noise, as discussed in section 2.2.4.2, which is expected to become a problem as seismic noise is mitigated.

There are a number of benefits to having multiple detectors located in the same facility. The noise properties across the facility will have a similar impact and so it should be possible to combine the signals from each detector in such a way as to generate a *null stream* that contains noise but not signal [158, 159]; this will be useful for the characterisation of noise sources and will be particularly beneficial for the new noise sources that may be interrogated due to the increased sensitivity. The arrangement of three detectors in a triangle also allows the facility to be optimally sensitive to gravitational waves from all directions [160], whereas existing single-interferometer detectors are sensitive only to incident signals out of the plane of the detector. The use of multiple detectors also allows upgrades to be made to some of the interferometers without losing sky coverage.

7.1.2 Xylophone configuration

To provide maximum astrophysical reach the facility is intended to provide sensitivity across an unprecedented bandwidth, from around 2 Hz to 10 kHz—a bandwidth significantly larger than that of existing detectors. It was realised that the most technically feasible option to obtain this bandwidth would be to implement two different types of detector, each optimised to provide good sensitivity in either low or high frequencies [161], an idea first proposed for Advanced LIGO [162]. In the proposed *ET-D* configuration [47], a low power, cryogenic interferometer optimises sensitivity to reduce radiation pressure noise at the expense of shot noise, while a high power interferometer optimises high frequency sensitivity through the reduction of shot noise. The projected sensitivity of this arrangement is shown in black in figure 7.1.

7.1.2.1 ET-LF

The low frequency detector consists of a dual-recycled Fabry-Perot Michelson interferometer configuration as introduced in section 2.3.3, but with a detuned signal recycling cavity. This detuning allows for enhanced sensitivity at the signal recycling cavity pole where the optomechanical dynamics create an optical spring that provides sensitivity below the SQL

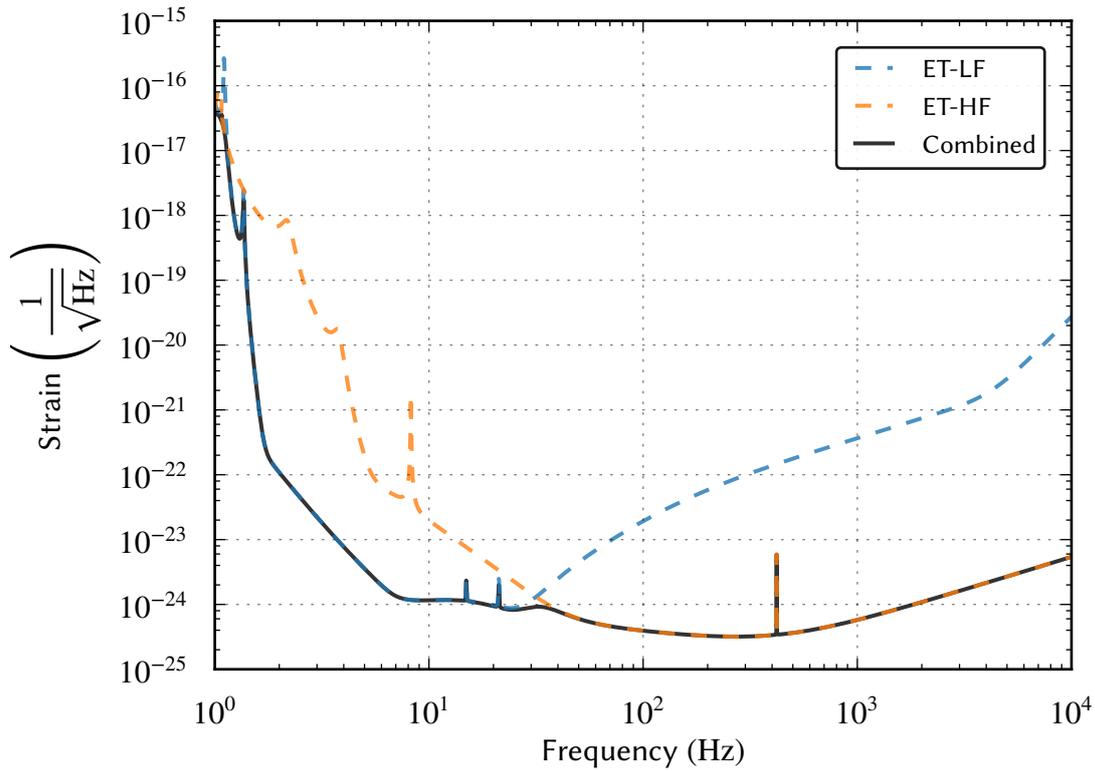


Figure 7.1: Sensitivity of the Einstein Telescope detectors, based on the ET-D design [47]. ET-LF is optimised for low frequencies, ET-HF is optimised for high frequencies, and the combination yields sensitivity between 2 Hz and 10 kHz.

at the spring frequency. The cavities will have 18 kW of light power, which is considerably lower than that of Advanced LIGO at design sensitivity (800 kW), leading to reduced quantum radiation pressure noise at low frequencies. Cryogenic test masses are to be used to facilitate a reduction in thermal noise, and the wavelength of the carrier will be changed from the standard 1064 nm to 1550 nm to utilise lower noise materials at such temperatures. Two filter cavities to facilitate frequency dependent squeezing for the further suppression of quantum noise are also included.

The suspension systems for the main test masses are based on those of the *superattenuator* in Virgo [163]. The proposed 17 m long pendulum system pushes the longitudinal resonant frequency down from around 1 Hz in existing detectors to 170 mHz, providing better attenuation of seismic noise above 2 Hz.

The sensitivity of ET-LF is shown in blue in figure 7.1. The sensitive frequency band is between around 2 Hz and 200 Hz.

7.1.2.2 ET-HF

ET-HF takes the designs of Advanced LIGO and Advanced Virgo and assumes improvements to the test mass coating loss, substrate absorption and available input laser power based on expectations for future research, and adds new technologies such as LG33 cavity modes [164] and frequency dependent squeezing [43] to reduce coating thermal and quantum noise. The combination of greater arm cavity power, heavier test masses, squeezing and improved coatings and materials will increase sensitivity at frequencies above a few hundred Hz beyond the current generation by a factor of around 10, as shown in the ET design study.

The sensitivity of ET-HF is shown in orange in figure 7.1.

7.2 Control challenges with the Einstein Telescope

Both ET-LF and ET-HF will present new challenges to the control of large-scale dual-recycled Fabry-Perot Michelson interferometers. Although ET-HF can to some extent be seen as a larger version of Advanced LIGO and Advanced Virgo, and it may therefore be possible to adapt much of the advance detectors' strategies for both longitudinal and angular control, some aspects such as the use of LG33 modes on a large scale and the presence of parametric instabilities at such high arm cavity powers [60] require extensive research to understand the implications they may have on control. ET-LF will also use a topology that resembles existing generation detectors, but it pushes the sensitivity at low frequencies further down and this presents additional challenges with sensing and noise. This chapter will discuss the longitudinal control of ET-LF, focusing in particular on the challenge of controlling the interferometer in its detuned state. The following subsection discusses the various configurations for tuned and detuned signal recycling.

7.2.1 Signal recycling and resonant sideband extraction

As introduced in section 2.3.3, signal recycling cavities can be used to enhance the sensitivity of a detector over a particular frequency band [165]. *Tuned* signal recycling involves holding the signal recycling cavity resonant by ensuring that the carrier light's transmitted and reflected phase difference is zero, giving an enhancement below the signal recycling cavity's pole frequency. *Detuned* signal recycling, meanwhile, involves configuring the signal recycling cavity's tuning in such a way as to provide the greatest sensitive bandwidth, by choosing to make the signal recycling cavity resonant for one of the signal sidebands instead of the carrier. This involves detuning the phase of the signal recycling mirror as seen by the carrier, ϕ , within the range $0 < \phi < \frac{\pi}{4}$ [166]. Signal recycling techniques enhance

the storage time of the signal within the interferometer. Another technique for changing the response of the interferometer is through the use of *resonant sideband extraction* [167]. This is typically used alongside topologies with arm cavities and the purpose is to *decrease* the storage time of the signal within the interferometer. The signal recycling cavity in this case is instead called the signal *extraction* cavity. *Tuned* resonant sideband extraction broadens the response of the interferometer beyond the bandwidth of the arm cavities, achieved by making the carrier anti-resonant ($\phi = \pi$) within the signal extraction cavity. The signal extraction cavity is then less reflective than the arm cavity ITMs, effectively reducing the arm cavity finesse for the signal and therefore increasing the sensitive bandwidth. *Detuned* resonant sideband extraction in contrast to the tuned variety provides greatest sensitivity at a non-zero signal frequency at which neither the arms nor the signal extraction cavity are resonant, and can be used to enhance the interferometer's sensitivity for a particular source. This involves tuning the signal extraction cavity slightly off anti-resonance, within the range $\frac{\pi}{4} < \phi < \pi$ [166].

Detuned techniques involve the use of radiation pressure induced dynamics. So-called *optical springs* are created when high laser power encounters optical cavities detuned from resonance, and the resulting optomechanical interactions can result in enhanced sensitivity at the spring frequency [76].

7.2.2 Resonant sideband extraction in ET-LF

The ability for a detuned signal recycling cavity to shift the most sensitive frequency has been demonstrated in GEO 600 between 200 Hz and 1 kHz [165], however the plan for ET-LF is to use detuned resonant sideband extraction with a detuning of around 25 Hz, a feat that has not been achieved before in a suspended audio-band detector. With tuned techniques, the sidebands used for control of the differential arm cavity mode are present within the signal recycling cavity with equal amplitude. In detuned operation, where the cavity is not resonant for the carrier, the sidebands have unequal amplitude and some of the interferometer's noise couplings that otherwise cancel at the output in the tuned case no longer cancel [168]. The phase modulation of the control sidebands, created by EOMs on the input path, also gets partially converted to amplitude modulation by the detuning, and this can lead to issues with the dynamic range of, and offsets upon, any photodetectors used to sense the readout [169]. Often it can be difficult to find a port at which to sense the motion of the signal recycling cavity decoupled from other cavities, and this effect can be exacerbated by large detunings such as in ET-LF where noise cross-couplings can become more significant [168].

In the ET design study the discussion for ET-LF stopped short of a control scheme. The rest of this chapter will discuss some control concepts, present a conceptual approach to the

control of ET-LF and highlight the future work that must be undertaken before a technical design for the control of ET-LF can be produced.

7.3 Longitudinal control of a dual-recycled Fabry-Perot Michelson interferometer

A successful control scheme for an interferometer must satisfy a number of requirements. When the interferometer is in its uncontrolled state, the control scheme must be able to bring it to the *operating point* where it has the desired response to incident gravitational wave signals in a process called *lock acquisition*. Once at the operating point, it must be robust against small perturbations by controlling the impact of noise and signal nonstationarities. Finally, the signal that represents the gravitational wave channel must have low noise, and therefore high sensitivity, to meet the design goals of the scheme.

The lock acquisition scheme is inextricably linked to the technical environment in which the interferometer will operate, and so it is inappropriate to discuss this while ET's technical design is subject to ongoing research. We will focus our efforts, therefore, on the second and third challenges above.

7.3.1 The dc readout technique

The standard readout technique for GEO 600 and the current generation detectors is dc readout [24, 124, 168], and the plan is for ET-LF and ET-HF to continue to use it. This technique is a form of homodyne readout that involves a compromise between the best sensitivity and technical complexity. The operating point is kept close to the *dark fringe* (see section 2.1.2) to optimise shot noise and reduce the coupling of technical noise, but a slight offset is introduced between the differential round trip phase of each arm in order to allow some of the carrier light to enter the readout port where it acts as a homodyne local oscillator to the signal sidebands (see section 2.1.3.2). In practice, this detuning—of the order pm at the arm cavities in Advanced LIGO—is sufficiently small to prevent the sensitivity from being diminished. The light that does enter the output port is, however, pre-filtered by the arm cavities and the impact of laser noise is suppressed by the differential detuning. As the source of local oscillator is a fraction of the light from the cavities to be controlled with the readout, the phase stability of the local oscillator is linked to that of the carrier light, as opposed to having a separate degree of freedom to control as with other homodyne techniques. The sensitivity of this readout is furthermore improved with respect to heterodyne techniques [24], and in squeezed interferometers it avoids the need to inject squeezing at RF sideband frequencies in addition to the carrier frequency.

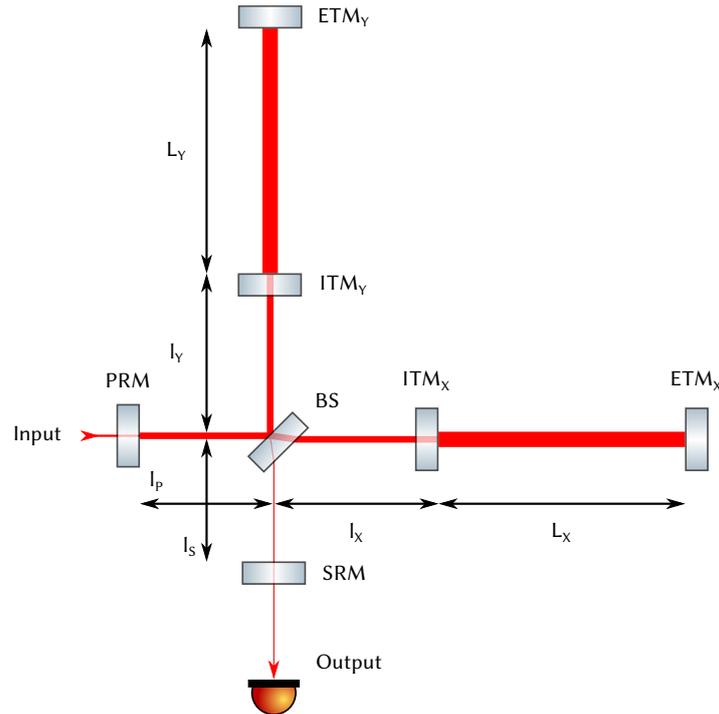


Figure 7.2: Cavity lengths in a dual-recycled Fabry-Perot Michelson interferometer. The differential and common arm cavity modes compose L_X and L_Y , while the auxiliary power and signal recycling and Michelson cavity lengths are composed of a subset of the distances between the beam splitter and each ITM and the distances between each recycling mirror and the beam splitter.

7.3.2 Degrees of freedom

The *degrees of freedom* of an interferometer are the non-degenerate ways in which the interferometer's mirrors may move away from the operating point. Each degree of freedom has a different precision requirement, with the most stringent typically being the degree of freedom corresponding to the gravitational wave channel. With dc readout there is no local oscillator phase or alignment to independently control and so in a dual-recycled Fabry-Perot Michelson interferometer the main degrees of freedom to consider are the arm cavity differential and common modes, the length between the beam splitter and the ITMs and the recycling cavity lengths². These are defined in the following sections and the lengths that compose each degree of freedom are shown in figure 7.2.

The motion of each degree of freedom must be witnessed by a sensor and fed back to actuators to control the relevant length. This is called *linear negative feedback* and the concept of shaping feedback dynamics is discussed in greater detail in the context of suspended interferometers in chapter 5.

²Note that throughout this chapter we refer to the signal extraction cavity as a *recycling* cavity to fit with convention.

7.3.2.1 Differential arm cavity length

Gravitational waves change the length of the arms in a Michelson interferometer differentially and due to the presence of the arm cavities in a dual-recycled Fabry-Perot Michelson interferometer changes to the differential arm cavity length represents the motion the main readout is most sensitive to. This length is held close to the dark fringe but with a small offset for dc readout. The differential arm cavity length signal can be fed back to the ETMs differentially to hold the length at the desired operating point.

We define the differential arm cavity length, $DARM$, in terms of the average differential length of the arms:

$$\delta L_{DARM} = \frac{L_X - L_Y}{2}. \quad (7.1)$$

7.3.2.2 Common arm cavity length

In-phase changes in the length of the arms of a Michelson interferometer do not primarily couple to the gravitational wave channel, but it is crucial to control this degree of freedom in order to keep the arm cavities at their operating point which, in a dual-recycled Fabry-Perot Michelson interferometer, is the state in which the power in the arm cavities is held near maximum. Noise common to the light entering both arms, particularly from the laser's amplitude and frequency fluctuations as discussed in section 2.2.5.1, can, unless corrected, change the resonant condition in the arms. Due to its speed a convenient actuator to control common arm length changes is the laser's crystal, whereupon the application of strain changes the laser's frequency as shown by equation (1.6). It is typical to use feedback to an EOM for corrections above around 100 kHz. It is also possible to correct slower, larger drifts with common feedback to the ETMs or input optics.

We can define the common arm cavity length, $CARM$, in terms of the average length of the arms:

$$\delta L_{CARM} = \frac{L_X + L_Y}{2}. \quad (7.2)$$

Any arm cavity length change can be expressed in terms of a linear combination of $DARM$ and $CARM$.

7.3.2.3 Power and signal recycling cavity lengths

The power recycling cavity should be resonant for the input light in order to optimally recycle light reflected from the beam splitter back towards the laser, to allow the arm cavity power with respect to the interferometer's input power to be maximised. The power recycling cavity length, $PRCL$ can be defined in terms of the average distance between the

power recycling mirror (PRM) and the two ITMs:

$$\delta L_{\text{PRCL}} = l_P + \frac{l_X + l_Y}{2}. \quad (7.3)$$

The signal recycling length, along with the signal recycling mirror (SRM) transmissivity, determines the bandwidth of the signal extraction and therefore needs controlled in order to keep the interferometer's response stationary in time. This length, *SRCL*, is defined similarly to PRCL in terms of the position of the signal recycling mirror:

$$\delta L_{\text{SRCL}} = l_S + \frac{l_X + l_Y}{2}. \quad (7.4)$$

Control of PRCL and SRCL can be achieved via corrective feedback to the position of the power and signal recycling mirrors, respectively.

7.3.2.4 Michelson length

The length between the beam splitter and the ITMs should be held in the dark fringe condition for the carrier to correctly couple the common and differential arm cavity modes to the input and output port of the beam splitter, respectively. In a dual-recycled Fabry-Perot Michelson interferometer it must also be held constant to keep the amount of carrier and sideband power in the signal recycling cavity stable, which avoids the need for complicated time-varying control signals. This length, *MICH*, can be expressed as the differential length between the beam splitter and the ITMs:

$$\delta L_{\text{MICH}} = \frac{l_X - l_Y}{2}. \quad (7.5)$$

The MICH length is kept constant by feeding back to the positions of the cavity mirrors. Alternatively this feedback can be applied to the beam splitter, but in this case the recycling mirrors must also be moved to avoid influencing the lengths of PRCL and SRCL.

There is an apparent degeneracy between the dark fringes produced by the MICH and DARM degrees of freedom, and so one might expect that the dark fringe offset required for dc readout could be applied to the former. While a DARM offset has the disadvantage that it involves the creation of an optical spring due to the high light power in the arms—mechanically coupling the CARM and DARM modes [170, 171]—it has favourable noise couplings compared to a MICH offset [143]. The MICH degree of freedom is filtered above the pole frequency of the power recycling cavity, which is typically at a high frequency, whereas the DARM mode is additionally filtered by the arm cavity poles which are usually much lower in frequency, and so the noise coupling is reduced.

	ITM X	ETM X	ITM Y	ETM Y	BS	PRM	SRM
CARM	0	0.5	0	0.5	0	0	0
DARM	0	0.5	0	-0.5	0	0	0
MICH	-0.5	0.5	0.5	-0.5	0	0	0
PRCL	0	0	0	0	0	1	0
SRCL	0	0	0	0	0	0	1

Table 7.1: Driving coefficients for each mirror and each degree of freedom of ET-LF. CARM and DARM involve the driving of the ETMs in- and out-of-phase, respectively. MICH involves moving the ITMs differentially, but to avoid sensing DARM effects the ETMs must be moved too. The driving coefficients are defined with respect to the direction of the HR surface of each optic, and as such the coefficients in MICH for the ETM and ITM in each cavity have opposite sign. PRCL and SRCL involve moving the power and signal recycling mirrors, respectively.

7.3.2.5 Driving coefficients

Table 7.1 lists the coefficients to be applied to the test masses to produce the error signals representing each of the degrees of freedom. Corrections to CARM and DARM are defined in terms of the positions of the ETMs, with the differences in coefficient sign between the two degrees of freedom being whether the correction is in-phase or out-of-phase, respectively. The PRCL and SRCL degrees of freedom are defined as corrections to the power and signal recycling mirrors, respectively. MICH is a little more tricky, and as explained in section 7.3.2.4 it can be achieved with feedback either to all of the arm cavity mirrors or to the beam splitter and recycling mirrors. For the purposes of this study we choose the former to avoid the need to correct the unit drive amplitude applied to the beam splitter due to its angle of incidence.

7.3.3 Decoupled control signals

With the exception of the main dc readout for DARM, the error signals representing the degrees of freedom are derived using heterodyne schemes, usually variations of the Pound-Drever-Hall technique discussed in section 3.3.4.1. For Michelson interferometers a number of techniques have been developed such as *internal* and *external* modulation, but the most prominent technique employed in all recent detectors is *Schnupp* modulation. This uses control sidebands imposed on the carrier by means of phase modulation before the light is coupled into the arms by the beam splitter [172].

In a heterodyne sensing scheme, for the control sidebands to enter the output port their lengths must be macroscopically mismatched; this is called a *Schnupp asymmetry*. For a Schnupp asymmetry Δl_{SCH} , the two arms in a Michelson interferometer have length $l + \frac{\Delta l_{\text{SCH}}}{2}$ and $l - \frac{\Delta l_{\text{SCH}}}{2}$ where the nominal length is l . The Schnupp asymmetry differs from the previous microscopic distances defined in equations (7.1) to (7.5) in that it is typically equivalent

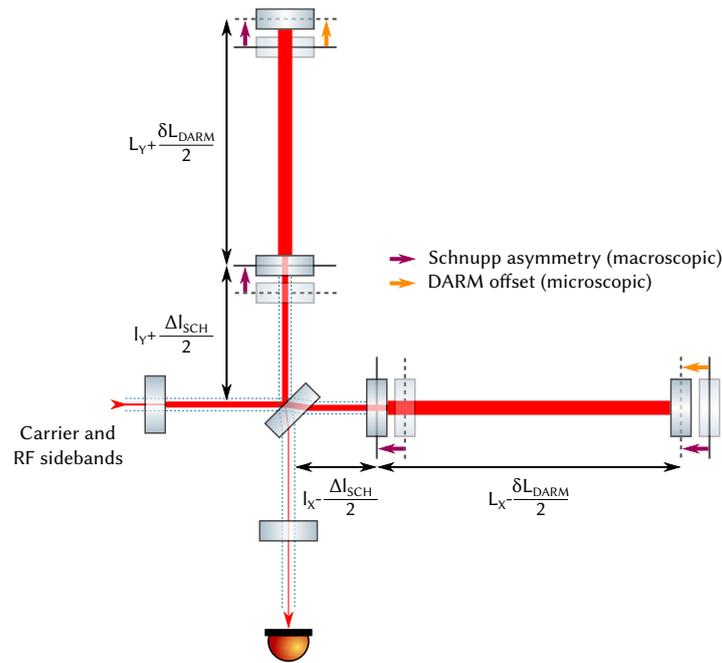


Figure 7.3: The Schnupp asymmetry and differential arm cavity offset in a dual-recycled Fabry-Perot Michelson interferometer. The Schnupp asymmetry allows the control sidebands resonant in the Michelson interferometer to leave the beam splitter's output port where they may sense the motion of the signal recycling cavity. This is typically an offset of the order of cm. The differential arm cavity (DARM) offset is a microscopic detuning of the arm cavities to allow a small amount of light to interfere constructively at the beam splitter in order for it to act as a local oscillator to the signal sidebands at the output port. This detuning is typically of the order of pm.

to many thousands of wavelengths. Although a dual-recycled Fabry-Perot Michelson interferometer will typically employ dc readout and therefore does not require heterodyne error signals for the control of DARM, this asymmetry is required for the control of SRCL. The application of a Schnupp asymmetry within a dual-recycled Fabry-Perot Michelson interferometer is shown in figure 7.3.

In a simple Michelson interferometer control sidebands at a single modulation frequency can be imposed upon the carrier to discriminate the phase of the arm cavities with respect to the input, and the modulation frequency is chosen such that the control sidebands will propagate in the interferometer but not enter the arm cavities so that they can act as a phase reference for the light that does. While the Schnupp modulation technique is applicable to dual-recycled Fabry-Perot Michelson interferometers [54], it is not trivial to decouple the five degrees of freedom when they are sensed by the carrier and a single modulation frequency. For example, the motion of the signal recycling mirror cannot be decoupled from the motion of the power recycling mirror in a single control sideband frequency; instead the signal will contain a linear combination of the motion of the two. By employing a second modulation frequency a suitable phase reference can be obtained for each of the degrees of freedom by carefully arranging for each modulation frequency to resonate in a subset of the

cavities to control. Methods to find and create decoupled error signals in a dual-recycled Fabry-Perot Michelson interferometer are discussed in the following subsections.

7.3.3.1 Gain hierarchy

Some degrees of freedom require a wider control bandwidth in order for the interferometer to be held at its operating point. For example, laser frequency and intensity noise is present across the measurement band at the input of the interferometer and must be suppressed by many orders of magnitude to be compatible with the sensitivity requirements of a audio-band gravitational wave interferometer. In order to provide high gain within the detection band the bandwidth of the controller (see appendix B.4.2) is high. Laser noise mainly couples to CARM and not to the other degrees of freedom, but control of CARM is necessary to maintain light power within the arm cavities, and hence sensitivity. In Advanced LIGO the control bandwidth for CARM is 65 kHz [173]. On the other hand, control of DARM and the auxiliary degrees of freedom tends to be driven by the presence of low frequency seismic noise and so typically requires much smaller control bandwidth (see for example the differential arm length control precision requirement for the Sagnac speed meter experiment in section 5.2.4). Through the appropriate selection of servos for each degree of freedom the cross-couplings present at each sensor from secondary degrees of freedom can be suppressed. This is called *gain hierarchy* and has been previously demonstrated in LIGO [174].

7.3.3.2 Combination of control sideband frequencies

Error signals with greater decoupling, particularly for the inner degrees of freedom, can in some cases be found by demodulating the light at some combination of the two sideband frequencies [175, 176]. Some of the control signals used in Advanced LIGO involve demodulation at the sum or difference of the two sideband frequencies [173] or the use of double-demodulation [177].

7.3.3.3 Control matrix operations

As the resonant conditions of the control sidebands can be different from the carrier, the magnitude and phase with which error signals representing each degree of freedom appear at each probe can be different.

The gradient of the error signal as witnessed by a sensor for motion of a given degree of freedom represents the interferometer's response for that degree of freedom, as shown for example in figure 5.19 for the Sagnac speed meter experiment. A *sensing matrix* can

be assembled with the collection of error signal slopes from each degree of freedom in an interferometer to each probe. With this matrix it is then possible to perform row and column operations to suppress cross-couplings in an operation that resembles Gaussian elimination. As the creation of linear combinations of signals can be performed in real time by the controller, the assembly of a *control matrix* based on simulations and used to feed back combinations of error signals to the interferometer's actuators can serve as an approximation to the eventual implementation. This is the starting point for further analysis of the effects of noise coupling from each degree of freedom to the sensors.

Some signals from different degrees of freedom appear with similar magnitude at a particular port and cannot easily be decoupled with gain hierarchy or row and column operations. Instead, the phase of the readout can be used to discriminate between the two error signals. By demodulating each light field at two phases separated by 0 and 90° with respect to the control sideband modulation frequencies, so-called *I* and *Q* quadrature error signals can be obtained which can later be mixed to suppress the effect of one error signal with respect to another.

7.4 Longitudinal sensing scheme for ET-LF

The following section presents a concept for the longitudinal control of ET-LF. Using the techniques discussed in section 7.3, we will devise a control scheme based on the approach taken for the second generation detectors, namely through the use of dc readout for DARM and heterodyne readout for CARM, PRCL, SRCL and MICH.

7.4.1 Scope and method

Given the complexity of the interferometer we choose a numerical approach to the modelling and for this we will employ Optickle (see appendix C.1.2) using its plane wave mode. Angular control of the interferometer is expected to present its own challenges given the higher cavity power, larger beams and *g*-factors closer to unity compared to current generation detectors, but the longitudinal degrees of freedom must be shown to be controllable before angular degrees of freedom can be considered. Future angular sensing and control simulations will be possible with the model developed over the course of this work.

The sensing scheme assumes that the interferometer has been brought close to its operating point by a lock acquisition routine, and here we primarily consider the control of the interferometer in terms of its sensing matrix. It is well known that noise coupling between longitudinal degrees of freedom of detuned dual-recycled Fabry-Perot Michelson interferometers can be significant [168], but a crucial initial step before any control noise simu-

lations can be undertaken is the development of a provisional sensing scheme. For ET-LF we loosely follow the approach taken for Advanced LIGO [173] and Advanced Virgo [178] given that these represent the most sensitive dual-recycled Fabry-Perot Michelson interferometers built to date. To simplify the steps required to produce the scheme the parameters are defined in order from the least to the most constrained. With these parameters fixed we can then search for ports at which error signals representing each degree of freedom can be extracted, and this allows us to define a sensing matrix that highlights cross-couplings between each of the degrees of freedom. This is the starting point for control loop noise studies employing gain hierarchy as discussed in section 7.3.3.1, which will be the subject of future work.

7.4.1.1 Optimal input coupling

Placing a power recycling mirror before a Michelson interferometer, as discussed in section 2.3.2, creates an additional cavity between the input and the ITMs. The intention of the power recycling mirror is to minimise the light reflected back towards the laser, and in order to do this the cavity it creates should be *impedance matched* (see, for example, ref. [23, section 5.1]). The loss within the interferometer in combination with the transmissivity of the power recycling mirror T_{PRM} determines the impedance matching. In the ET design study the loss per surface is assumed to be 35 ppm and the transmissivity of the ETMs is 6 ppm; both contribute to the total loss. It also suggests the transmissivity of the power recycling mirror to be 4.6 %. Figure 7.4 shows the ratio of the light power leaving the power recycling mirror heading back towards the laser to the input light power. The minimum reflected power corresponds to a power recycling mirror transmissivity of 4.6 %, validating the choice from the ET design study. If the scatter or substrate loss of the optics is changed in the future, for instance due to the development of new coatings or the introduction of additional steering mirrors, this model can be used to recalculate the optimum power recycling mirror transmissivity.

7.4.2 Control sideband frequencies

In the following subsections we define the constraints on the control sidebands before calculating appropriate frequencies. For a given interferometer it will be possible to find a set of control sideband frequencies any of which will be acceptable for a particular length, but in order for two sideband frequencies to be resonant in a combination of lengths it may be necessary to compromise the resonant conditions such that neither one is optimal for its particular purpose. The case is further complicated by the presence of a detuned signal recycling cavity in ET-LF, where the upper and lower control sidebands created with

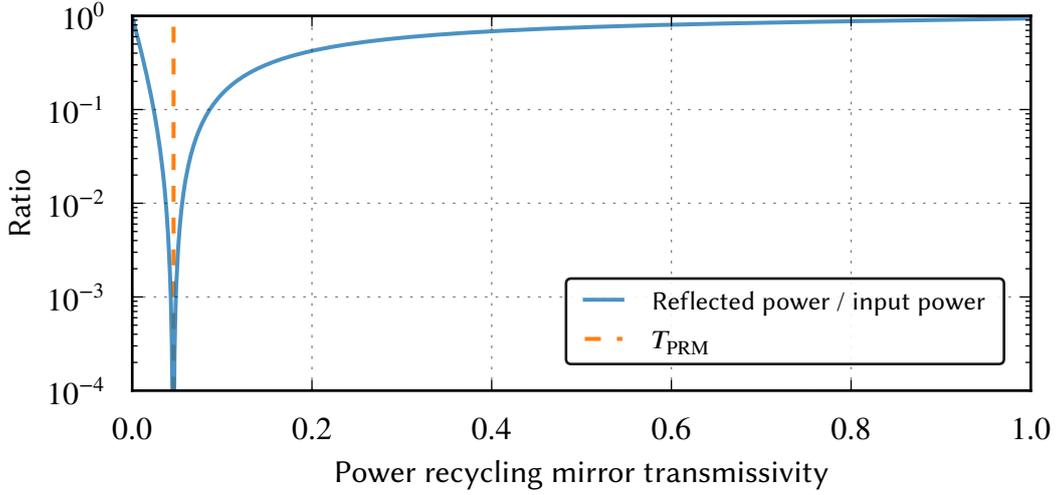


Figure 7.4: Reflected power from ET-LF as a function of power recycling mirror transmissivity. For optimal coupling of the input laser light to the interferometer, the transmissivity of the power recycling mirror must be set to balance the input light with the total loss from the interferometer. For ET-LF with loss as per the ET design study, this transmissivity should be 4.6 %.

frequencies $-f$ and $+f$ offset from the carrier (see figure A.4) have different resonant conditions. These issues will be addressed in the following subsections.

7.4.2.1 Control sideband resonance in the recycling cavities

The 310 m recycling cavity lengths defined in the ET design study have free spectral range (FSR, see appendix A.2.1):

$$\text{FSR} = \frac{c_0}{2L_{\text{PRCL}}} = \frac{c_0}{2L_{\text{SRCL}}} = 483.5 \text{ kHz}. \quad (7.6)$$

To start, we can try to make the first sideband frequency resonate in the power recycling cavity. As the arm cavities at the operating point reflect the light back towards the power recycling cavity, the resonant condition is a half-integer multiple of the power recycling cavity FSR, i.e.:

$$f_1 = \left(A + \frac{1}{2}\right) \frac{c_0}{2L_{\text{PRC}}}, \quad (7.7)$$

for positive integer A .

We can repeat this step for the signal recycling cavity, instead making the second sideband frequency resonant and the first anti-resonant in order to create discrimination. In this case

the resonant condition is an integer multiple of the cavity's FSR, i.e.:

$$\begin{aligned} f_1 &\neq B \frac{c_0}{2L_{\text{SRC}}} \\ f_2 &= C \frac{c_0}{2L_{\text{SRC}}}, \end{aligned} \quad (7.8)$$

where B and C are again positive integers.

Any control sideband frequencies used in ET-LF must be outside the arm cavity resonances spaced by integer multiples of the arm cavity FSR in order to allow them to act as a phase discriminant for CARM and DARM. The 10 km arm cavity FSR is 14.99 kHz. As an integer multiple of the arm cavity FSR would allow optimal coupling of the sidebands into the arm cavities, one might assume that an odd half-integer multiple would be optimally anti-resonant; however, in this scenario the lower higher-order control sidebands, necessarily created by the phase modulation upon the EOM (see equation (A.20)), would become resonant, and so we choose to offset the sideband frequency slightly from the anti-resonant condition. We therefore stipulate two further requirements in addition to equations (7.7) and (7.8):

$$f_1 = (D_1 + \delta_1) \frac{c_0}{2L_{\text{Arm}}} \quad (7.9)$$

$$f_2 = (D_2 + \delta_2) \frac{c_0}{2L_{\text{Arm}}}, \quad (7.10)$$

for positive integers D_i and small perturbations $|\delta_i| \ll \frac{1}{2}$.

7.4.2.2 Control sideband frequencies

The control sideband frequencies should ideally be RF, at least around 10 MHz, to benefit from the noise advantages discussed in section 3.3.4.1. An upper limit of 100 MHz is reasonable given that quadrant photodetectors requiring large surface area will eventually be required for alignment control. Larger surface areas typically lead to greater stray capacitance, limiting the ability of the device to register signals at higher frequencies [23].

We chose the first sideband frequency f_1 to be 11 363 101 Hz which for $A = 23$ and $D_1 = 758$ satisfy the requirements and falls within the suitable range. As the light in the signal recycling cavity must first pass through the power recycling cavity, we must ensure that f_2 also resonates in the power recycling cavity. This is achieved by choosing the second sideband to be an integer multiple of the first, which is already resonant in the power recycling cavity. To provide ample difference between the first and second sideband frequencies, we chose $f_2 = 5f_1 = 56 815 505$ Hz. This is separated far enough in frequency from f_1 that we can investigate the use of beats between f_1 and f_2 for control purposes as discussed in

section 7.3.3.2.

We assume a modulation depth of 0.1 rad for f_1 and f_2 to keep the power in higher modulation orders low whilst allowing for a reasonable amount of light power in the first order. This parameter has little impact on others and can be tuned later to provide better separation between the two sideband signals on the sensors.

7.4.3 Schnupp asymmetry

In addition to facilitating the control of the signal recycling cavity, the Schnupp asymmetry governs whether the second control sideband frequency couples to both recycling cavities, or just one. A small offset of a few cm between the Michelson lengths, l_X and l_Y , allows for the signal recycling cavity to be resonant for only one of the sideband frequencies, whereas a larger offset of a few tens of cm makes both sideband frequencies resonant there [178]. Both methods of control are feasible, with the former being implemented in Advanced LIGO [173] and the latter in KAGRA [179].

Figure 7.5 shows the power of the upper and lower sideband fields with respect to the carrier in the recycling cavities of ET-LF given the Schnupp asymmetry and recycling cavity lengths with detuned signal recycling. Since the Schnupp asymmetry is a macroscopic length, it is not easily adjusted during operation and so it is necessary to designate this length in the design phase. Here we choose a Schnupp asymmetry that attempts to maximise the difference in power between the two sideband frequencies in the recycling cavities during detuned operation. This is around 0.08 m, as shown by the black dashed line in figure 7.5. Also observe that a Schnupp asymmetry of 0 results in no sideband power in the signal recycling cavity, and that asymmetries of around 0.5 m result in a situation where both sideband frequencies are resonant in both recycling cavities. In the latter case it would be possible to obtain decoupled control signals through control matrix operations as discussed in section 7.3.3.3.

7.4.3.1 Optimisation of the signal recycling cavity length

Note the discrepancy between two of the control sideband frequency constraints in section 7.4.2: the power and signal recycling cavities cannot both be simultaneously resonant and anti-resonant to f_1 and f_2 given that $f_2 = 5f_1$. To resolve this discrepancy we can scan the length of the signal recycling cavity in order to find a position where f_2 is resonant and f_1 is not, as shown in figure 7.6. This could otherwise have been achieved by changing f_2 by a fraction of the signal recycling cavity's FSR given the relation between frequency and length shown in equation (1.6). We can see that changing the signal recycling cavity length from 310 m to 311.585 m results in the desired sideband resonance condition for the

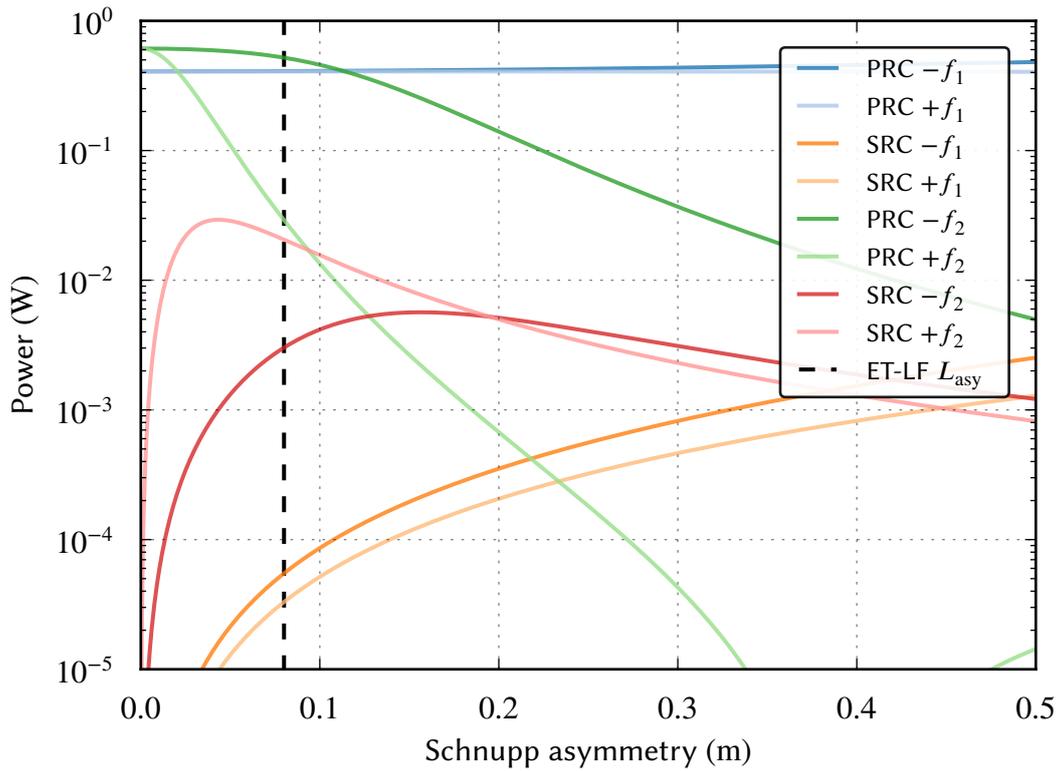


Figure 7.5: Power of the control sidebands in the cavities of ET-LF during detuned operation as a function of Schnupp asymmetry. A macroscopic offset—the *Schnupp asymmetry*—is intentionally introduced to the Michelson length in order to allow the coupling of the sidebands $\pm f_1$ and $\pm f_2$ into the signal recycling cavity for the purposes of control whilst maintaining the dark fringe condition for the carrier. Here, we choose to allow $\pm f_2$, but not $\pm f_1$, to couple to the signal recycling cavity. The power is a reasonable estimate for the relative sensitivity of the sidebands in the recycling cavities, and we choose an asymmetry which gives good separation of the power of the sidebands in each cavity, 0.08 m.

$+f_2$ sideband. The power of f_2 in the power recycling cavity drops as the lower and upper sidebands get critically coupled into the signal recycling cavity. As the signal recycling cavity detuning in ET-LF is large, the signal recycling cavity is not resonant for both the upper and lower f_2 sidebands (in the tuned signal recycling cavity configuration the upper and lower sidebands are degenerate).

In the power recycling cavity, f_1 provides an error signal for the power recycling cavity that is a factor of 13 larger than that of f_2 . At the same time, f_2 provides an error signal for the length that is 650 times larger than the equivalent for f_1 in the signal recycling cavity.

7.4.4 Dark fringe offset

As described in section 2.1.3.2, dc readout at the output port of a dual-recycled Fabry-Perot Michelson interferometer requires carrier light to be present to act as a phase reference for

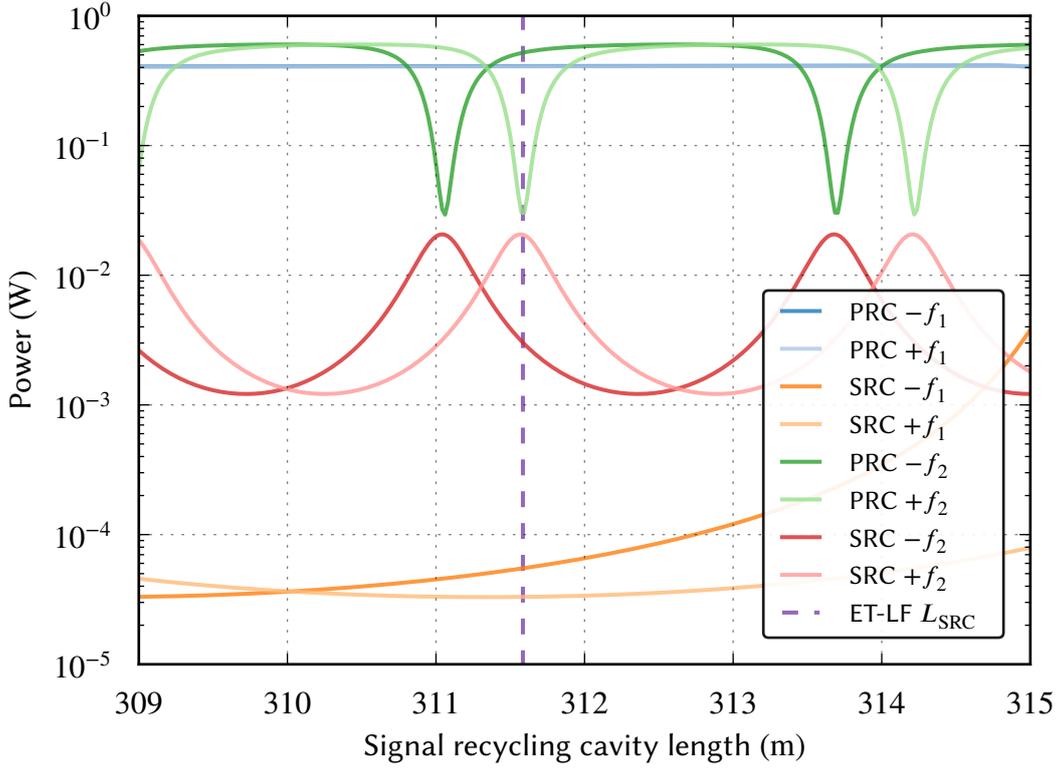


Figure 7.6: The power in each sideband in each recycling cavity as a function of signal recycling cavity length. The nominal signal recycling cavity length defined by the ET design study, 310 m, is not resonant for the second sideband as intended. As the sideband frequencies are different, we can fix this situation by scanning the signal recycling cavity to find a length that is resonant for f_2 but not for f_1 . For the given choice of Schnupp asymmetry, this occurs at a length of 311.585 m for the upper sideband. The choice to optimise either the upper or lower sideband is arbitrary since both sidebands contribute signal at the readout.

the signal sidebands. In an interferometer with matched arms there is no classical light at the output port and so a phase asymmetry must be introduced by differentially detuning the arms by a small amount to create the appropriate *dark fringe offset*. In practice, asymmetries within the arms are already present, for example arising from mismatched arm cavity finesse or asymmetric beam splitter reflectivity. As these effects change the amplitude of the light, they appear in a different quadrature to the signal at the output port and so as long as the loss is small the effect on the sensitivity is minimal.

For our model we define the DARM offset as a microscopic detuning of the arm cavity lengths, and it is differential such that one arm has length $L + \frac{\delta L_{\text{DARM}}}{2}$ while the other has length $L - \frac{\delta L_{\text{DARM}}}{2}$, where L is the average length. This is depicted in figure 7.3 alongside the Schnupp asymmetry.

Figure 7.7 shows the power at the output port and in the arm cavities as a function of DARM offset in the detuned configuration. Standard photodetectors used in Advanced LIGO and

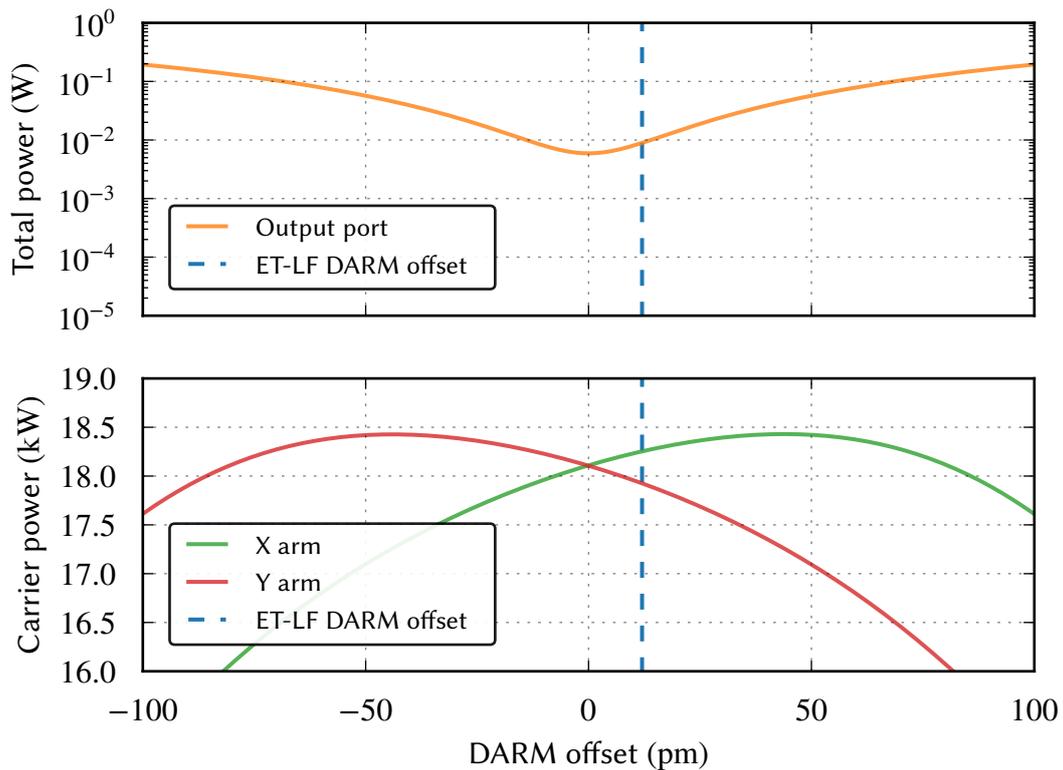


Figure 7.7: Carrier power at the output port of ET-LF in detuned configuration as a function of differential arm cavity length (DARM) offset. The differential arm detuning required to allow carrier light to enter the dark port for dc readout involves an increase or decrease in the microscopic length of each arm cavity, and this changes the circulating power. The compromise must be made between the power available to the photodetector for sensing whilst maintaining reasonably balanced arm cavities to prevent optical springs from influencing the sensitive band and creating additional noise coupling.

Advanced Virgo can handle up to a few 10s of mW and this should ideally be the power incident upon the photodetector at the output port to maximise the signal to dark noise ratio. In ET-LF, however, the DARM offset required to reach this figure would create a significant mismatch in the power of each arm leading to a strong optical spring effect. For an offset of 12 pm, the power at the output can be set to around 10 mW with a difference of around 3% in the power in the arms, which should be tolerable in terms of sensitivity and noise coupling.

7.4.5 Power in each light field

The power in each field within each relevant space or cavity of the interferometer is shown in table 7.2 for the detuned interferometer.

	-68 MHz	-57 MHz	-45 MHz	-11 MHz	Carrier	11 MHz	45 MHz	57 MHz	68 MHz	Total
Input from laser	0	0	0	0	3 W	0	0	0	0	3 W
After modulators	19 μ W	7 mW	19 μ W	7 mW	3 W	7 mW	19 μ W	7 mW	19 μ W	3 W
Power recycling cavity	220 nW	521 mW	220 nW	410 mW	65 W	407 mW	220 nW	34 mW	220 nW	66 W
Power recycling pick-off	33 pW	78 μ W	33 pW	61 μ W	9 mW	61 μ W	33 pW	4 μ W	33 pW	10 mW
Michelson cavity	131 nW	228 mW	112 nW	208 mW	33 W	204 mW	107 nW	27 mW	115 nW	33 W
Arm cavity X	261 pW	557 μ W	375 pW	9 mW	18 kW	9 mW	357 pW	65 μ W	230 pW	18 kW
Arm cavity Y	182 pW	722 μ W	361 pW	9 mW	18 kW	9 mW	379 pW	68 μ W	210 pW	18 kW
Signal recycling cavity	3 nW	4 mW	379 pW	69 μ W	15 mW	41 μ W	411 pW	26 mW	928 pW	45 mW
Reflected back to laser	19 μ W	6 mW	19 μ W	7 mW	62 μ W	7 mW	19 μ W	2 mW	19 μ W	23 mW
Output	558 pW	756 μ W	76 pW	14 μ W	3 mW	8 μ W	82 pW	5 mW	186 pW	9 mW

Table 7.2: Powers in various parts of ET-LF in the detuned configuration. The input light is passed through EOMs which impart control sidebands at frequencies offset from the carrier. As the arm cavity FSR is almost optimally separated from the control sideband frequencies, the sideband light power in the arms is vastly smaller than the carrier power. The first sidebands at ± 11 MHz are resonant within the power recycling cavity. The second sidebands at ± 57 MHz is greatest within the power and signal recycling cavities, with the -57 MHz sideband almost exactly resonant in the power recycling cavity and the $+57$ MHz sideband exactly resonant in the signal recycling cavity. The power reflected back towards the laser is composed mainly of light at the two sideband frequencies, as the transmissivity of the power recycling mirror minimises the reflected carrier light. At the output port, the carrier power is present due to the DARM offset, and acts as a local oscillator to the signal sidebands there.

7.4.6 Readout ports

Figure 7.8 shows some available readout ports for ET-LF where the sidebands and carrier may be measured for the purposes of control:

- **REFL** (*reflected*) senses the light reflected from the interferometer back towards the input laser, rejected for example by a Faraday isolator;
- **POP** (*pick off PRCL*) senses the light in the power recycling cavity using a small pick-off mirror with reflectivity 150 ppm;
- **AS** (*asymmetric*) senses the light at the output port of the dual-recycled Fabry-Perot Michelson interferometer.

For the purposes of the control simulations, each readout port contains incident fields at the carrier frequency and offsets of $\pm f_1$ (± 11 MHz), $\pm f_2$ (± 57 MHz), $-f_1 - f_2$ (-68 MHz), $f_1 - f_2$ and $-f_2 + f_1$ (-45 MHz³), $-f_1 + f_2$ and $f_2 - f_1$ (45 MHz), and $f_1 + f_2$ (68 MHz). There are therefore 9 light fields propagated through the interferometer that are combined into signals at 5 frequencies at 3 ports. We furthermore demodulate the heterodyne sensors in the I and Q quadratures to be able to calculate the optimal readout phase at each port given the propagation delay between each mirror and each sensor. These are shown in table 7.3.

³This is not -46 MHz due to rounding.

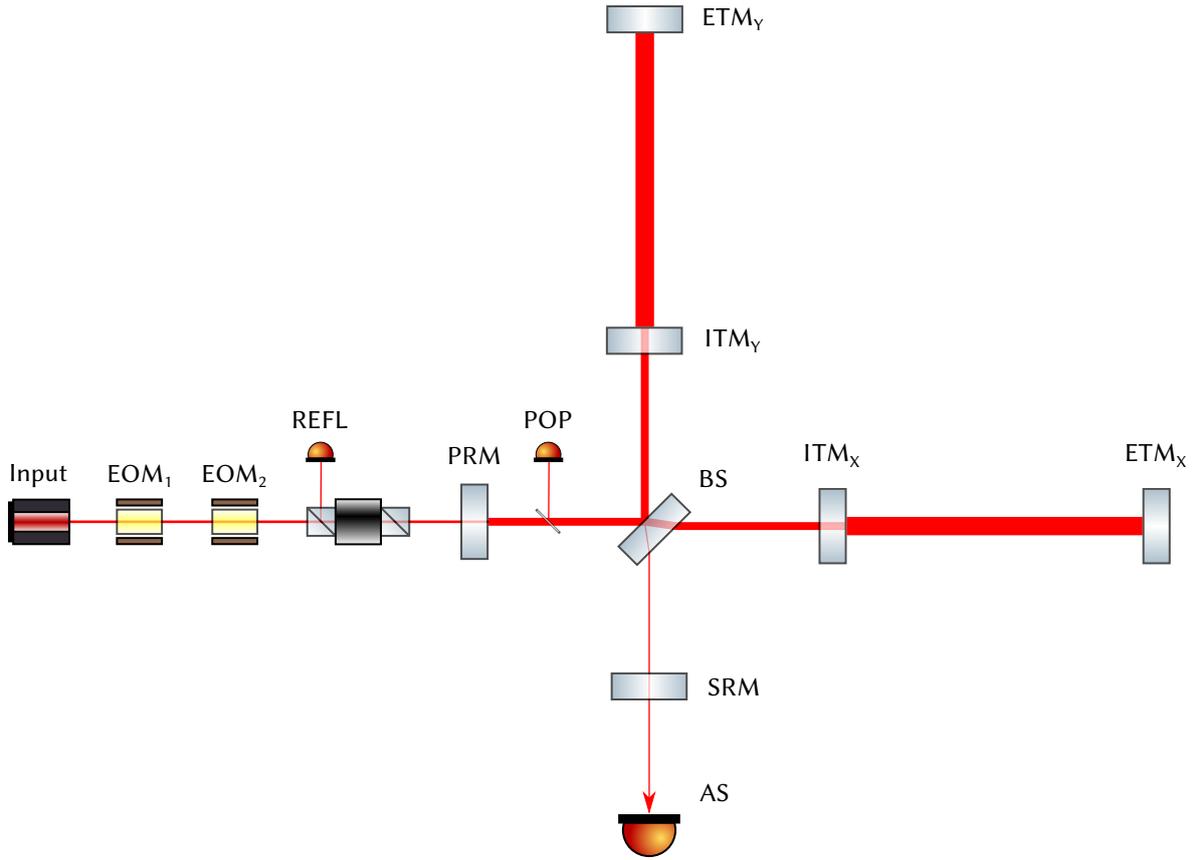


Figure 7.8: Some available readout ports for sensing motion of the degrees of freedom in ET-LF in detuned configuration. Phase modulation sidebands are imparted upon the input laser light at two primary frequencies before it enters the interferometer. Photodetectors are placed at three ports to sense the carrier and these sidebands for the purposes of sensing and control. The most important readout is AS, which strongly senses DARM, and this is placed in transmission of the signal recycling mirror. Light reflected back towards the laser is sensed via a Faraday isolator at the REFL port, and a pick-off—POP—senses light in the power recycling cavity.

Offset	Output port	Power recycling cavity pick-off	Reflected light pick-off
0	AS_{DC}	—	—
$\pm f_1$	AS_{11}^I, AS_{11}^Q	POP_{11}^I, POP_{11}^Q	$REFL_{11}^I, REFL_{11}^Q$
$\pm f_2$	AS_{57}^I, AS_{57}^Q	POP_{57}^I, POP_{57}^Q	$REFL_{57}^I, REFL_{57}^Q$
$\pm (f_2 - f_1)$	AS_{45}^I, AS_{45}^Q	POP_{45}^I, POP_{45}^Q	$REFL_{45}^I, REFL_{45}^Q$
$\pm (f_1 + f_2)$	AS_{68}^I, AS_{68}^Q	POP_{68}^I, POP_{68}^Q	$REFL_{68}^I, REFL_{68}^Q$

Table 7.3: Probes that sense the light fields propagating within ET-LF with respect to the carrier frequency. The carrier is sensed by AS_{DC} at the interferometer’s output port, which is the light that propagates through the signal recycling mirror. The control sidebands, and the beats between the sidebands, are sensed at the same port demodulated at each relevant frequency, along with similar readouts sensing a small transmission of light through a folding mirror in the power recycling cavity (POP) and the light reflected from the interferometer (REFL).

7.4.7 Control signals

The sensing matrix, as introduced in section 7.3.3.3, can be calculated for ET-LF by exciting each degree of freedom with the driving coefficients shown in table 7.1 and measuring the response at each probe shown in table 7.3.

At this preliminary stage, a reasonable choice of error signals from the sensing matrix to use for control can be determined through a heuristic approach building upon knowledge gained from the control of the second generation detectors. As f_2 resonates in both recycling cavities, it samples the motion of the mirrors that influence MICH and SRCL as well as PRCL. Conversely, f_1 only strongly samples the motion of MICH and PRCL. Using demodulations at these frequencies, and combinations thereof, it should be possible to find a set of reasonably decoupled error signals for each degree of freedom.

7.4.7.1 Combination of readout quadratures

Table 7.3 shows signals at both the I and Q quadratures for all of the RF sideband frequencies considered. These quadratures can be combined electronically to produce an error signal with optimal gradient. The exact phase corresponding to the greatest magnitude is not important, as this can be influenced by technical factors such as the length of RF transmission lines, but the relative phase between maximum error signals from different degrees of freedom on the same sensor is. If error signals from different degrees of freedom have the same maximum gradient at the same or opposite phase, one cannot be minimised with respect to the other through appropriate choice of demodulation phase. On the other hand, if two degrees of freedom couple to a pair of I and Q sensors with equal magnitude but separate phase, they can be used to sense both degrees of freedom. In practice due to temperature drifts and other time-varying effects it is difficult to maintain the demodulation phase of a set of sensors to a precision better than around 1° [180], and so the readout quadratures chosen for each of the degrees of freedom of ET-LF should ideally be separated by many degrees.

7.4.7.2 Sensing matrix for ET-LF in detuned configuration

The sensing matrix for the detuned configuration given the readout ports defined in table 7.3 is shown in table 7.4 for mirror perturbations at dc. The suggested readouts for each degree of freedom are highlighted in bold and are described in the following text.

The gravitational wave signal will primarily affect DARM, and this is by design sensed by AS_{DC} . The common mode can be sensed at $REFL_{11}$. The MICH, PRCL and SRCL cavity error signals are difficult to separate due to the cross-coupling of control sidebands via the

Schnupp asymmetry and the sideband asymmetry created by the presence of a detuned signal recycling cavity. As the control system will be implemented in a LIGO CDS-style system [142], it will be possible to define in software error signals formed from linear combinations of different sensor signals that decouple other degrees of freedom from a particular readout, as discussed in section 7.3.3.3. In table 7.4, PRCL is dominant over MICH and SRCL at POP₁₁ and so this represents a good extraction point for the motion of the power recycling cavity. MICH couples strongly to a number of ports but alongside strong signals from the other degrees of freedom. Its strongest ports, AS_{DC} and AS₅₇, contain much larger signals from DARM due to the difference in finesse between the Michelson and the arm cavities, and this would be difficult to remove electronically given the phase degeneracy. A better readout could be to use POP₅₇ with the dominant PRCL signal suppressed with careful tuning of the demodulation phase. Residual PRCL coupling can be subtracted electronically.

SRCL is difficult to sense with $\pm f_1$ because they are not resonant in the signal recycling cavity, nor a single demodulation at $\pm f_2$ because these error signals are contaminated with contributions from MICH or PRCL given the cross-coupling facilitated by the Schnupp asymmetry. A possible sensing strategy could utilise the beats between sidebands found at REFL₄₅ or REFL₆₈ where the contribution from MICH and PRCL can be suppressed with suitable choice of demodulation phase. All SRCL signals contain an offset from zero due to the detuning, since $-f_2$ is not resonant when the signal recycling mirror is at its operating point (see figure 7.6). To use REFL₆₈ for SRCL control an offset equivalent to -4.2 mW will be necessary.

Table 7.5 shows the gradient of each degree of freedom's error signal at each of the suggested sensors. The values have been normalised with respect to the degree of freedom to be read out. This table shows that the hierarchical control techniques discussed in section 7.3.3 will be necessary in order to decouple the individual degrees of freedom.

Error signals corresponding to the suggested readouts for each degree of freedom are shown in figure 7.9. These are produced by calculating the power on each sensor at the relevant demodulation frequency as the mirrors are driven as shown in table 7.1 and represent the low-frequency limit of the transfer functions of each degree of freedom to each sensor.

The gradient of the AS_{DC} readout is zero when the arm cavities are tuned, consistent with figure 7.7, showing that some classical carrier light power is always present for functional dc readout. The SRCL error signal shows the operating point crossing at a power of 4.2 mW, necessitating an offset.

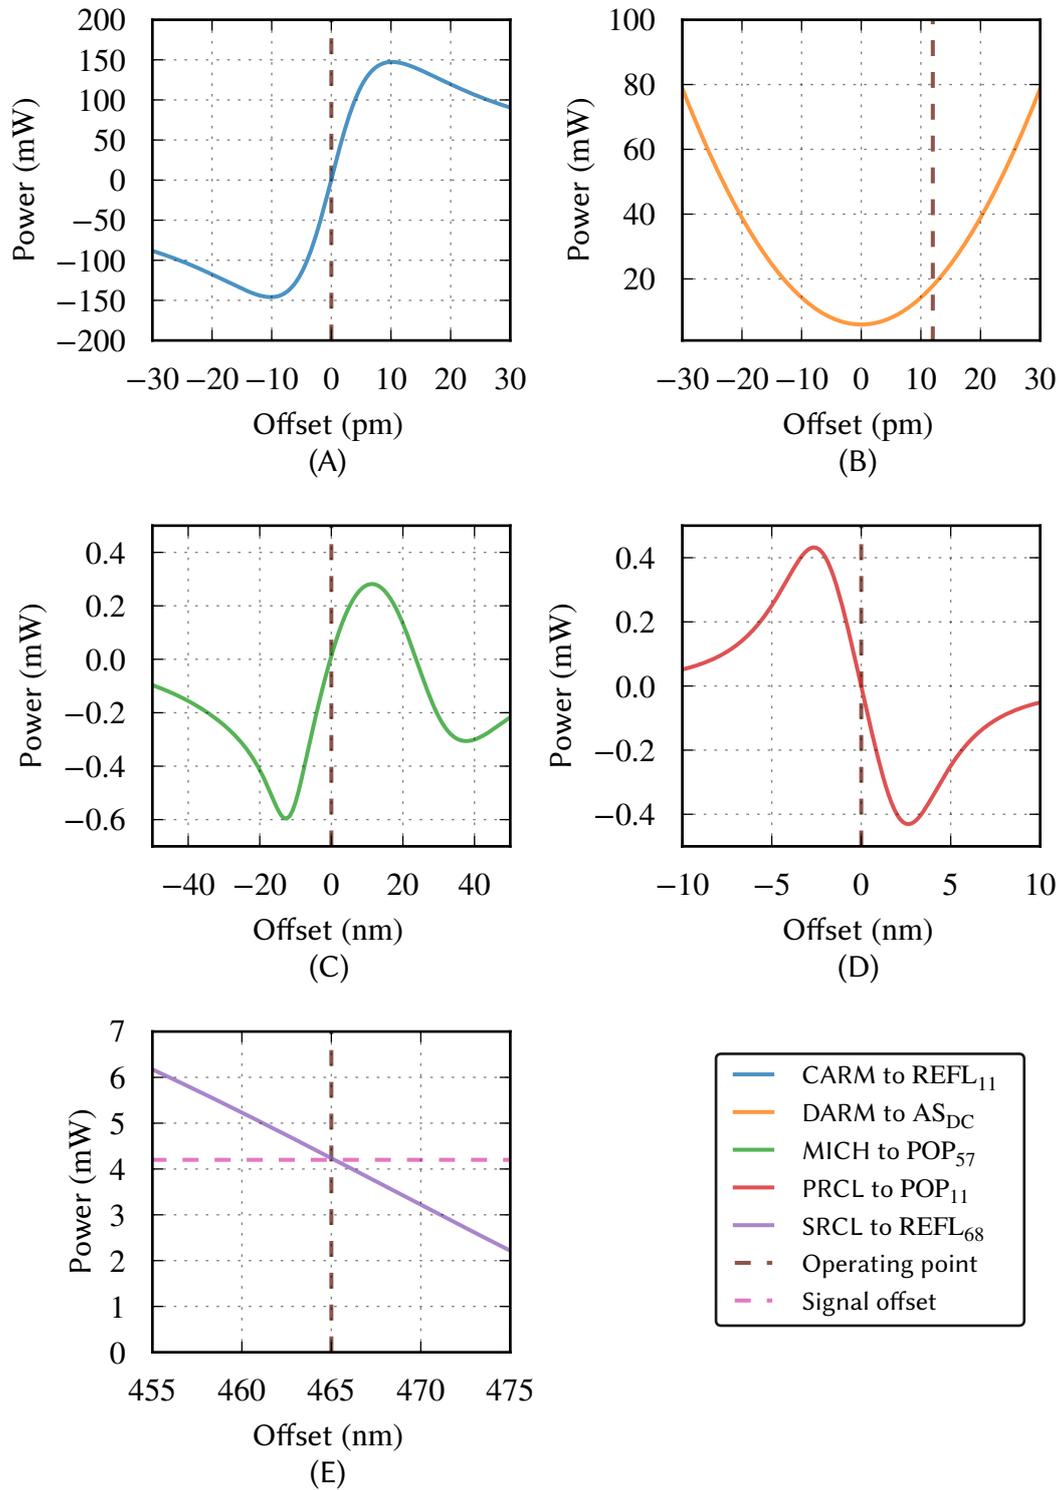


Figure 7.9: Sweeps through the zero-crossings of the chosen error signals in ET-LF in the detuned configuration. Each error signal is linear about the operating point, which ensures a simple, bipolar error signal is available for the purposes of controlling each associated degree of freedom. This linearity has a different range for each readout, with CARM and DARM requiring the greatest precision.

	CARM	DARM	MICH	PRCL	SRCL
AS_{DC}	1.44×10^6	4.99×10^8	1.17×10^6	1.21×10^5	1.39×10^4
AS₁₁	2.25×10^6 (-32.56°)	3.00×10^7 (55.83°)	6.25×10^4 (56.09°)	7.54×10^3 (77.99°)	1.09×10^3 (58.46°)
AS₅₇	1.76×10^8 (-112.30°)	3.98×10^8 (-162.68°)	8.39×10^5 (-163.43°)	6.03×10^5 (-78.31°)	6.73×10^4 (94.85°)
AS₄₅	4.71×10^4 (178.86°)	3.52×10^3 (107.72°)	3.43×10^4 (85.09°)	6.44×10^4 (61.41°)	3.41×10^3 (-146.01°)
AS₆₈	1.00×10^5 (-177.46°)	4.43×10^4 (-83.81°)	5.97×10^4 (43.44°)	5.13×10^4 (103.97°)	4.67×10^3 (-22.78°)
POP₁₁	7.61×10^7 (-144.47°)	6.88×10^5 (-144.49°)	8.70×10^2 (-26.41°)	2.55×10^5 (-36.40°)	3.08×10^1 (-81.49°)
POP₅₇	5.27×10^7 (99.38°)	4.72×10^5 (92.38°)	2.95×10^4 (50.72°)	1.26×10^5 (-130.50°)	1.08×10^4 (120.38°)
POP₄₅	2.01×10^2 (-51.56°)	9.34×10^1 (-127.43°)	1.24×10^3 (1.97°)	1.61×10^4 (88.80°)	1.07×10^3 (167.69°)
POP₆₈	4.94×10^2 (-22.60°)	9.93×10^1 (-15.51°)	1.07×10^3 (45.23°)	8.34×10^3 (-145.33°)	1.07×10^3 (-149.81°)
REFL₁₁	1.44×10^{10} (-0.01°)	1.31×10^8 (-0.01°)	2.34×10^5 (4.53°)	5.01×10^7 (-0.41°)	2.25×10^4 (45.05°)
REFL₅₇	4.63×10^9 (-75.19°)	5.69×10^7 (-86.52°)	1.47×10^5 (-83.24°)	1.58×10^7 (-76.05°)	1.67×10^4 (30.68°)
REFL₄₅	3.90×10^6 (0.94°)	1.31×10^6 (0.04°)	2.19×10^5 (163.21°)	1.28×10^6 (-171.17°)	2.03×10^5 (-25.15°)
REFL₆₈	4.04×10^6 (-101.97°)	1.32×10^6 (-101.79°)	2.25×10^5 (-76.80°)	1.40×10^6 (-53.54°)	2.03×10^5 (89.52°)

Table 7.4: Gradients of the error signals from each degree of freedom to each probe, in units of W m^{-1} , in ET-LF at dc. The suggested readout probes for each degree of freedom are shown in bold red. The I and Q quadratures of each heterodyne readout have been combined into a single magnitude and the phase representing the greatest slope and the phase at which it is achieved. They have maximum gradient at phase angles determined by the propagation of the control sidebands through the interferometer. Probes can be optimised to sense the motion of a particular degree of freedom by adjusting the angle at which the sensor I and Q quadratures are combined, but signals on sensors that contain strong signals from other degrees of freedom at nearby phase angles are difficult to use.

7.4.8 Sensitivity of the scheme

The sensitivity for ET-LF shown in figure 7.1 assumes that the interferometer contains squeezed vacuum input via two filter cavities in addition to the presence of seismic and other noise, and these features have not been modelled in this work. A comparison is shown in figure 7.10 between the quantum noise limited sensitivity for ET-LF in the absence of squeezing, as calculated by the tool used to present the results in the ET-D study [47], GWINC [26], and the quantum noise limited sensitivity of the AS_{DC} readout in this scheme modelled with Optickle. The reference curve from GWINC has been scaled to remove the

	CARM	DARM	MICH	PRCL	SRCL
REFL₁₁	1.0	9.1×10^{-3}	1.6×10^{-5}	3.5×10^{-3}	1.6×10^{-6}
AS_{DC}	2.8×10^{-3}	1.0	2.3×10^{-3}	2.4×10^{-4}	2.8×10^{-5}
POP₅₇	1.8×10^3	1.6×10^1	1.0	4.3	3.7×10^{-1}
POP₁₁	3.0×10^2	2.7	3.4×10^{-3}	1.0	1.2×10^{-4}
REFL₆₈	2.0×10^1	6.5	1.1	6.9	1.0

Table 7.5: Normalised gradients of the suggested sensors to be used for the control of ET-LF's degrees of freedom. Each row from table 7.4 corresponding to a sensor used to control a degree of freedom has been scaled by the gradient of the error signal corresponding to the respective degree of freedom. This shows the prominence of the other degrees of freedom on each sensor without having applied any of the other control techniques as discussed in section 7.3.3.

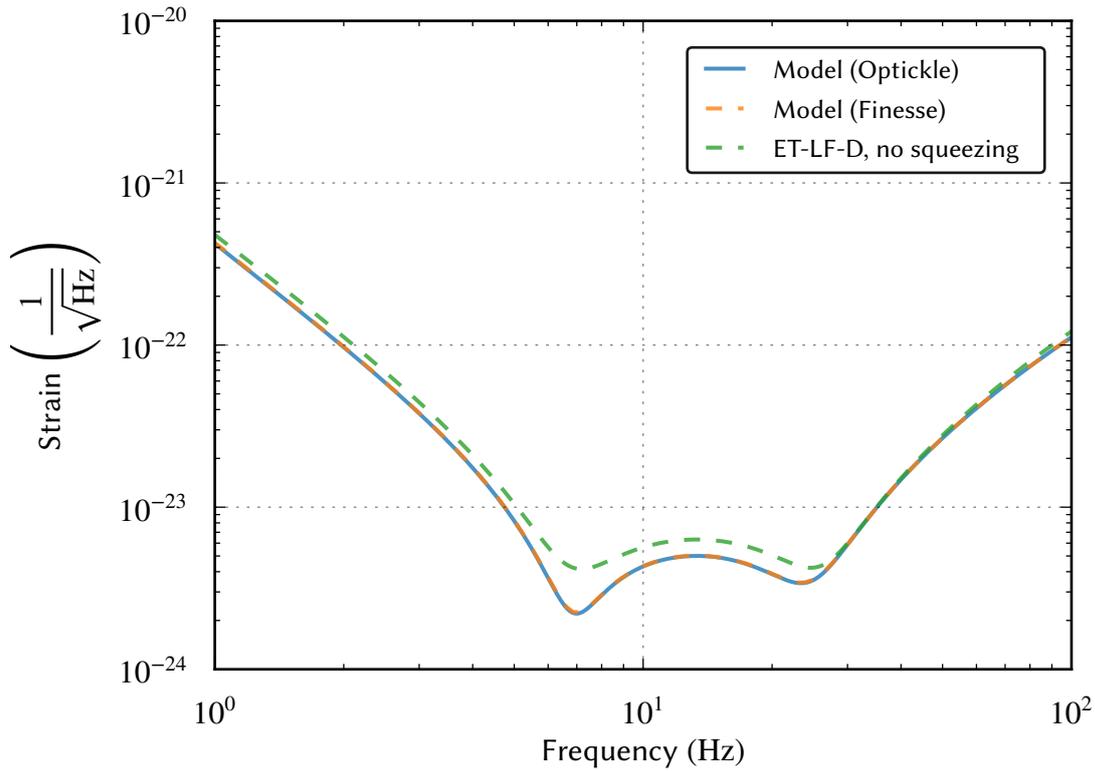


Figure 7.10: ET-LF quantum noise limited sensitivity using the conceptual control scheme, with no squeezed light injection. The reference curve from the ET-D study is shown next to the sensitivity calculated with the Optickle model developed in this chapter, and also a Finesse curve generated using identical parameters. The reference and simulation curves roughly agree, showing that the chosen parameters do not have a significant impact upon the sensitivity, though a difference in the noise calculations and the DARM offset assumed in this work creates the slight mismatch in the most sensitive region.

effect of squeezing injection, for comparison to the work presented here. Also shown is a simulated curve generated using *Finesse* (see appendix C.1.1) with identical parameters to those used in this work, showing good agreement with that of Optickle. The simulated curves show that the choice of parameters in this work do not negatively impact upon the design sensitivity of the interferometer. The difference in sensitivity between the new simulations and the reference in the region of the cavity pole (7 Hz) and the optical spring from the detuned signal recycling cavity (25 Hz) is due to the chosen DARM offset, resulting in different quantum shot noise at the photodetector, and the inclusion of additional losses arising from the filter cavities in the reference curve.

The parameters used in the proposed control scheme are shown in table 7.6 alongside the pre-existing parameters from the ET design study. It should be noted that the proposed scheme is not rigorously optimised, and does not consider a number of other control possibilities such as the use of secondary reflections arising from the anti-reflective coatings on the beam splitter and ITMs or the use of higher order combinations of f_1 and f_2 such as $3f$

Parameter	Symbol in text	Design study value	Updated value
Laser wavelength			1550 nm
Input power			3 W
Arm power			18 kW
ITM transmissivity			7000 ppm
ETM transmissivity			6 ppm
PRM transmissivity	T_{PRM}		4.6 %
SRM transmissivity			10 %
Signal recycling detuning			0.6 rad
Arm cavity lengths	L_X, L_Y		10 km
Power recycling cavity length	L_{PRCL}		310 m
Signal recycling cavity length	L_{SRCL}	310 m	311.585 m
Schnupp asymmetry	Δl_{SCH}	—	0.08 m
DARM offset	δL_{DARM}	—	12 pm
Control sideband frequencies	f_1, f_2	—	11 363 101 Hz, 56 815 505 Hz
Control sideband modulation depths		—	0.1 rad
Demodulation frequencies		—	$f_1, f_2, f_2 - f_1, f_1 + f_2$

Table 7.6: Updated parameters for ET-LF in the detuned configuration following the development of the conceptual sensing scheme in this chapter.

signals used in the lock acquisition sequence of Virgo [181] and Advanced LIGO [177]. It serves, however, as a first concept for the control of ET-LF proving that it can in principle be controlled in its detuned state.

7.5 Outlook and future work

By utilising an RF phase modulation scheme we have shown in this chapter that ET-LF can in principle be controlled with the presence of a 25 Hz signal recycling cavity detuning. Control noise issues remain unaddressed, as are more exotic control schemes such as the use of additional modulation frequencies or carriers. This section describes some future work that will be necessary to refine and improve the results presented here towards a comprehensive technical design.

7.5.1 Optimising the sensing matrix

In Advanced LIGO the conceptual control scheme was first tested at the Caltech 40 m prototype and it is probable that any technical design for ET-LF will require a similar test. At this stage a quantitative assessment of the performance that a particular control scheme might provide might take the form of a technique presented by Mantovani and Freise [182] developed for alignment control in Virgo, but suitable for longitudinal control. This involves the calculation of a *quality parameter* representing the controllability of a given set of sensors and degrees of freedom. This approach only makes sense when M represents the interferometer at its operating point, which means that the residual motion of controlled

degrees of freedom does not create a significant cross-coupling. This approach requires hierarchical gain to be simulated as part of a lock acquisition sequence, and so some effort will be required to design some feedback servos.

7.5.2 Switching between tuned to detuned operation

Transition from tuned to detuned signal recycling operating points and vice versa involves a technique which can maintain control of the interferometer as it transitions between two desired set points. When dual recycling was first demonstrated in suspended optics in the Garching prototype, it involved a varying frequency offset applied to the RF modulation sidebands as the tuning was changed [55]. This control technique was evolved in GEO 600 where a complicated sequence of actions [57] including an uncontrolled “jump” between two operating points [168] were performed to reach tuned mode from a detuned start point. In ET-LF it is expected that the signal recycling cavity finesse will be too high to allow for a previously demonstrated transition scheme, and so investigations are ongoing to model the impact that combinations of phase- and amplitude-modulated control sidebands added to the input light or *subcarriers* added to the squeezing injection port have on the lock acquisition and control of the signal recycling cavity at arbitrary detunings. Another possibility is to adapt the *arm length stabilisation* system developed for the lock acquisition of Advanced LIGO [177, 183], whereby a second carrier at a different wavelength is used to lock cavities. This takes advantage of the lower finesse of the second carrier’s wavelength in the cavities, allowing for a wider locking range. The cavities are first pre-stabilised using this second carrier before the main carrier is brought to resonance.

7.5.3 Sensing and control of seismic and gravity gradient noise

The Einstein Telescope facility will be located and designed to minimise the impact of seismic noise, but due to the sensitivity requirement for ET-LF the microseism must be suppressed from around 10^{-6} to 10^{-8} m $\sqrt{\text{Hz}}^{-1}$ at frequencies between 0.1 and 1 Hz to below 10^{-18} m $\sqrt{\text{Hz}}^{-1}$ by 2 Hz. This represents a signal difference of around 10^{10} , and current sensor electronics can typically only provide dynamic range in the region of 130 dB $\approx 3 \times 10^6$ without significant design effort. To control seismic noise in addition to being able to sense displacements at the required level, a sensor hierarchy will need to be developed.

There is some precedent for seismic isolation from the work carried out in current and past detectors. In Virgo, the seismic coupling in the superattenuator was suppressed with a local controller in addition to the global feedback from the main interferometric read-outs [184]. In Advanced LIGO, seismic pre-isolation is performed through the use of a

series of displacement and velocity sensors [185], and in the AEI 10 m prototype the addition of a *suspension platform interferometer* [186] is able to reduce seismic noise to the level of around $100 \text{ pm } \sqrt{\text{Hz}^{-1}}$ between 0.1 and 1 Hz [187]. Such a system would allow a high dynamic range global sensor to control the remaining motion. The particularly challenging aspect for ET-LF is that the suppression of this motion must occur over a bandwidth below 2 Hz, which makes the implementation of stable control filters extremely challenging. The results from the control of the AEI 10 m prototype and the advanced detectors will provide input to the technical design of the seismic isolation system for ET-LF.

7.6 Summary

The Einstein Telescope interferometers will create new challenges not previously encountered in the control of gravitational wave interferometers. In particular, the low frequency ET-LF detector, a dual-recycled Fabry-Perot Michelson interferometer, will have significant cross-couplings between signals from each of the degrees of freedom due to the detuned signal recycling cavity. We have introduced some of the techniques employed in the state of the art Advanced LIGO and Advanced Virgo detectors for the control of the five longitudinal degrees of freedom in a dual-recycled Fabry-Perot Michelson interferometer, and we have shown with simulations that these methods can also be applied to the ET-LF interferometer. The presented sensing matrix for ET-LF's longitudinal degree of freedom has reasonable error signals compatible with the experiment's sensitivity requirement. This result will be the basis for future work investigating the dynamics of the control system and the noise present at the site eventually selected for the facility.

Chapter 8

Conclusions and future work

At their most sensitive frequencies, current generation detectors are limited in sensitivity primarily by quantum and thermal noise. Improvements beyond the level already achieved in Advanced LIGO and soon to be achieved in Advanced Virgo and KAGRA will require significant research and development, with the “low hanging fruit” all but gone. This thesis presented techniques to potentially reduce the thermal and quantum noise for future detector facilities.

In chapter 3 we discussed the thermal noise arising from dielectric mirrors used in all previous and current gravitational wave detectors. This is a consequence of the presence of Brownian noise within the many layers of material forming the reflective coatings of the test masses. While reduction of this noise can be made through the development of new coating materials, progress in recent years has been frustrating. We introduced a grating mirror with a resonant waveguide structure as a possible alternative to dielectric coating layers, with calculations showing that the thermal noise these mirrors would produce at cryogenic temperatures would offer a factor 10 improvement over the coatings employed in existing detectors. One potential show-stopper with this type of mirror is, however, the potential for phase noise coupling arising from motion of the mirror transverse to the beam axis as seen with previous grating mirrors. The addition in this design of a resonant waveguide is intended to mitigate this phase coupling. We developed an experiment to verify that the phase noise coupling was suppressed. The results showed that, if indeed the coupling is present, the measurement errors place an upper limit on the coupling at one part in 17 000. While further experiments will be necessary to fully assess their suitability, these results suggest that waveguide mirrors are a potential option for thermal noise reduction in future detectors.

Chapters 4 to 6 discussed the *speed meter* topology as a way in which to reduce quantum noise in future gravitational wave detectors. Theory has shown that this design can provide a vast reduction in quantum radiation pressure noise beyond that of ubiquitous position

meters. Alongside some theory highlighting the reduced noise in this design, the ongoing proof-of-concept Sagnac speed meter experiment in Glasgow was introduced in chapter 4 with the goal of the experiment being to demonstrate a reduction in quantum radiation pressure noise over an equivalent Michelson interferometer.

One of the key experimental challenges faced with the Sagnac speed meter is the control of its longitudinal degree of freedom given its lack of sensitivity at low frequencies. Due to the speed meter's velocity response, the lack of a signal towards dc leads to long term drift from the set point due to seismic motion, creating a significant drop in sensitivity. This problem will only become worse in full-scale detectors utilising this topology. A solution to the control of this drift was presented in chapter 5 by blending a displacement signal with the main velocity signal at low frequencies, and a realistic simulation of the full control loop was presented to show that the drift was effectively removed without harming the reduction in radiation pressure noise. The implementation of this control system in the experiment once commissioned will serve as a test for its effectiveness in the correction of long term drifts, informing the technical implementation of detector-scale speed meters in the future.

Another aspect of research connected to the Sagnac speed meter was discussed in chapter 6 with regards to the use of a novel electrostatic drive for direct actuation upon the test masses within the experiment. This design has a number of advantages over the approach taken in current detectors, the foremost being the potential for reduced seismic noise coupling to the gravitational wave channel. This chapter outlined the development of the control apparatus for the actuator, with a particular focus on the high voltage electronics. The development of an amplifier with adequate voltage, dynamic range and noise and safety standards was presented, and calculations of the expected maximum displacement noise in the Sagnac speed meter experiment due to the noise of the electronics showed that the design meets the requirements. Tests with the electronics in-situ will help to characterise this actuator's performance.

The last chapter, 7, discussed the challenges faced with the control of the low frequency interferometer, ET-LF, as part of the Einstein Telescope detector. This is a proposed facility utilising two colocated *dual-recycled Fabry-Perot Michelson interferometers*, each respectively optimised for low and high frequencies, to be built by a European collaboration in the distant future. The detector will have a longer baseline and will use advanced materials and special techniques to reduce thermal and quantum noise beyond the current generation. The conceptual design for ET-LF recommended that its signal recycling cavity be detuned in order to enhance the sensitivity within a particular frequency band, but stopped short of discussing the control strategy. The use of detuning within dual-recycled Fabry-Perot Michelson interferometers has been shown to create control challenges not present with the equivalent tuned varieties, and so the purpose of the work presented in this chapter

was to show that the interferometer can indeed be controlled. Using an approach similar to that of the Advanced LIGO and Advanced Virgo detectors, we showed using numerical simulations that the interferometer is indeed controllable, albeit with the need for a number of control procedures to obtain suitable error signals. Future work will investigate the effect that environmental and technical noise sources will have on the controllability of the interferometer, and the investigate more exotic control approaches.

Appendix A

Interferometric fundamentals

Basic interferometers can be modelled algebraically with only a few simple rules. This appendix introduces some of the mathematics describing Fabry-Perot cavities and modulation used in the description of some of the control techniques discussed in the main text.

A.1 Interferometer fields

We can define the field amplitude of light across dimension x as

$$E = E_0 e^{-ikx}, \quad (\text{A.1})$$

using the same terms as shown in equation (2.1). Upon reflection from an optic, the reflected field is equal to the input field, E_{in} , scaled by the optic's field reflectivity, r :

$$E_r = r E_{\text{in}}. \quad (\text{A.2})$$

The transmitted field is similarly scaled by the optic's field transmissivity, t :

$$E_t = it E_{\text{in}}. \quad (\text{A.3})$$

Note the presence of complex coefficient i , equivalent to a multiplication by $e^{i\frac{\pi}{2}}$. This shows that the phase difference between reflected and transmitted fields is $\frac{\pi}{2}$, which is required for energy conservation. Here we choose to apply this phase shift to the transmitted field, but it is equally valid to apply, for instance, a negative, real coefficient to the reflected field (see appendix C.2.1.1).

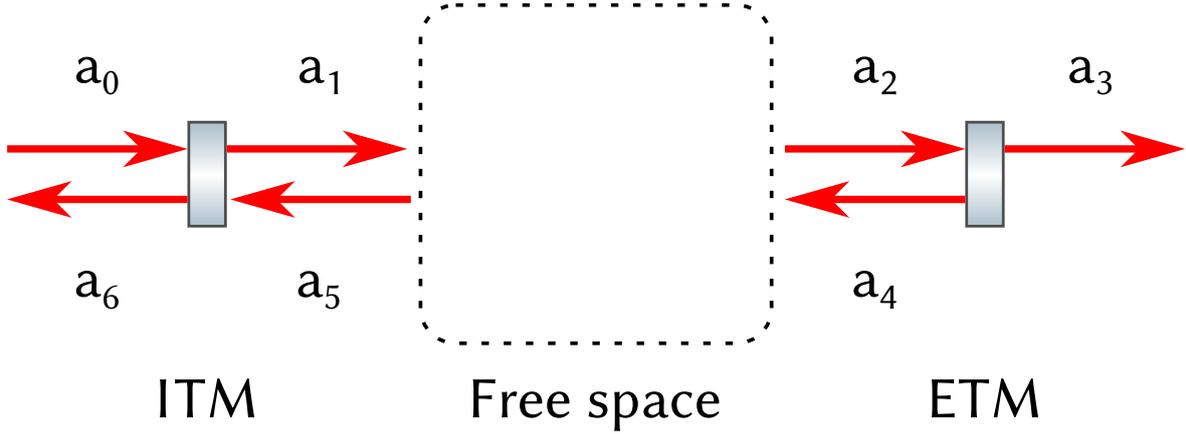


Figure A.1: Fabry-Perot cavity with its mirror input and output coefficients. Coefficients a_0 and a_6 are for the fields entering and leaving the cavity from the side of the ITM, while a_3 represents the field leaving the cavity on the side of the ETM. We assume the ETM has high reflectivity and so neglect the field entering at the ETM for simplicity. Coefficients a_1 , a_2 , a_4 and a_5 denote the circulating field at various points. The mirrors in the cavity are separated by free space through which the fields leaving each mirror's inside surface propagate.

A.2 Simple cavities

Combining equations (A.1) to (A.3) allows us to determine the field amplitude at different ports of a collection of mirrors. Combinations of mirrors produce *optical cavities*, which possess the property that they can, under certain conditions, accumulate light power. If we take the simplest of examples, the two-mirror Fabry-Perot cavity (see figure A.1), we can determine the fields coefficients present at the output nodes of the mirrors to be

$$\begin{aligned}
 a_1 &= it_1 a_0 + r_1 a_5 \\
 a_3 &= it_2 a_2 \\
 a_4 &= r_2 a_2 \\
 a_6 &= r_1 a_0 + it_1 a_5.
 \end{aligned}
 \tag{A.4}$$

Similarly, the coefficients at the input nodes within the cavity can be determined from the output nodes of the opposite mirrors and the separation x :

$$\begin{aligned}
 a_2 &= a_1 e^{-ikx} \\
 a_5 &= a_4 e^{-ikx}.
 \end{aligned}
 \tag{A.5}$$

Note the lack of a field input from the right side of the cavity in figure A.1. In general this field is present and has a small effect on the coefficients, but in gravitational wave detectors it is typically small by design. As this additional term adds mathematical complexity but

not significant additional clarity, it has been neglected.

The coefficients in equations (A.4) and (A.5) can be used to determine the field amplitude at different points of the interferometer. As there is only one input field in this example, the coefficients can be reduced to depend only on a_0 , r_1 and r_2 , t_1 and t_2 and x . The light reflected from the Fabry-Perot is an important field to know as this is used in many experiments to assist with interferometer sensing and control. Coefficient a_6 determines the reflected field, but this depends on a_0 and a_5 , the latter being one of the coefficients inside the cavity. Coefficient a_5 can have terms progressively substituted as follows:

$$\begin{aligned}
 a_5 &= a_4 e^{-ikx} \\
 &= a_2 r_2 e^{-ikx} \\
 &= a_1 r_2 e^{-2ikx} \\
 &= r_2 (it_1 a_0 + r_1 a_5) e^{-2ikx},
 \end{aligned} \tag{A.6}$$

where we find that the equation for a_5 depends on itself. This shows how the cavity operates: on each subsequent round trip, the cavity field is enhanced with addition field in the form of light transmitting through the first mirror. The field builds up until the amount of light entering the cavity is equal to the amount leaving. We can manipulate equation (A.6) to show this:

$$a_5 = a_0 \frac{it_1 r_2 e^{-2ikx}}{1 - r_1 r_2 e^{-2ikx}}, \tag{A.7}$$

where it is clear to see that the cavity coefficient depends on the input coefficient a_0 . The coefficient representing the reflected light from the interferometer is then just a_5 scaled by t_1 , with the addition of the light reflected before entering the cavity:

$$a_6 = a_0 \left(r_1 - \frac{t_1^2 r_2 e^{-2ikx}}{1 - r_1 r_2 e^{-2ikx}} \right). \tag{A.8}$$

With a_6 expressed in terms of a_0 , we can then calculate the reflected field amplitude as a function of the input field. The output field as a function of input field E_{in} is then simply

$$\begin{aligned}
 E_{\text{out}} &= a_6 E_{\text{in}} \\
 &= E_{\text{in}} \left(r_1 - \frac{t_1^2 r_2 e^{-2ikx}}{1 - r_1 r_2 e^{-2ikx}} \right) \\
 &= E_{\text{in}} R.
 \end{aligned} \tag{A.9}$$

The term R is the *transfer function* of the cavity, i.e. the ratio of the output field with respect to the input field. R behaves like the reflectivity of a simple mirror, but with the property that it is frequency dependent. Transfer functions are important figures of merit for cavities

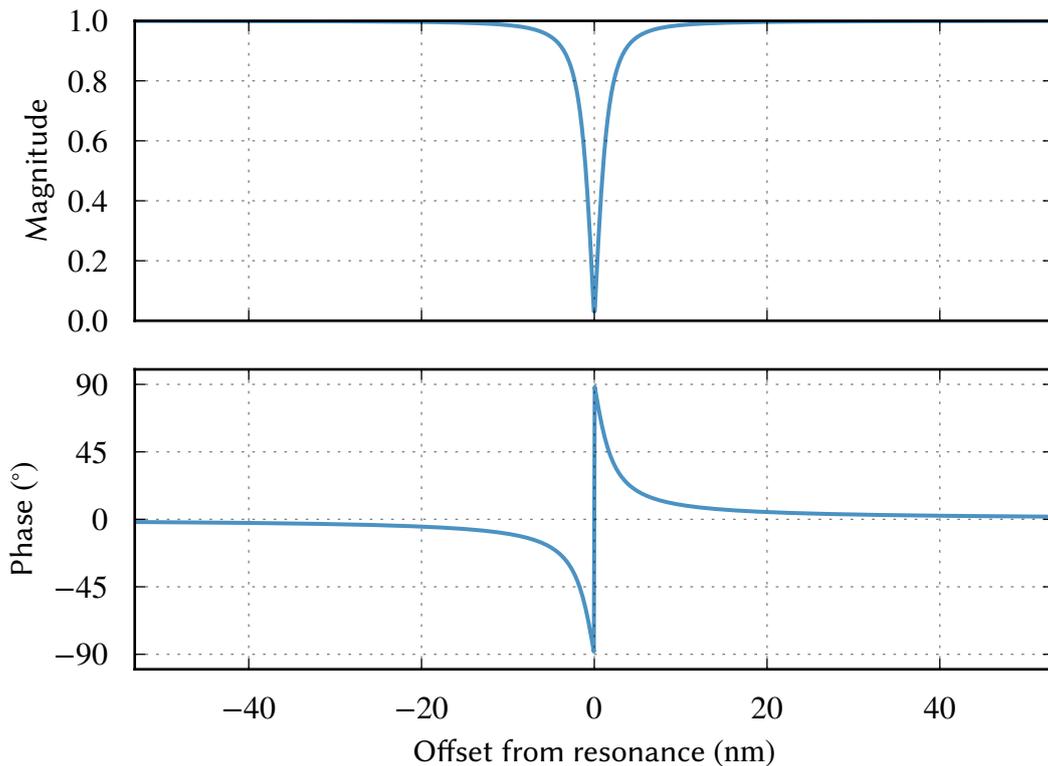


Figure A.2: Reflected field transfer function for a simple Fabry-Perot cavity. When the cavity is on resonance, the reflected field drops to almost zero as the field coupling coefficients of the ITM favour transmission over reflection. Away from resonance, the coupling coefficients favour reflection over transmission. In this example the mirror reflectivities are $r_1 = r_2 = 0.99$, which is a configuration known as an *impedance matched* cavity. A more complete description of the behaviour of cavities with different mirror reflectivities can be found in, for example, [23].

and they are used extensively throughout the main body of this work.

The reflected field transfer function R is shown in figure A.2. This shows that when the cavity resonance condition is not met, the field incident upon the ITM is mainly reflected. Close to resonance, the phase of the incident light and the ITM is favourable for transmission, leading to the majority of the light being coupled into the cavity. The width of the trough is determined by the cavity mirror transmissivities and reflectivities.

For asymmetric cavities, cavities with non-unity end mirror transmissivity and more complicated arrangements such as compound mirrors, the algebra involved in calculating transfer functions quickly becomes unwieldy and it is beneficial to utilise simulation tools (see appendix C).

A.2.1 Cavity figures of merit

Resonance is achieved within a Fabry-Perot cavity by making the microscopic length equivalent to an integer number of half-wavelengths for a given carrier. For this we can look at the ratio of the field in the cavity to the field entering it:

$$\frac{a_1}{a_0} = \frac{it}{1 - r_1 r_2 e^{-2ikx}}, \quad (\text{A.10})$$

but we see that the denominator contains a complex exponential with maxima in both the sine and cosine quadratures. In reality, we measure the light power with photodetectors, so taking the square of the modulus leads to

$$\begin{aligned} \left| \frac{a_1}{a_0} \right|^2 &= \frac{t^2}{1 - r_1 r_2 e^{2ikx} - r_1 r_2 e^{-2ikx} + r_1^2 r_2^2} \\ &= \frac{t^2}{1 - 2r_1 r_2 \cos(2kD) + r_1^2 r_2^2}, \end{aligned} \quad (\text{A.11})$$

which occurs when kD becomes an integer multiple of π . This means that either the wave number, representing carrier wavelength, or the cavity length can be controlled to achieve cavity resonance. Both of these techniques are used in interferometry, sometimes simultaneously.

As the resonance condition requires only an integer multiple of π , it is clear to see that resonance within the cavity is periodic as a function of laser frequency or length. A plot of the resonance condition as a function of length and frequency is shown in figure A.3. The resonant peaks occur for every half wavelength offset from the nominal length (10 m). Conversely, the resonant peaks also occur in terms of frequency offsets from the nominal carrier frequency ($f_0 = \frac{c_0}{\lambda}$). For this cavity, the frequency offset between resonant peaks is approximately 15 MHz. For a much longer cavity, such as that of Advanced LIGO—4 km—the frequency offset reduces to approximately 37.5 kHz. The difference arises from the fact that a change in laser frequency will change the wavelength, and a longer cavity can fit more waves, so a smaller frequency offset is required for the same change in effective length.

The frequency offset between successive resonant peaks is termed the *free spectral range* (FSR). This figure of merit provides some idea of the bandwidth a control system may need to be able to hold a cavity resonant, and it can be defined in terms of the speed of light c_0 and the cavity length L :

$$\text{FSR} = \frac{c_0}{2L}. \quad (\text{A.12})$$

Another figure of merit for a Fabry-Perot cavity is the *full-width at half-maximum* (FWHM), also known as the *linewidth*, which represents the width of a resonant peak at half of its maximum power, in units of frequency. As the width of the resonant peak is a function of

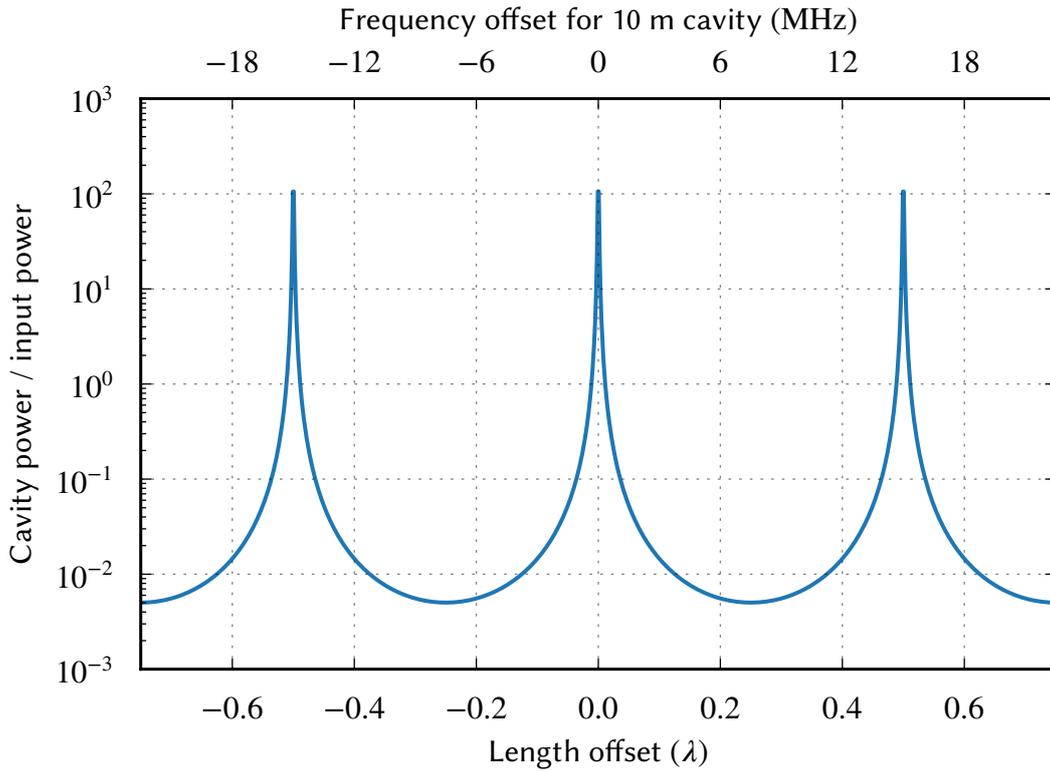


Figure A.3: Resonant enhancement of input light in a cavity. When the length offset from zero is an integer number of half-wavelengths, the cavity enhances the input power many times over. When the offset is between integer half-wavelengths, it is anti-resonant and the cavity power is much smaller than the input power. In this example, the mirror properties are $r_1 = 0.99$, $r_2 = 1$, $t = \sqrt{1 - r_1^2}$ and the cavity's macroscopic length is 10 m.

mirror reflectivity, the FWHM is defined as

$$\text{FWHM} = \frac{c_0}{\pi L} \sin^{-1} \left(\frac{1 - r_1 r_2}{2\sqrt{r_1 r_2}} \right). \quad (\text{A.13})$$

For Advanced LIGO the ITMs and ETMs have (power) reflectivities 1.4 % and 5 ppm, respectively, and so the 4 km arm cavity FWHM is around 83 Hz.

Sometimes used in the literature is the *half-bandwidth*, γ , equivalent to half of the FWHM:

$$\gamma = \frac{\text{FWHM}}{2} \approx \frac{c_0 T_{\text{ITM}}}{4L_{\text{RT}}}. \quad (\text{A.14})$$

In the frequency domain γ denotes the cavity's corner frequency, otherwise known as the *pole frequency*. In a Fabry-Perot cavity the response is essentially flat below this frequency, and degrading proportionally to frequency above it.

The ratio of the FSR to the FWHM defines the cavity *fineness*, \mathcal{F} , given as

$$\mathcal{F} = \frac{\text{FSR}}{\text{FWHM}} = \frac{\pi}{2 \sin^{-1} \frac{1-r_1r_2}{2\sqrt{r_1r_2}}}, \quad (\text{A.15})$$

which indicates the cavity's ability to store photons. It is closely related to the quality factor Q , namely the ratio of the energy stored in the cavity to the energy lost per radian of oscillation, via the relation [188]

$$Q = \frac{f_0}{\text{FSR}} \mathcal{F}. \quad (\text{A.16})$$

The finesse and quality factor indicate the time it would take for light to escape a resonant cavity in the event that the input light source were to be removed. A more practical use for the finesse figure is to quickly approximate the stored power in a cavity on resonance by multiplying it by the input power.

A.3 Signal sidebands

The field due to motion of a mirror Δx at a distance x can be determined by combining equations (A.1) and (A.2). As the light must travel distance x to get to the optic, the field picks up a factor of $e^{-ik(x+\Delta x)}$ in phase. Upon reflection, it gets a coefficient of r in amplitude:

$$E_r(x) = rE_0 e^{-ik(x+\Delta x)}. \quad (\text{A.17})$$

For clarity, we can express the constant x term as a static phase representing the carrier, separated from Δx via the wave number:

$$E_r(t) = rE_0 e^{-i(\omega_0 t + k\Delta x)}. \quad (\text{A.18})$$

It is possible to express any particular motion in the form of a series of sinusoidal functions. Expressing mirror motion as a single frequency sinusoid modulating the impinging field, we see

$$E_r(t) = rE_0 e^{-i(\omega_0 t + m \cos(\omega t))}, \quad (\text{A.19})$$

where we introduce m as the *modulation depth*, a dimensionless number expressing the strength of the mirror motion. With some algebraic manipulation involving Bessel functions, we can express this as [23]

$$E_r(t) = rE_0 e^{-i\omega_0 t} \sum_{n=-\infty}^{\infty} i^n J_n(m) e^{in\omega t}. \quad (\text{A.20})$$

When m is 0 all Bessel functions J_n are 0 except the first, which is 1; this shows that the field contains only the carrier component when there is no external modulation applied by a moving mirror. Conversely, when a non-zero modulation index and signal frequency is present, there exists a sum of sinusoidal functions as a product of the carrier, which we call *signal sidebands*. For small modulation depths $m \ll 1$, equation (A.20) can be approximated to

$$E_r = E_0 e^{-i\omega_0 t} \left(1 - \frac{m^2}{4} + i \frac{m}{2} (e^{-i\omega t} + e^{i\omega t}) \right). \quad (\text{A.21})$$

Here it is clear to see the presence of upper and lower sidebands at frequencies $\omega_0 \pm \omega$.

Amplitude modulation has a similar but not identical effect to phase modulation. Devices such as piezoelectric transducers for trimming laser outputs can perform amplitude modulation on a light field, and the effect can be expressed again in terms of the modulation depth and frequency:

$$E_r = E_0 e^{-i\omega_0 t} (1 + m \cos \omega t), \quad (\text{A.22})$$

and we can manipulate this expression to show the presence of exactly one upper and one lower sideband:

$$E_r = E_0 e^{-i\omega_0 t} \left(1 + \frac{m}{2} e^{i\omega t} + \frac{m}{2} e^{-i\omega t} \right). \quad (\text{A.23})$$

An example of the sideband structure created by amplitude and phase modulation is shown in figure A.4.

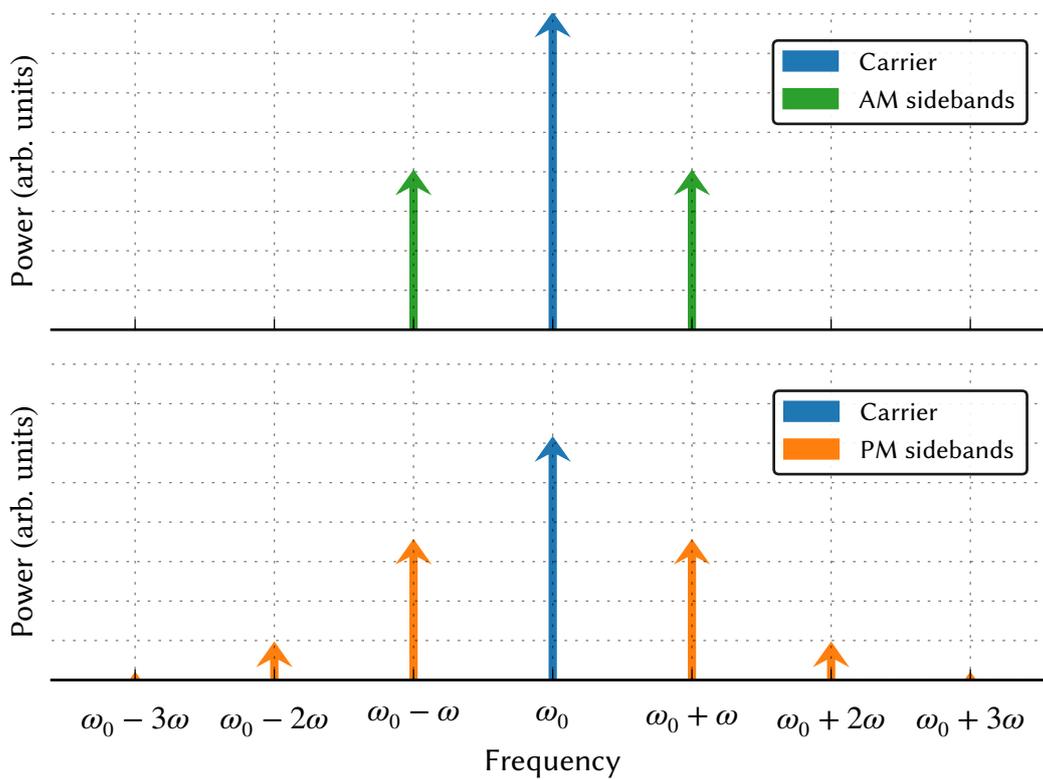


Figure A.4: Sideband structure for amplitude and phase modulation. The carrier, at frequency ω_0 , is present in the centre. Upper and lower sidebands appear offset from the carrier due to modulation. Amplitude modulation (AM) produces exactly two sidebands, whereas phase modulation (PM) produces infinitely many, though with drastically decreasing amplitude for higher orders. The vertical axis denotes the (arbitrary) power of each sideband, though in this example the modulation depth has been greatly exaggerated over that which would normally be used.

Appendix B

Control in the frequency domain

Control is an essential aspect of gravitational wave interferometry. This appendix contains some background information on some control topics covered within the main text.

B.1 Signal to noise ratio

Signal cannot be measured below the noise present on a photodetector. Maximum sensitivity can be achieved by maximising the ratio of signal power, S , to noise power, N , in the frequency band of interest. This is expressed as the *signal-to-noise* ratio (SNR),

$$\text{SNR} = \frac{S}{N}. \quad (\text{B.1})$$

When discussing controls, the standard representation of signal to noise is in units of *decibels* (dB), defined for signal power as

$$\text{SNR}_{\text{dB}} = 10 \log_{10} \left(\frac{S}{N} \right). \quad (\text{B.2})$$

As this representation is logarithmic, it is a useful for expressing both small and large signals and is therefore suitable for the presentation of noise sources across many orders of magnitude, as with interferometry.

For certain aspects of controls, for example with the expression of gain—an amplitude and not a power—the following equation is more suitable:

$$\text{Gain}_{\text{dB}} = 20 \log_{10} \left(\frac{\text{Output}}{\text{Input}} \right). \quad (\text{B.3})$$

The factor 20 is present here instead of 10 in the case of power quantities due to the lack of a square law dependency in the terms forming the ratio. Whether the expressed quantity

in dB refers to an amplitude or power must be made clear from the context.

B.2 Frequency representation of signals and noise

B.2.1 Fourier transform

A Fourier transform represents a time-varying signal in terms of its frequency components, i.e. as a series of sinusoidal waves of different frequency and amplitude. The Fourier transform $x(f)$ of $x(t)$ can be defined as

$$x(f) = \int_{-\infty}^{\infty} x(t) e^{-2i\pi ft} dt. \quad (\text{B.4})$$

Noise transients that appear in a detector have finite energy over a finite time, and can be entirely characterised by a Fourier transform of the time series in which the event occurred. Other forms of noise, however, cannot be represented in this way.

B.2.2 Spectral density

Apart from noise transients, the remaining noise sources within the interferometer tend to arise from stationary, random processes. This means that the noise source's *autocorrelation*—its self-similarity—is zero for all measurement times greater than zero. The energy of this noise approaches infinity as measurement time approaches infinity. In this circumstance, the Fourier transform of the underlying time-domain signal, as shown in appendix B.2.1, does not strictly exist. An alternative representation of a noise process is to represent the amount of work it performs per unit time: its power. The *power spectral density* is a representation of the power present within each frequency of a signal in the steady state.

The infinite time Fourier transform in equation (B.4) can be truncated to instead represent the frequency components within a certain window:

$$x(f) \approx \frac{1}{\sqrt{T}} \int_0^T x(t) e^{-2i\pi ft} dt, \quad (\text{B.5})$$

and the power spectrum of the signal measured by a photodetector is then

$$S_{xx} = \frac{1}{T} \int_0^T x(f)^2 dt. \quad (\text{B.6})$$

Remembering that photodetectors measure power (section 2.1.2), this means that the photodetector's power spectrum would have units of W^2/Hz . At the operating point, the photodetector's power is arranged in such a way as to be a linear with cavity mirror displacement, and so a more useful unit is the *amplitude spectral density* $A(f)$, which is simply the square root:

$$A(f) = \sqrt{S_{xx}(f)}, \quad (\text{B.7})$$

which for the photodetector would have units $W \sqrt{Hz}^{-1}$.

B.2.3 Estimation of spectral density

Equation (B.6) gives the formal definition of the power spectral density but not a practical means to measure it. To estimate the frequency components of a measured photodetector signal, we employ *spectral density estimation* techniques. The standard in experimental interferometry is Welch's method [189], which splits a measured time series into a series of segments which can overlap with adjacent segments before calculating Fourier transforms on each individual segment. The resulting Fourier transforms are recombined to produce the spectral density estimate.

The number of samples in a given period determines the lowest frequency resolved by the calculation. For instance, $N = 1000$ samples at a frequency $f_s = 1$ Hz would result in a lowest resolved frequency of $\frac{f_s}{N} = 10^{-3}$ Hz. *Segmentation* is a technique that can be used to trade bandwidth for resolution. Instead of using the full duration of the recorded data for one Fourier transform, therefore achieving resolution down to the lowest possible frequency, segments of shorter duration can be combined to produce better resolution at higher frequencies at the expense of lower frequencies. Specifically, the lowest resolved frequency becomes $\frac{f_s x}{N}$, where x is the number of segments the recorded data is divided into.

Once a segment is created, the resulting Fourier transform is applied to the time series with the assumption that the end cycles back to the start. If the data is noisy, or if the period is not an integer number of wavelengths of all the frequency components, then this creates discontinuities which lead to unphysical frequency domain content (*spectral leakage*). A *window function* can be applied to emphasise the signal in the middle of the segment at the expense of that at the edges. The window function typically used in the field is the Hanning window, and the effect that this has compared to a flat (*boxcar*) window is shown in figure B.1. A sine wave of frequency 2499 Hz and unity amplitude is recorded for a period of 1 s, sampled at a frequency of 10 kHz. The signal frequency is intentionally chosen to avoid an integer multiple of the sample frequency, which would remove the discontinuities at the edge of each segment. The power spectral density has been estimated using Welch's

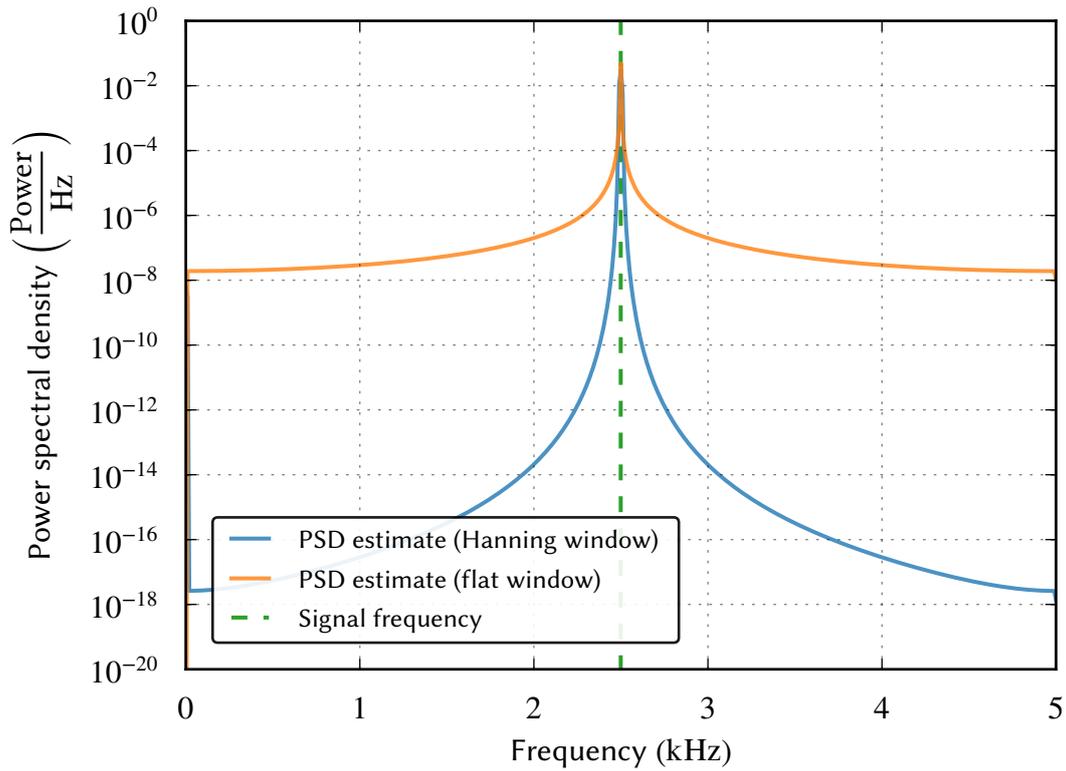


Figure B.1: The effect of windowing on a power spectral density estimate. The underlying signal time series is a sine wave of frequency 2499 Hz, and the sample rate is 10 kHz. The power spectral density estimates have been made using both Hanning windows, which de-emphasise the start and end of each segment to suppress discontinuities, and a flat window which performs no relative scaling of the data points. Both methods recover the signal's frequency, but in the former case the noise floor is greatly reduced.

method, with both Hanning and flat windows. The estimate for the noise floor in each case is drastically different because of the effect of segment discontinuities.

B.3 Root mean square amplitude

A real actuator or sensor has finite range, and the sum of the signal spectral density must be within this range to avoid clipping. The *root-mean-square* (rms) signal is equal to the sum of the absolute values of each of the frequency components. The rms representation is a useful form for the calculation of required actuator and sensor dynamic range (see, for example, chapter 5).

The following equation converts a power spectral density defined between f_{\min} and f_{\max} to an equivalent rms value:

$$x_{\text{rms}}^2 = \int_{f_{\min}}^{f_{\max}} S_{xx}(f) df. \quad (\text{B.8})$$

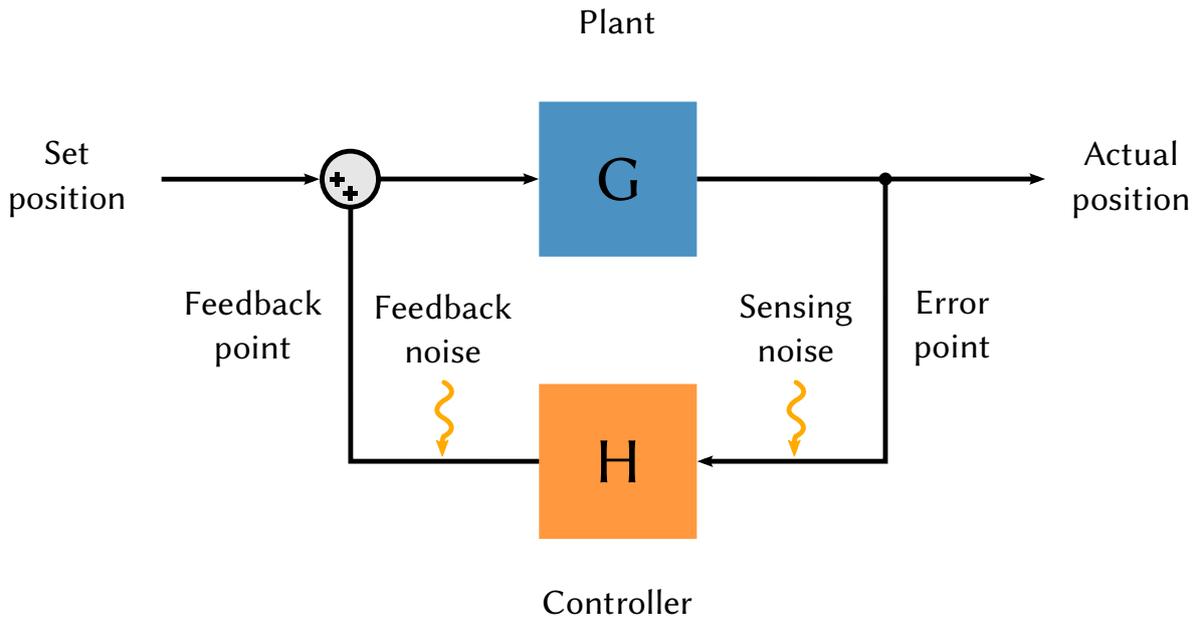


Figure B.2: A basic control loop. The plant contains the dynamics of the device to be controlled, such as an interferometer. The set position determines the desired point at which the plant should be held, and the error point shows the real position of the plant. The controller generates a corrective signal from the error signal input, and this is the feedback point. Noise enters the system at both the sensing and feedback points.

The rms representation of a spectral density is sometimes quoted with $f_{\min} = 1$ Hz and f_{\max} set to the sampling rate, and as such this value represents the signal “in a 1 Hz band”.

It is advisable to avoid pushing the rms signal applied to a sensor or actuator too close to its limit. Stationary random noise follows a well defined mean but can contain infrequent, larger noise transients allowed by Gaussian statistics. In such a case the instantaneous signal on a sensor or actuator might be greater than the rms. A good rule of thumb is to keep the rms signal expected at a sensor or actuator a factor of about 10 below its range to account for such events.

B.4 Control loops

A control loop can be used to sense the error in an interferometer from its operating point, and feed back signals to the actuators to correct it. We can in general split a control loop into two distinct parts: the *plant* G , which is the device under control, and the *controller* H , which is the device that senses the plant’s error and generates the corrective feedback. Noise entering the loop between the plant and the controller’s input (the *error point*) is termed *sensing* noise, and noise entering between the controller’s output and the plant (the *feedback point*) is called *feedback* noise, or, more commonly when discussing interferometers, *displacement* noise. Figure B.2 shows this scenario.

The controller cannot measure errors below its sensing noise and so sensing noise is not suppressed by the control loop.

B.4.1 Control loop figures of merit

A functioning control loop will suppress feedback noise by a level determined by the *open-loop gain*, defined as the product of G and H . This can be calculated by breaking the loop and taking a transfer function between the broken edges, and it shows the combined effect the system under control, its actuators and sensors have on their inputs at their outputs.

The *closed-loop gain* is the effect that the control loop has on the system when negative feedback is being applied. If the controller is able to sense errors and fully correct them, then the closed-loop gain is unity. The frequency domain representation of the closed-loop gain is useful to visualise at which frequencies the gain from the controller is not being applied: closed-loop gain higher than 1 shows that the system is not controlling the error point by matching it with equal magnitude and opposite sign, but rather following it. The closed-loop gain becomes particularly useful when comparing the effect of gain hierarchy, where multiple actuators are used to correct a single error point, as shown in section 5.2.8.3.

The neither the open- nor closed-loop gain figures show the explicit effect the controller has on the plant, which in the case of an interferometer would be the positions of the test masses. The *out-of-loop gain* provides this information, and is equal to the error point of the plant when the control loop is enabled. This figure is what is typically plotted in noise budgets such as the one shown in figure 2.1.

B.4.2 Control bandwidth

As discussed in appendix B.3, actuators and sensors have finite range. When designing a control system for an interferometer, or indeed any plant, the decision must be made between magnitude of the corrective feedback at some frequencies of interest, and the bandwidth over which the plant is to be controlled. For example, a ground-based interferometer typically oscillates with greatest amplitude at frequencies below 1 Hz due to seismic noise, as discussed in section 2.2.4.1. Meanwhile, the shot noise at higher frequencies is small enough such that the test masses do not move away from the operating point. Ideally, the finite actuator range on the test masses should be used to correct for displacements at low frequencies, where it is needed to keep the interferometer at the operating point. To achieve this, the control loop must limit the bandwidth to prevent feedback at high frequencies and enhance feedback at low frequencies. Figure B.3 shows the effect that two control servos have given the same ability (e.g. actuator range). By shaping the controller to increase the low frequency gain, the maximum frequency at which feedback is provided (the unity gain

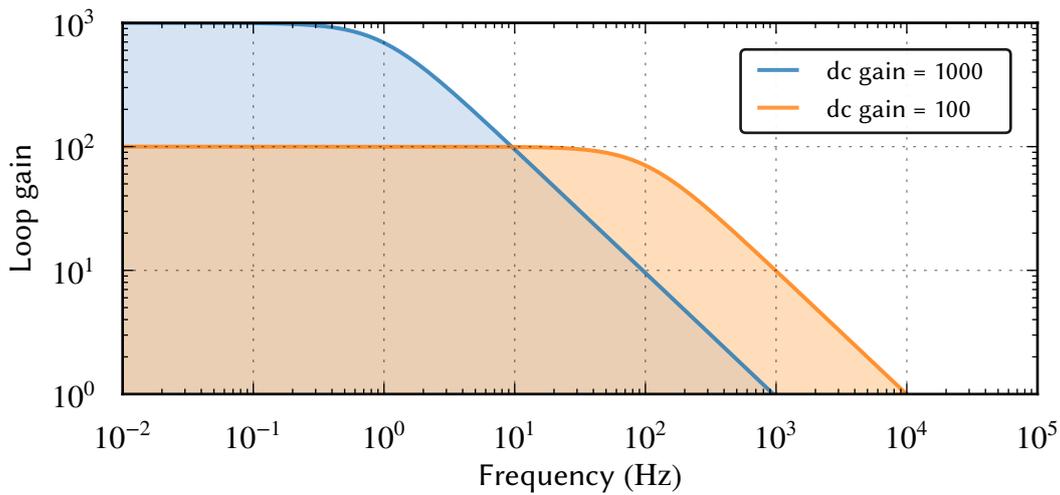


Figure B.3: Limiting a servo to enhance gain in a certain band. The figure shows the loop gain of two servos, each consisting of a simple low pass filter. The area under each transfer function is equal, and this represents the ability of the controller to make corrections to the system (for example, the range of an actuator). To enhance the gain by a factor of 10 at dc, the unity gain frequency has to be reduced by the same factor, meaning that the system will be controlled over a smaller bandwidth but with greater effort.

frequency) is necessarily reduced. The only way to enhance feedback effort whilst retaining bandwidth is to enhance the range of the actuators and/or sensors.

B.4.3 Stable loops

The stability of a control loop is determined by the controller's ability to generate a corrective signal that is opposite in sign to the disturbance. The interaction between the controller and the plant and its actuators and sensors can in some circumstances create situations where the feedback signal has the same sign as the error signal. If the feedback is of similar magnitude to the error signal, or greater, then this situation leads to positive feedback which makes the system uncontrollable, or *unstable*. In terms of magnitude and phase, this means that any points of unity gain in the transfer function's magnitude must not be coupled with a corresponding phase of -180° . To allow for some uncertainty in the system dynamics, a good rule of thumb in the implementation of the control system is to allow for a *phase margin* of around 35° [148], meaning that the phase at each unity gain point should not be lower than -155° .

Appendix C

Interferometer modelling

This appendix describes the process of modelling an interferometer in terms of its optical response and noise and introduces the two numerical simulation tools used throughout this thesis, Finesse and Optickle, introduced in appendix C.1. Appendix C.2 gives an introduction into how these tools work, and appendix C.3 describes modifications made to Optickle in order to calculate a particular type of noise amplitude.

C.1 Software

C.1.1 Finesse

Finesse is an extensive tool for simulating complex optical environments in terms of their transverse electromagnetic (TEM) modes. It was developed originally for use in GEO 600 [190] but has since been used for checks and tests in the design and commissioning of the advanced detectors [191, 192] and a variety of other experiments¹. Finesse's syntax allows for a number of different outputs to be generated for given excitations, for instance the transfer function from a set of optics to a set of sensors given an arbitrary signal. Finesse's strengths are its numerous supported optical components, speed of computation, and support for large numbers of paraxial modes.

C.1.2 Optickle

Optickle was originally created for the design of Advanced LIGO but has been extended to support a feature set capable of simulating arbitrary topologies. It is implemented in Matlab and is primarily intended to simulate plane wave behaviour of interferometers, however it

¹A list of known uses of Finesse for scientific research is maintained at <http://www.gwoptics.org/finesse/impact.php>.

also contains extensions to its code base for first order modes necessary to simulate misalignment effects. The output from Optickle is a series of matrices representing the transfer functions from optical degrees of freedom to sensors placed within the system, and the quantum noise upon those sensors.

C.2 Modelling an interferometer

The analytical calculation of the behaviour of interferometers beyond all but the most trivial examples is a complicated process and has to be performed with a particular configuration in mind. For example, adding or removing an optic from an analytical model of an interferometer may involve the addition of many new terms to the equations describing the main readout signals. Models for dual-recycled Fabry-Perot Michelson interferometers have been available for a number of years [167, 175, 193] but cannot be easily modified to account for optics beyond the ones considered in the model and the equations representing the readout signals have only been developed for the most important ports of the interferometer. In order to be able to calculate the signals present at any optic or probe within an interferometer, the most straightforward approach for an experimentalist is to use a numerical simulation tool.

Both Finesse and Optickle use broadly the same approach to compute results based on a technique used to model electronic circuits as for example with LISO, where the interferometer is described by an often large set of simultaneous equations. The tools take advantage of decades of development of numerical linear algebra tools to quickly solve these systems, with the results typically available in seconds.

The primary output from the tools is the calculation of field amplitudes or powers at ports of the interferometer at its operating point given a set of connected optics. In order to calculate these signals the simulation must map the effect that each optic's motion has onto the light fields within the interferometer, propagate these fields to each output and then calculate the corresponding electrical signals. These processes are described in more detail below in the context of Optickle, but the approach taken in Finesse is similar.

C.2.1 Optics

An *optic* refers to any component within the interferometer which has an effect on the light's amplitude, phase or frequency. Apart from mirrors and beam splitters, components such as lasers, EOMs and Faraday isolators can all be handled in the same way via matrices which translate inputs to outputs. Such *transfer matrices* define the signal behaviour of the light fields entering and exiting the component from arbitrary directions. For example, the

transfer matrix of a simple two-surface mirror can be defined as [23]

$$M = \begin{bmatrix} it & r \\ r & it \end{bmatrix}, \quad (\text{C.1})$$

where r and t represent the amplitude reflectivity and transmissivity of the mirror, with the condition $r^2 + t^2 = 1$, assuming no loss.

Using the first mirror in figure A.1 as an example, the inputs a_0 and a_5 map to outputs a_1 and a_6 as

$$\begin{bmatrix} a_1 \\ a_6 \end{bmatrix} = M \begin{bmatrix} a_0 \\ a_5 \end{bmatrix}, \quad (\text{C.2})$$

which can be re-expressed as individual transfer functions identical to those shown in equation (A.4):

$$\begin{aligned} a_1 &= ita_0 + ra_5 \\ a_6 &= ra_0 + ita_5. \end{aligned} \quad (\text{C.3})$$

C.2.1.1 Reflection phase convention

To conserve energy, a phase change must be applied to either the reflected or transmitted fields at an optical surface. In Finesse, the convention is such that this phase is added to the transmitted beam in the form of an imaginary coefficient. Optickle uses the convention that the reflection coefficient from the front of a mirror is $-r$, and from the rear it is r , and transmission is always t . Both conventions are valid, with Finesse's definition closer to the real effect that dielectric coatings would have on the transmitted light, and Optickle's more consistent with an infinitely thin optic.

This difference has no effect on the fidelity of the simulations because only the relative phase between reflected and transmitted fields is important for the calculation of transfer functions, noise and control signals. One practical difference is that in Finesse one arm of a Michelson interferometer must be tuned by 90° with respect to the other to achieve cancellation of the carrier at the output port of the beam splitter. In Optickle, due to the reflection phase convention this happens without special tuning.

C.2.2 Propagation

Propagation through free space is in general defined by equation (A.1), and in matrix form it is:

$$M = e^{-ikx} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (\text{C.4})$$

for wave vector k and distance x . In Finesse and Optickle, however, this behaviour is slightly different. Interferometers have path lengths of many metres, whereas the wavelength of the light being modelled is typically microscopic. For a cavity to be resonant its mirrors must be separated by an integer number of half-wavelengths, and so the 4 km Fabry-Perot cavities of Advanced LIGO would actually have to be defined with length $3\,759\,398\,497\lambda = 4000.000\,000\,81$ m given its 1064 nm wavelength. To make the creation of interferometer configurations easier, the simulation tools instead take the macroscopic propagation length and round it to the nearest integer number of wavelengths for the carrier, whereupon the phase difference from propagation is zero². To model the effects of non-zero phase propagation, such as a detuned cavity, the optics have additional phase tuning factors present within their transfer matrices. By separating these macroscopic and microscopic phase effects, issues with numerical precision can be avoided.

C.2.3 Fields

As spaces are defined as zero-phase propagation, the light between optics can be modelled with a single amplitude for each mode within the interferometer. In addition to the carrier, any signal sideband or control sideband present within the interferometer is considered a mode, as well as vacuum fields entering at points of loss as described in section 2.2.1, and so there can be many tens of figures representing the light between any two optics. These can be considered as the interferometer's *degrees of freedom*, and the propagation of each field through the interferometer can be modelled individually. A point in the interferometer represented by a field is termed a *field evaluation point*.

C.2.4 Drive and field maps

The calculation of signals at photodetectors requires the calculation of the field amplitudes within the interferometer, which can be determined by computing the steady-state solution of the optical system defined within a matrix mapping each field to each other field.

The process of calculating this *interferometer* matrix starts with the creation of the *field to field* matrix, which maps the transfer function between each of the fields within the interferometer with stationary optics without the presence of signal or control sidebands. This matrix allows the propagation of input light from lasers or vacuum injection to an arbitrary part of the interferometer to be calculated.

The effects of the mechanical degrees of freedom of a mirror or the electrical degrees of freedom of for instance an EOM on the light can be described by a *drive to field* map.

²In *most* cases. This is different for certain modes in which Finesse can run, controlled by the *phase* command.

Encompassed within this map are the amplitude and phase effects upon the carrier and sidebands cause by for example the motion of a mirror in the longitudinal direction. Similarly, the *field to drive* map, encompassing the effect of fields on the mechanical degrees of freedom of optics, allows the effect of radiation pressure to be handled properly.

Once the various maps between the light inputs, the mechanical and electrical drives and the optics have been calculated, they can be combined together in the form of a block-diagonal matrix \mathbf{M}_{AC} representing the transfer functions between the carrier, signal and control sidebands at each field evaluation point to each other field evaluation point.

C.2.5 Calculation of field amplitudes

The field amplitudes within the interferometer are of course determined by the excitation of the interferometer by external light injection, but in general they are also influenced by the signal sidebands produced by the modulation of optics within the interferometer at non-zero frequencies. The field amplitudes within the interferometer therefore depend not only on the excitation but also on the existing field amplitudes, analogous to feedback systems. In the initial state these fields are zero and so the interferometer's field amplitude vector is simply equal to the excitation vector, i.e. $\vec{v}_{AC} = \vec{v}_{exc}$. The stored light will increase until eventually the injected excitation is equal to the light power lost in the interferometer. Once this condition is reached the interferometer is in its steady-state, and the matrix of field equations this represents is the required input to the calculation of readout signals.

The steady state condition can be solved numerically using matrix inversion. As described above, the field amplitudes can be described as the current amplitudes plus the input:

$$\vec{v}_{AC} = \mathbf{M}_{AC} \vec{v}_{AC} + \vec{v}_{exc}, \quad (C.5)$$

where \mathbf{M}_{AC} is the interferometer matrix specified earlier. This equation can be solved as such:

$$\vec{v}_{AC} = \frac{\vec{v}_{exc}}{\mathbb{1} - \mathbf{M}_{AC}}. \quad (C.6)$$

Since \mathbf{M}_{AC} is a matrix and \vec{v}_{AC} and \vec{v}_{exc} are vectors, the problem can be represented as the equation

$$\vec{v}_{AC} = (\mathbb{1} - \mathbf{M}_{AC})^{-1} \vec{v}_{exc}, \quad (C.7)$$

where $\mathbb{1}$ is the identity matrix. The calculation of the field amplitudes in the interferometer therefore becomes a task of finding the inverse of $\mathbb{1} - \mathbf{M}_{AC}$, which is a problem for which many optimised algorithms have been developed.

C.2.6 Probe signals

With the steady-state field amplitudes, the signals produced by the interferometer can be determined with the application of a *probe matrix* $\mathbf{M}_{\text{probe}}$ which maps the fields in the interferometer to its probes. Since the fields amplitudes are determined for every wavelength under consideration, it is possible to calculate the signals that would appear on photodetector circuits implementing RF demodulation. The probe matrix contains complex amplitudes to transform the fields at the location of the probe by the required amount given the demodulation frequencies and phase angles. The probe signals are therefore defined as

$$\vec{v}_{\text{probe}} = \mathbf{M}_{\text{probe}} (\mathbb{I} - \mathbf{M}_{\text{AC}})^{-1} \vec{v}_{\text{exc}}. \quad (\text{C.8})$$

C.2.7 Calculation of transfer functions

The operation of calculating the probe signals from the field amplitudes in the interferometer can be repeated for arbitrary frequencies of excitation to produce a three-dimensional drive-to-probe transfer matrix. This represents the transfer function from each optic's degree of freedom to each probe. As such, the signal from a particular set of mirror excitations can be constructed via a linear combination of the transfer functions representing the degrees of freedom of individual optics. The differential arm degree of freedom transfer function for a Michelson interferometer to its asymmetric port, for instance, can be calculated by extracting the transfer function of each end test mass to a probe situated at the asymmetric port and taking the difference of the two, which gives the signal due to the length change equivalent to that shown in equation (7.1).

C.2.8 Probe quantum noise and sensitivity

The quantum noise calculations are similar in Finesse and Optickle, and they are both based on the work by Corbitt et al. [194], which ultimately derives from the two-photon formalism by Caves and Schumaker [129, 130]. Quantum noise is calculated using this technique by propagating two photons through the interferometer from each point of noise entry, in much the same way as signals are propagated. One photon represents the amplitude quadrature and the other represents the phase quadrature, and so appropriate quantum noise limited signals can be properly derived for any readout quadrature. A simpler, "one-photon" calculation involving the propagation of a single complex number representing the amplitude and phase photons is used in Optickle and Finesse [195, 196].

The sensitivity of a probe within an interferometer model is simply the noise on that probe divided by the relevant transfer function. When only quantum noise is considered, the sensitivity is *quantum noise limited*.

C.3 Calculation of field transfer matrices in Optickle

In the process of determining a probe signal in Optickle, the quadrature sum of the field amplitudes immediately in front of the probe is computed and the phase information contained within these fields is lost. Similarly, transfer functions from drives to probes are provided, but not transfer functions from drives to fields.

In order to calculate the cross-correlation spectral density required for the calculation of the optimal filter in section 5.5, the complex field and drive transfer matrices, \mathbf{M}^{ff} and \mathbf{R} , respectively, must be extracted from Optickle indirectly. Optickle's calculation of the quantum noise at each probe within the interferometer uses field to field and drive to field matrices, but because the quantum noise and drive excitations are not necessarily unity, these matrices are not transfer matrices. In order to obtain \mathbf{M}^{ff} the code which computes the quantum noise at each probe has to be modified to instead inject quantum noise at open ports with unity amplitude. Similarly, \mathbf{R} can be computed by setting the drive amplitudes to unity. The modified source code is publicly available [197].

Acronyms

ac Alternating current, though sometimes this refers to non-zero frequencies. 72, 89, 125, 127, 128

ADC Analogue-to-digital converter. 28, 75, 85–89, 97, 99, 118

BHD Balanced homodyne detector. 71, 72, 74, 78–85, 87, 88, 94–97, 99, 101–103, 106–108

BS Beam splitter. 147

CARM Common arm cavity degree of freedom. 145–147, 149, 150, 153, 162

CDS LIGO Control and Data Acquisition System. 47, 75, 84–87, 89, 90, 94, 96, 97, 103, 104, 109, 118–126, 129–131, 133, 161

CMOS Complementary metal-oxide semiconductor. 122, 126

CMRR Common-mode rejection ratio. 119

DAC Digital-to-analogue converter. 28, 75, 86–89, 97, 100, 120, 131, 133

DARM Differential arm cavity degree of freedom. 11, 15, 145–150, 153, 156–162, 164, 165

dc Direct current, though sometimes this refers to zero frequency, i.e. 0 Hz. 14–16, 23, 32, 45, 59, 61–63, 67, 72, 81–83, 94, 96, 101, 115, 123, 131, 143–148, 150, 155, 157, 160, 161, 163, 170, 189

eLISA Evolved Laser Interferometer Space Antenna. 8

EM Electromagnetic. 3

ENOB Effective number of bits. 86

EOM Electro-optic modulator. 43, 44, 46, 47, 142, 145, 153, 158, 192, 194

ESD Electrostatic drive. 74, 76, 90–96, 113–119, 126, 128, 132–134

- ET** Einstein Telescope. xiii, xviii, xix, 26, 36, 140–143, 147, 150–160, 162–167, 170
- ETM** End test mass. 26, 35, 36, 38, 40–46, 48–53, 56, 58, 68, 70, 71, 74, 82, 90–94, 96, 109, 114–118, 132, 134, 145, 147, 151, 165, 174, 178
- Finesse** Frequency domain interferometer simulation software. 51, 73, 78, 164, 191–194, 196
- FSR** Free spectral range. 152–154, 158, 177, 179
- FWHM** Full width at half maximum. 46, 81–83, 111, 177–179
- GPS** Global positioning system. 75
- GW** Gravitational wave. xv, xxii, 1, 2, 29, 138
- GWINC** Gravitational wave interferometer noise calculator. 18, 163
- HF** High frequency. 140, 141, 143
- HR** Highly reflective. 147
- HV** High voltage. 116, 118–136
- ITM** Input test mass. 26, 43, 46, 49, 67, 68, 70, 71, 75, 82, 142, 144, 146, 147, 151, 164, 165, 174, 176, 178
- KAGRA** Kamioka Gravitational Wave Detector. 1, 7, 8, 20, 21, 36, 154, 169
- LED** Light-emitting diode. 122
- LF** Low frequency. xiii, xviii, xix, 26, 36, 140–143, 147, 150–155, 157–160, 162–167, 170
- LIGO** Laser Interferometer Gravitational-Wave Observatory. xxii, 1, 2, 4, 7, 8, 17, 18, 21, 26, 30–32, 34, 36, 74, 75, 91, 113, 114, 137, 139–141, 143, 149, 151, 154, 156, 161, 165–167, 169, 171, 178, 191, 194
- LISO** Linear Simulation and Optimization of Analog Electronic Circuits. 120, 121, 128, 131, 192
- MICH** Michelson degree of freedom. 146, 147, 150, 160, 161
- MOSFET** Metal-oxide-semiconductor field-effect transistor. 123
- PDH** Pound-Drever-Hall. 39, 43, 45–47, 80, 101–104, 106–110

PRCL Power recycling cavity degree of freedom. 145–147, 150, 158, 160, 161

PRM Power recycling mirror. 146, 147, 165

PZT Piezo-electric transducer. 43, 47, 48

QND Quantum non-demolition. 32, 64, 68, 70, 108

RF Radio frequency. 43, 46, 47, 101, 143, 153, 160, 165, 166, 196

rms Root-mean-square. 48, 83, 91, 92, 96, 99, 103, 105, 186, 187

SNR Signal-to-noise ratio. 14, 17, 183

SQL Standard Quantum Limit. 24, 25, 32, 33, 60, 64, 69, 73, 139

SRCL Signal recycling cavity degree of freedom. 146–148, 150, 160, 161

SRM Signal recycling mirror. 146, 147, 165

TEM Transverse electromagnetic. 191

UTC Coordinated Universal Time. 1

WGM Waveguide mirror. 36–43, 48, 49, 51–56, 58

Bibliography

- [1] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116 (6 2016), p. 061102. DOI: 10 . 1103/PhysRevLett . 116 . 061102.
- [2] B. P. Abbott et al. “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence”. In: *Phys. Rev. Lett.* 116 (24 2016), p. 241103. DOI: 10 . 1103/PhysRevLett . 116 . 241103.
- [3] B. P. Abbott et al. “Tests of General Relativity with GW150914”. In: *Phys. Rev. Lett.* 116 (22 2016), p. 221101. DOI: 10 . 1103/PhysRevLett . 116 . 221101.
- [4] B. P. Abbott et al. “Astrophysical Implications of the Binary Black-hole Merger GW150914”. In: *Astrophys. J.* 818.2 (2016), p. L22. DOI: 10 . 3847/2041-8205/818/2/L22.
- [5] B. P. Abbott et al. “Binary Black Hole Mergers in the First Advanced LIGO Observing Run”. In: *Phys. Rev. X* 6 (4 2016), p. 041015. DOI: 10 . 1103 / PhysRevX . 6 . 041015.
- [6] The LIGO Scientific Collaboration and The Virgo Collaboration. “Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo”. In: *Living Rev. Relat.* 19.1 (2016). DOI: 10 . 1007/1rr-2016-1.
- [7] B. P. Abbott et al. “Localization and Broadband Follow-up of the Gravitational-wave Transient GW150914”. In: *Astrophys. J.* 826.1 (2016), p. L13. DOI: 10 . 3847/2041-8205/826/1/L13.
- [8] B. S. Sathyaprakash and B. F. Schutz. “Physics, Astrophysics and Cosmology with Gravitational Waves”. In: *Living Rev. Relat.* (2009). DOI: 10 . 12942/1rr-2009-2.
- [9] J. Weber. “Detection and Generation of Gravitational Waves”. In: *Phys. Rev.* 117 (1 1960), pp. 306–313. DOI: 10 . 1103/PhysRev . 117 . 306.

- [10] M. Pitkin, S. Reid, S. Rowan, and J. Hough. “Gravitational Wave Detection by Interferometry (Ground and Space)”. In: *Living Rev. Relat.* 5 (2011). DOI: 10 . 12942 / 1rr-2011-5.
- [11] G. M. Harry, J. L. Houser, and K. A. Strain. “Comparison of advanced gravitational-wave detectors”. In: *Phys. Rev. D* 65 (8 2002), p. 082001. DOI: 10 . 1103 / PhysRevD . 65 . 082001.
- [12] V. I. Pustovoit and M. E. Gertsenshtein. “Gravitational Radiation by a Relativistic Particle”. In: *Ĵ. Exp. Theor. Phys.* 15.1 (1962), p. 116.
- [13] G. E. Moss, L. R. Miller, and R. L. Forward. “Photon-Noise-Limited Laser Transducer for Gravitational Antenna”. In: *Appl. Opt.* 10.11 (1971), pp. 2495–2498. DOI: 10 . 1364 / AO . 10 . 002495.
- [14] R. Weiss. *Electromagnetically coupled broadband gravitational antenna*. Tech. rep. MIT, 1972.
- [15] P. R. Saulson. “If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves?” In: *Am. Ĵ. Phys.* 65.6 (1997), pp. 501–505. DOI: 10 . 1119 / 1 . 18578.
- [16] The Einstein Telescope Collaboration. *Einstein gravitational wave telescope conceptual design study*. Tech. rep. European Gravitational Observatory, 2011.
- [17] S. Dwyer et al. “Gravitational wave detector with cosmological reach”. In: *Phys. Rev. D* 91 (8 2015), p. 082001. DOI: 10 . 1103 / PhysRevD . 91 . 082001.
- [18] The LIGO Scientific Collaboration. “Exploring the Sensitivity of Next Generation Gravitational Wave Detectors”. In peer review.
- [19] P. Amaro-Seoane et al. “Low-frequency gravitational-wave science with eLISA/NGO”. In: *Class. Quant. Grav.* 29.12 (2012), p. 124016. DOI: 10 . 1088 / 0264 - 9381 / 29 / 12 / 124016.
- [20] B. J. Meers and K. A. Strain. “Modulation, signal, and quantum noise in interferometers”. In: *Phys. Rev. A* 44 (7 1991), pp. 4693–4703. DOI: 10 . 1103 / PhysRevA . 44 . 4693.
- [21] B. Willke et al. “The GEO 600 gravitational wave detector”. In: *Class. Quant. Grav.* 19.7 (2002), p. 1377. DOI: 10 . 1088 / 0264 - 9381 / 19 / 7 / 321.
- [22] T. M. Niebauer, R. Schilling, K. Danzmann, A. Rüdiger, and W. Winkler. “Nonstationary shot noise and its effect on the sensitivity of interferometers”. In: *Phys. Rev. A* 43 (9 1991), pp. 5022–5029. DOI: 10 . 1103 / PhysRevA . 43 . 5022.
- [23] A. Freise and K. A. Strain. “Interferometer Techniques for Gravitational-Wave Detection”. In: *Living Rev. Relat.* 13.1 (2010). DOI: 10 . 1007 / 1rr-2010-1.

- [24] T. T. Fricke et al. “DC readout experiment in Enhanced LIGO”. In: *Class. Quant. Grav.* 29.6 (2012), p. 065005. DOI: 10.1088/0264-9381/29/6/065005.
- [25] B. T. Gard et al. “Optimal measurement scheme for Gravitational Wave Detection with Coherent and Squeezed Vacuum”. In peer review. 2016.
- [26] *Gravitational wave interferometer noise calculator, v3.1*. 2016.
- [27] H. B. Callen and T. A. Welton. “Irreversibility and Generalized Noise”. In: *Phys. Rev.* 83.1 (1951), pp. 34–40. DOI: 10.1103/PhysRev.83.34.
- [28] G. M. Harry et al. “Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings”. In: *Class. Quant. Grav.* 19.5 (2002), p. 897. DOI: 10.1088/0264-9381/19/5/305.
- [29] G. M. Harry et al. “Titania-doped tantala/silica coatings for gravitational-wave detection”. In: *Class. Quant. Grav.* 24.2 (2007), p. 405. DOI: 10.1088/0264-9381/24/2/008.
- [30] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. “Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae”. In: *Phys. Lett. A* 264.1 (1999), pp. 1–10. ISSN: 0375-9601. DOI: 10.1016/S0375-9601(99)00785-9.
- [31] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. “Thermo-refractive noise in gravitational wave antennae”. In: *Phys. Lett. A* 271.5-6 (2000), pp. 303–307. ISSN: 0375-9601. DOI: 10.1016/S0375-9601(00)00389-3.
- [32] K. Somiya. “Detector configuration of KAGRA: the Japanese cryogenic gravitational-wave detector”. In: *Class. Quant. Grav.* 29.12 (2012), p. 124007. DOI: 10.1088/0264-9381/29/12/124007.
- [33] G. D. Cole, W. Zhang, M. J. Martin, J. Ye, and M. Aspelmeyer. “Tenfold reduction of Brownian noise in high-reflectivity optical coatings”. In: *Nature Photon.* 7.8 (2013), pp. 644–650. ISSN: 1749-4885. DOI: 10.1038/nphoton.2013.174.
- [34] L. Mashev and E. Popov. “Zero order anomaly of dielectric coated gratings”. In: *Opt. Commun.* 55.6 (1985), pp. 377–380. ISSN: 0030-4018. DOI: 10.1016/0030-4018(85)90134-8.
- [35] S. Leavey et al. “Upper limit to the transverse to longitudinal motion coupling of a waveguide mirror”. In: *Class. Quant. Grav.* 32.17 (2015), p. 175005. DOI: 10.1088/0264-9381/32/17/175005.
- [36] B. W. Barr et al. “Silica research in Glasgow”. In: *Class. Quant. Grav.* 19.7 (2002), p. 1655. DOI: 10.1088/0264-9381/19/7/357.

- [37] G. D. Hammond et al. “Reducing the suspension thermal noise of advanced gravitational wave detectors”. In: *Class. Quant. Grav.* 29.12 (2012), p. 124009. DOI: 10 . 1088/0264-9381/29/12/124009.
- [38] N. A. Robertson et al. “Quadruple suspension design for Advanced LIGO”. In: *Class. Quant. Grav.* 19.15 (2002), p. 4043. DOI: 10 . 1088/0264-9381/19/15/311.
- [39] B. Sorazu, K. A. Strain, I. S. Heng, and R. Kumar. “Violin mode amplitude glitch monitor for the presence of excess noise on the monolithic silica suspensions of GEO 600”. In: *Class. Quant. Grav.* 27.15 (2010), p. 155017. DOI: 10 . 1088/0264-9381/27/15/155017.
- [40] S. L. Danilishin and F. Y. Khalili. “Quantum Measurement Theory in Gravitational-Wave Detectors”. In: *Living Rev. Relat.* 15.5 (2012). DOI: 10 . 1007/lrr-2012-5.
- [41] V. B. Braginsky. “Classical and quantum restrictions on the detection of weak disturbances of a macroscopic oscillator.” In: *Ĵ. Exp. Theor. Phys.* 26 (1967), pp. 831–834.
- [42] V. B. Braginsky and F. Y. Khalili. “Quantum nondemolition measurements: the route from toys to tools”. In: *Rev. Mod. Phys.* 68 (1 1996), pp. 1–11. DOI: 10 . 1103/RevModPhys . 68 . 1.
- [43] H. J. Kimble, Y. Levin, A. B. Matsko, K. S. Thorne, and S. P. Vyatchanin. “Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics”. In: *Phys. Rev. D* 65 (2 2001), p. 022002. DOI: 10 . 1103/PhysRevD . 65 . 022002.
- [44] C. M. Caves. “Quantum-mechanical noise in an interferometer”. In: *Phys. Rev. D* 23 (8 1981), pp. 1693–1708. DOI: 10 . 1103/PhysRevD . 23 . 1693.
- [45] R. Adhikari. “Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae”. PhD thesis. MIT, 2004.
- [46] S. M. Aston et al. “Update on quadruple suspension design for Advanced LIGO”. In: *Class. Quant. Grav.* 29.23 (2012), p. 235004. DOI: 10 . 1088/0264-9381/29/23/235004.
- [47] S. Hild et al. “Sensitivity studies for third-generation gravitational wave observatories”. In: *Class. Quant. Grav.* 28.9 (2011), p. 094013. DOI: 10 . 1088 / 0264 - 9381/28/9/094013.
- [48] J. Harms. “Terrestrial Gravity Fluctuations”. In: *Living Rev. Relat.* 18.3 (2015). DOI: 10 . 1007/lrr-2015-3.
- [49] J. Harms and S. Hild. “Passive Newtonian noise suppression for gravitational-wave observatories based on shaping of the local topography”. In: *Class. Quant. Grav.* 31.18 (2014), p. 185011. DOI: 10 . 1088/0264-9381/31/18/185011.

- [50] B. Allen and P. Brady. *Quantization Noise in Ligo Interferometers*. Tech. rep. California Institute of Technology, 1997.
- [51] B. J. Meers. “Recycling in laser-interferometric gravitational-wave detectors”. In: *Phys. Rev. D* 38 (8 1988), pp. 2317–2326. DOI: 10 . 1103 / PhysRevD . 38 . 2317.
- [52] A. Buonanno and Y. Chen. “Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors”. In: *Phys. Rev. D* 64 (4 2001), p. 042006. DOI: 10 . 1103 / PhysRevD . 64 . 042006.
- [53] K. A. Strain and B. J. Meers. “Experimental demonstration of dual recycling for interferometric gravitational-wave detectors”. In: *Phys. Rev. Lett.* 66 (11 1991), pp. 1391–1394. DOI: 10 . 1103 / PhysRevLett . 66 . 1391.
- [54] G. Heinzl et al. “Experimental Demonstration of a Suspended Dual Recycling Interferometer for Gravitational Wave Detection”. In: *Phys. Rev. Lett.* 81 (25 1998), pp. 5493–5496. DOI: 10 . 1103 / PhysRevLett . 81 . 5493.
- [55] A. Freise et al. “Demonstration of detuned dual recycling at the Garching 30 m laser interferometer”. In: *Phys. Lett. A* 277.3 (2000), pp. 135–142. ISSN: 0375-9601. DOI: 10 . 1016 / S0375 - 9601 (00) 00704 - 0.
- [56] G. Heinzl, A. Freise, H. Grote, K. Strain, and K. Danzmann. “Dual recycling for GEO 600”. In: *Class. Quant. Grav.* 19.7 (2002), p. 1547. DOI: 10 . 1088 / 0264 - 9381 / 19 / 7 / 343.
- [57] H. Grote et al. “Dual recycling for GEO 600”. In: *Class. Quant. Grav.* 21.5 (2004), S473. DOI: 10 . 1088 / 0264 - 9381 / 21 / 5 / 013.
- [58] B. Sathyaprakash et al. “Scientific objectives of Einstein Telescope”. In: *Class. Quant. Grav.* 29.12 (2012), p. 124013. DOI: 10 . 1088 / 0264 - 9381 / 29 / 12 / 124013.
- [59] *Single-frequency photonic crystal fiber amplifier with 148-W output power*. Vol. 6453. 2007, 64531J–64531J–8. DOI: 10 . 1117 / 12 . 700268.
- [60] M. Evans et al. “Observation of Parametric Instability in Advanced LIGO”. In: *Phys. Rev. Lett.* 114 (16 2015), p. 161102. DOI: 10 . 1103 / PhysRevLett . 114 . 161102.
- [61] J. Steinlechner et al. “Optical absorption of ion-beam sputtered amorphous silicon coatings”. In: *Phys. Rev. D* 93 (6 2016), p. 062005. DOI: 10 . 1103 / PhysRevD . 93 . 062005.
- [62] V. Braginsky and F. Khalili. *Quantum Measurement*. Ed. by K. Thorne. Cambridge University Press, 1995.

- [63] S. P. Vyatchanin and E. A. Zubova. “Quantum variation measurement of force”. In: *Phys. Lett. A* 201 (1995), pp. 269–274. DOI: 10 . 1016 / 0375 - 9601 (95) 00280 -G.
- [64] S. P. Vyatchanin and A. B. Matsko. “Quantum variation scheme of measurement of force and compensation of back action in interferometric meter of position”. In: *J. Exp. Theor. Phys.* 83.4 (1996), pp. 690–697.
- [65] V. B. Braginsky and F. Y. Khalili. “Gravitational wave antenna with QND speed meter”. In: *Phys. Lett. A* 147.5 (1990), pp. 251–256. ISSN: 0375-9601. DOI: 10 . 1016 / 0375 - 9601 (90) 90442 -Q.
- [66] Y. Chen, S. Danilishin, F. Khalili, and H. Müller-Ebhardt. “QND measurements for future gravitational-wave detectors”. In: *Gen. Rel. Gravit.* 43 (2 2011), pp. 671–694. DOI: 10 . 1007 / s10714 - 010 - 1060 -y.
- [67] B. Willke et al. “The GEO-HF project”. In: *Class. Quant. Grav.* 23.8 (2006), S207. DOI: 10 . 1088 / 0264 - 9381 / 23 / 8 / S26.
- [68] C. Affeldt et al. “Advanced techniques in GEO 600”. In: *Class. Quant. Grav.* 31.22 (2014), p. 224002. DOI: 10 . 1088 / 0264 - 9381 / 31 / 22 / 224002.
- [69] P. Kwee, J. Miller, T. Isogai, L. Barsotti, and M. Evans. “Decoherence and degradation of squeezed states in quantum filter cavities”. In: *Phys. Rev. D* 90 (6 2014), p. 062006. DOI: 10 . 1103 / PhysRevD . 90 . 062006.
- [70] H. Grote et al. “First Long-Term Application of Squeezed States of Light in a Gravitational-Wave Observatory”. In: *Phys. Rev. Lett.* 110 (18 2013), p. 181101. DOI: 10 . 1103 / PhysRevLett . 110 . 181101.
- [71] J. Miller et al. “Prospects for doubling the range of Advanced LIGO”. In: *Phys. Rev. D* 91 (6 2015), p. 062005. DOI: 10 . 1103 / PhysRevD . 91 . 062005.
- [72] A. Buonanno and Y. Chen. “Improving the sensitivity to gravitational-wave sources by modifying the input-output optics of advanced interferometers”. In: *Phys. Rev. D* 69 (10 2004), p. 102004. DOI: 10 . 1103 / PhysRevD . 69 . 102004.
- [73] J. Harms et al. “Squeezed-input, optical-spring, signal-recycled gravitational-wave detectors”. In: *Phys. Rev. D* 68 (4 2003), p. 042001. DOI: 10 . 1103 / PhysRevD . 68 . 042001.
- [74] S. Steinlechner et al. “Local-oscillator noise coupling in balanced homodyne readout for advanced gravitational wave detectors”. In: *Phys. Rev. D* 92 (7 2015), p. 072009. DOI: 10 . 1103 / PhysRevD . 92 . 072009.
- [75] V. B. Braginsky and F. Y. Khalili. “Low noise rigidity in quantum measurements”. In: *Phys. Lett. A* 257 (1999), pp. 241–246. DOI: 10 . 1016 / S0375 - 9601 (99) 00337 -0.

- [76] A. Buonanno and Y. Chen. “Signal recycled laser-interferometer gravitational-wave detectors as optical springs”. In: *Phys. Rev. D* 65 (4 2002), p. 042001. DOI: 10 . 1103/PhysRevD . 65 . 042001.
- [77] T. Corbitt et al. “An All-Optical Trap for a Gram-Scale Mirror”. In: *Phys. Rev. Lett.* 98 (15 2007), p. 150802. DOI: 10 . 1103/PhysRevLett . 98 . 150802.
- [78] H. Rehbein et al. “Double optical spring enhancement for gravitational-wave detectors”. In: *Phys. Rev. D* 78 (6 2008), p. 062003. DOI: 10 . 1103/PhysRevD . 78 . 062003.
- [79] N. Gordon. “Characterisation and control of coupled optical springs for future gravitational wave detectors”. PhD thesis. University of Glasgow, 2015.
- [80] F. Y. Khalili et al. “Negative optical inertia for enhancing the sensitivity of future gravitational-wave detectors”. In: *Phys. Rev. D* 83.6 (2011), p. 062003. DOI: 10 . 1103/PhysRevD . 83 . 062003.
- [81] N. V. Voronchev, S. L. Danilishin, and F. Y. Khalili. “Negative optical inertia in optomechanical systems”. In: *Opt. Spectrosc.* 112.3 (2012), pp. 377–385. ISSN: 1562-6911. DOI: 10 . 1134/S0030400X12030216.
- [82] V. B. Braginsky, M. L. Gorodetsky, and F. Y. Khalili. “Optical bars in gravitational wave antennas”. In: *Phys. Lett. A* 232 (1997), pp. 340–348. DOI: 10 . 1016/S0375-9601(97)00413-1.
- [83] F. Y. Khalili. “The “optical lever” intracavity readout scheme for gravitational-wave antennae”. In: *Phys. Lett. A* 298 (2002), pp. 308–314. DOI: 10 . 1016/S0375-9601(02)00578-9.
- [84] S. L. Danilishin and F. Y. Khalili. “Practical design of the optical lever intracavity topology of gravitational-wave detectors”. In: *Phys. Rev. D* 73.2 (2006), p. 022002. DOI: 10 . 1103/PhysRevD . 73 . 022002.
- [85] V. B. Braginsky, M. L. Gorodetsky, F. Y. Khalili, and K. S. Thorne. “Dual-resonator speed meter for a free test mass”. In: *Phys. Rev. D* 61 (2000), p. 044002. DOI: 10 . 1103/PhysRevD . 61 . 044002.
- [86] P. Purdue and Y. Chen. “Practical speed meter designs for quantum nondemolition gravitational-wave interferometers”. In: *Phys. Rev. D* 66 (12 2002), p. 122004. DOI: 10 . 1103/PhysRevD . 66 . 122004.
- [87] Y. Chen. “Sagnac interferometer as a speed-meter-type, quantum-nondemolition gravitational-wave detector”. In: *Phys. Rev. D* 67 (12 2003), p. 122004. DOI: 10 . 1103/PhysRevD . 67 . 122004.

- [88] S. L. Danilishin. “Sensitivity limitations in optical speed meter topology of gravitational-wave antennas”. In: *Phys. Rev. D* 69 (10 2004), p. 102003. DOI: 10 . 1103/PhysRevD . 69 . 102003.
- [89] S. L. Danilishin et al. “Quantum noise of non-ideal Sagnac speed meter interferometer with asymmetries”. In: *New J. Phys.* 17.4 (2015), p. 043031. DOI: 10 . 1088 / 1367 - 2630 / 17 / 4 / 043031.
- [90] M. Wang et al. “Realistic polarizing Sagnac topology with DC readout for the Einstein Telescope”. In: *Phys. Rev. D* 87 (9 2013), p. 096008. DOI: 10 . 1103/PhysRevD . 87 . 096008.
- [91] S. Huttner et al. “Candidates for a possible third-generation gravitational wave detector: comparison of ring-Sagnac and sloshing-Sagnac speedmeter interferometers”. Accepted for publication. 2016.
- [92] D. Pascucci et al. “Quantum noise limitations due to coherent scattering in Sagnac speed meter with triangular cavities”. In preparation. 2016.
- [93] Y. Levin. “Internal thermal noise in the LIGO test masses: A direct approach”. In: *Phys. Rev. D* 57 (1998), pp. 659–663. DOI: 10 . 1103/PhysRevD . 57 . 659.
- [94] N. Nakagawa, A. M. Gretarsson, E. K. Gustafson, and M. M. Fejer. “Thermal noise in half-infinite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror”. In: *Phys. Rev. D* 65 (10 2002), p. 102001. DOI: 10 . 1103/PhysRevD . 65 . 102001.
- [95] D. R. M. Crooks et al. “Excess mechanical loss associated with dielectric mirror coatings on test masses in interferometric gravitational wave detectors”. In: *Class. Quant. Grav.* 19.5 (2002), p. 883. DOI: 10 . 1088/0264 - 9381/19/5/304.
- [96] R. Flaminio et al. “A study of coating mechanical and optical losses in view of reducing mirror thermal noise in gravitational wave detectors”. In: *Class. Quant. Grav.* 27.8 (2010), p. 084030. DOI: 10 . 1088/0264 - 9381/27/8/084030.
- [97] R. Bassiri et al. “Correlations between the mechanical loss and atomic structure of amorphous TiO₂-doped Ta₂O₅ coatings”. In: *Acta Mater.* 61.4 (2013), pp. 1070–1077. ISSN: 1359-6454. DOI: 10 . 1016/j . actamat . 2012 . 10 . 009.
- [98] R. Dannenberg. *Advanced LIGO End Test Mass Coating Specification*. Tech. rep. E0900068-v3. LIGO, 2009.
- [99] Y. Sakakibara et al. “Progress on the cryogenic system for the KAGRA cryogenic interferometric gravitational wave telescope”. In: *Class. Quant. Grav.* 31.22 (2014), p. 224003. DOI: 10 . 1088/0264 - 9381/31/22/224003.

- [100] T. Uchiyama et al. “Reduction of Thermal Fluctuations in a Cryogenic Laser Interferometric Gravitational Wave Detector”. In: *Phys. Rev. Lett.* 108.14 (2012), pp. 141101–141105. DOI: 10.1103/PhysRevLett.108.141101.
- [101] M. Punturo et al. “The third generation of gravitational wave observatories and their science reach”. In: *Class. Quant. Grav.* 27.8 (2010), p. 084007. DOI: 10.1088/0264-9381/27/8/084007.
- [102] I. W. Martin et al. “Effect of heat treatment on mechanical dissipation in Ta2O5 coatings”. In: *Class. Quant. Grav.* 27.22 (2010), p. 225020. DOI: 10.1088/0264-9381/27/22/225020.
- [103] M. R. Abernathy et al. “Cryogenic mechanical loss measurements of heat-treated hafnium dioxide”. In: *Class. Quant. Grav.* 28.19 (2011), p. 195017. DOI: 10.1088/0264-9381/28/19/195017.
- [104] The LIGO Scientific Collaboration. *Instrument Science White Paper*. Tech. rep. T1500290-v2. LIGO, 2015.
- [105] M. Granata et al. “Cryogenic measurements of mechanical loss of high-reflectivity coating and estimation of thermal noise”. In: *Opt. Lett.* 38.24 (2013), pp. 5268–5271. DOI: 10.1364/OL.38.005268.
- [106] B. Mours, E. Tournefier, and J.-Y. Vinet. “Thermal noise reduction in interferometric gravitational wave antennas: using high order TEM modes”. In: *Class. Quant. Grav.* 23.20 (2006), p. 5777. DOI: 10.1088/0264-9381/23/20/001.
- [107] E. D’Ambrosio et al. “Advanced LIGO: non-Gaussian beams”. In: *Class. Quant. Grav.* 21.5 (2004), S867. DOI: 10.1088/0264-9381/21/5/074.
- [108] M. Bondarescu and K. Thorne. “New family of light beams and mirror shapes for future LIGO interferometers”. In: *Phys. Rev. D* 74.74 (8 2006), p. 082003. DOI: 10.1103/PhysRevD.74.082003.
- [109] B. Sorazu et al. “Experimental test of higher-order Laguerre-Gauss modes in the 10 m Glasgow prototype interferometer”. In: *Class. Quant. Grav.* 30.3 (2013), p. 035004. DOI: 10.1088/0264-9381/30/3/035004.
- [110] F. Brückner et al. “Monolithic dielectric surfaces as new low-loss light-matter interfaces”. In: *Opt. Lett.* 33.3 (2008), pp. 264–266. DOI: 10.1364/OL.33.000264.
- [111] F. Brückner et al. “Encapsulated subwavelength grating as a quasi-monolithic resonant reflector”. In: *Opt. Express* 17.26 (2009), pp. 24334–24341. DOI: 10.1364/OE.17.024334.
- [112] F. Brückner et al. “Realization of a Monolithic High-Reflectivity Cavity Mirror from a Single Silicon Crystal”. In: *Phys. Rev. Lett.* 104 (16 2010), p. 163903. DOI: 10.1103/PhysRevLett.104.163903.

- [113] D. Friedrich et al. “Waveguide grating mirror in a fully suspended 10 meter Fabry-Perot cavity”. In: *Opt. Express* 19.16 (2011), pp. 14955–14963. DOI: 10.1364/OE.19.014955.
- [114] A. Bunkowski, O. Burmeister, D. Friedrich, K. Danzmann, and R. Schnabel. “High reflectivity grating waveguide coatings for 1064 nm”. In: *Class. Quant. Grav.* 23.24 (2006), p. 7297. DOI: 10.1088/0264-9381/23/24/007.
- [115] A. Sharon, D. Rosenblatt, and A. A. Friesem. “Resonant grating-waveguide structures for visible and near-infrared radiation”. In: *J. Opt. Soc. Am.* 14.11 (1997), pp. 2985–2993. DOI: 10.1364/JOSAA.14.002985.
- [116] D. Brown et al. “Invariance of waveguide grating mirrors to lateral displacement phase shifts”. In: *Opt. Lett.* 38.11 (2013), pp. 1844–1846. DOI: 10.1364/OL.38.001844.
- [117] D. Heinert et al. “Calculation of thermal noise in grating reflectors”. In: *Phys. Rev. D* 88 (4 2013), p. 042001. DOI: 10.1103/PhysRevD.88.042001.
- [118] S. Wise et al. “Phase Effects in the Diffraction of Light: Beyond the Grating Equation”. In: *Phys. Rev. Lett.* 95 (1 2005), p. 013901. DOI: 10.1103/PhysRevLett.95.013901.
- [119] A. Freise, A. Bunkowski, and R. Schnabel. “Phase and alignment noise in grating interferometers”. In: *New J. Phys.* 9.12 (2007), p. 433. DOI: 10.1088/1367-2630/9/12/433.
- [120] B. W. Barr et al. “Translational, rotational, and vibrational coupling into phase in diffractively coupled optical cavities”. In: *Opt. Lett.* 36.14 (2011), pp. 2746–2748. DOI: 10.1364/OL.36.002746.
- [121] R. W. P. Drever et al. “Laser phase and frequency stabilization using an optical resonator”. English. In: *Appl. Phys. B* 31.2 (1983), pp. 97–105. ISSN: 0946-2171. DOI: 10.1007/BF00702605.
- [122] M. V. Plissi et al. “GEO 600 triple pendulum suspension system: Seismic isolation and control”. In: *Rev. Sci. Instrum.* 71.6 (2000), pp. 2539–2545. DOI: 10.1063/1.1150645.
- [123] A. White. “Frequency stabilization of gas lasers”. In: *IEEE J. Quant. Electron.* 1.8 (1965), pp. 349–357. ISSN: 0018-9197. DOI: 10.1109/JQE.1965.1072246.
- [124] R. L. Ward et al. “dc readout experiment at the Caltech 40m prototype interferometer”. In: *Class. Quant. Grav.* 25.11 (2008), p. 114030. DOI: 10.1088/0264-9381/25/11/114030.

- [125] D. A. Shaddock, M. B. Gray, and D. E. McClelland. “Frequency locking a laser to an optical cavity by use of spatial mode interference”. In: *Opt. Lett.* 24.21 (1999), pp. 1499–1501. DOI: 10.1364/OL.24.001499.
- [126] E. D. Black. “An introduction to Pound-Drever-Hall laser frequency stabilization”. In: *Am. J. Phys.* 69.1 (2001), pp. 79–87. DOI: 10.1119/1.1286663.
- [127] J. Macarthur. “Towards surpassing the standard quantum limit using optical springs”. PhD thesis. University of Glasgow, 2014.
- [128] W. K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1 (1970), pp. 97–109. DOI: 10.1093/biomet/57.1.97.
- [129] C. M. Caves and B. L. Schumaker. “New formalism for two-photon quantum optics. I. Quadrature phases and squeezed states”. In: *Phys. Rev. A* 31 (5 1985), pp. 3068–3092. DOI: 10.1103/PhysRevA.31.3068.
- [130] B. L. Schumaker and C. M. Caves. “New formalism for two-photon quantum optics. II. Mathematical foundation and compact notation”. In: *Phys. Rev. A* 31 (5 1985), pp. 3093–3111. DOI: 10.1103/PhysRevA.31.3093.
- [131] A. R. Wade et al. “Polarization speed meter for gravitational-wave detection”. In: *Phys. Rev. D* 86 (6 2012), p. 062001. DOI: 10.1103/PhysRevD.86.062001.
- [132] H. Müller-Ebhardt et al. “Quantum-state preparation and macroscopic entanglement in gravitational-wave detectors”. In: *Phys. Rev. A* 80.4 (2009), p. 043802. DOI: 10.1103/PhysRevA.80.043802.
- [133] N. V. Voronchev, S. P. Tarabrin, and S. L. Danilishin. “Broadband detuned Sagnac interferometer for future generation gravitational wave astronomy”. In preparation. 2015.
- [134] D. A. Shaddock, M. B. Gray, and D. E. McClelland. “Experimental demonstration of resonant sideband extraction in a Sagnac interferometer”. In: *Appl. Opt.* 37.34 (1998), pp. 7995–8001. DOI: 10.1364/AO.37.007995.
- [135] P. T. Beyersdorf, R. L. Byer, and M. M. Fejer. “Results from the Stanford 10 m Sagnac interferometer”. In: *Class. Quant. Grav.* 19.7 (2002), p. 1585. DOI: 10.1088/0264-9381/19/7/348.
- [136] T. Eberle et al. “Quantum Enhancement of the Zero-Area Sagnac Interferometer Topology for Gravitational Wave Detection”. In: *Phys. Rev. Lett.* 104 (25 2010), p. 251102. DOI: 10.1103/PhysRevLett.104.251102.
- [137] C. Gräf et al. “Design of a speed meter interferometer proof-of-principle experiment”. In: *Class. Quant. Grav.* 31.21 (2014), p. 215009. DOI: 10.1088/0264-9381/31/21/215009.

- [138] S. Leavey et al. “Control of a velocity-sensitive audio-band quantum non-demolition interferometer”. In preparation. 2016.
- [139] R. Weiss. *Collection of Reports on Barkhausen Noise*. Tech. rep. LIGO, 2008.
- [140] M. Hewitson et al. “Charge measurement and mitigation for the main test masses of the GEO 600 gravitational wave observatory”. In: *Class. Quant. Grav.* 24.24 (2007), p. 6379. DOI: 10.1088/0264-9381/24/24/013.
- [141] A. Gläfke. “Measurement beyond the Quantum Limit: Sensing and Control of a Sagnac Speed Meter”. MA thesis. Leibniz Universität Hannover, 2015.
- [142] R. Bork. *Advanced LIGO CDS Design Overview*. Tech. rep. LIGO Laboratory, 2010.
- [143] G. Vajente. *Advanced Virgo Steady State Length Sensing and Control Simulation*. Tech. rep. VIR-0449D-11. Virgo, 2011.
- [144] P. Gray, P. Hurst, S. Lewis, and R. Meyer. *Analysis and Design of Analog Integrated Circuits*. 5th ed. Analysis and Design of Analog Integrated Circuits. John Wiley & Sons, 2009. ISBN: 9780470245996.
- [145] F. Seifert. *Resistor Current Noise Measurements*. Tech. rep. Max-Planck-Institut für Gravitationsphysik, 2009.
- [146] P. Horowitz and W. Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015. ISBN: 9780521809269.
- [147] B. N. Shapiro. “Adaptive modal damping for advanced LIGO suspensions”. PhD thesis. Massachusetts Institute of Technology, 2012.
- [148] A. Freise. “The Next Generation of Interferometry: Multi-Frequency Optical Modelling, Control Concepts and Implementation”. PhD thesis. Universität Hannover, 2003.
- [149] J. R. Smith et al. “Linear projection of technical noise for interferometric gravitational-wave detectors”. In: *Class. Quant. Grav.* 23.2 (2006), p. 527. DOI: 10.1088/0264-9381/23/2/016.
- [150] H. Müller-Ebhardt et al. *Review of quantum non-demolition schemes for the Einstein Telescope*. Tech. rep. ET-010-09. Max-Planck-Institut für Gravitationsphysik, 2009.
- [151] H. Miao et al. “Probing macroscopic quantum states with a sub-Heisenberg accuracy”. In: *Phys. Rev. A* 81.1 (2010), p. 012114. DOI: 10.1103/PhysRevA.81.012114.
- [152] R. Dolesi et al. “Brownian force noise from molecular collisions and the sensitivity of advanced gravitational wave observatories”. In: *Phys. Rev. D* 84(6 2011), p. 063007. DOI: 10.1103/PhysRevD.84.063007.

- [153] S Margulies. “Force on a dielectric slab inserted into a parallel-plate capacitor”. In: *Am. J. Phys.* 52.6 (1984), pp. 515–518.
- [154] H. Wittel, S. Hild, G. Bergmann, K. Danzmann, and K. A. Strain. “New design of electrostatic mirror actuators for application in high-precision interferometry”. In: *Class. Quant. Grav.* 32.17 (2015), p. 175021. DOI: 10 . 1088 / 0264 - 9381 / 32 / 17 / 175021.
- [155] EPL/66. *BS EN 61010-1:2010: Safety requirements for electrical equipment for measurement, control, and laboratory use*. Tech. rep. European Commission, 2010.
- [156] D. A. Simakov. “Time-domain analysis of a dynamically tuned signal recycled interferometer for the detection of chirp gravitational waves from coalescing compact binaries”. In: *Phys. Rev. D* 90 (10 2014), p. 102003. DOI: 10 . 1103 / PhysRevD . 90 . 102003.
- [157] K. Somiya and K. Yamamoto. “Coating thermal noise of a finite-size cylindrical mirror”. In: *Phys. Rev. D* 79 (10 2009), p. 102004. DOI: 10 . 1103 / PhysRevD . 79 . 102004.
- [158] M. Hewitson and P. Ajith. “Using the null-stream of GEO 600 to veto transient events in the detector output”. In: *Class. Quant. Grav.* 22.22 (2005), p. 4903. DOI: 10 . 1088 / 0264 - 9381 / 22 / 22 / 015.
- [159] P. Ajith, M. Hewitson, and I. S. Heng. “Null-stream veto for two co-located detectors: implementation issues”. In: *Class. Quant. Grav.* 23.19 (2006), S741. DOI: 10 . 1088 / 0264 - 9381 / 23 / 19 / S13.
- [160] *Plans for a large gravitational wave antenna in Germany*. 1985.
- [161] S. Hild et al. “A xylophone configuration for a third-generation gravitational wave detector”. In: *Class. Quant. Grav.* 27.1 (2010), p. 015003. DOI: 10 . 1088 / 0264 - 9381 / 27 / 1 / 015003.
- [162] G. Conforto and R. DeSalvo. “Proposal for lower frequency companions for the advanced LIGO Gravitational Wave Interferometric Detectors”. In: *Nucl. Instr. Meth. Phys. Res. A* 518.1-2 (2004). Frontier Detectors for Frontier Physics: Proceedin, pp. 228–232. ISSN: 0168-9002. DOI: 10 . 1016 / j . nima . 2003 . 10 . 068.
- [163] F. Acernese et al. “Measurements of Superattenuator seismic isolation by Virgo interferometer”. In: *Astropart. Phys.* 33.3 (2010), pp. 182 –189. ISSN: 0927-6505. DOI: 10 . 1016 / j . astropartphys . 2010 . 01 . 006.
- [164] L. Carbone, C. Bogan, P. Fulda, A. Freise, and B. Willke. “Generation of High-Purity Higher-Order Laguerre-Gauss Beams at High Laser Power”. In: *Phys. Rev. Lett.* 110 (25 2013), p. 251101. DOI: 10 . 1103 / PhysRevLett . 110 . 251101.

- [165] S. Hild and The LIGO Scientific Collaboration. “The status of GEO 600”. In: *Class. Quant. Grav.* 23.19 (2006), S643. DOI: 10 . 1088/0264-9381/23/19/S02.
- [166] K. Somiya et al. “Development of a frequency-detuned interferometer as a prototype experiment for next-generation gravitational-wave detectors”. In: *Appl. Opt.* 44.16 (2005), pp. 3179–3191. DOI: 10 . 1364/AO . 44 . 003179.
- [167] J. E. Mason and P. A. Willems. “Signal extraction and optical design for an advanced gravitational-wave interferometer”. In: *Appl. Opt.* 42.7 (2003), pp. 1269–1282. DOI: 10 . 1364/AO . 42 . 001269.
- [168] S. Hild et al. “Demonstration and comparison of tuned and detuned signal recycling in a large-scale gravitational wave detector”. In: *Class. Quant. Grav.* 24.6 (2007), p. 1513. DOI: 10 . 1088/0264-9381/24/6/009.
- [169] H. Grote. “High power, low-noise, and multiply resonant photodetector for interferometric gravitational wave detectors”. In: *Rev. Sci. Instrum.* 78.5, 054704 (2007). DOI: 10 . 1063/1 . 2735559.
- [170] A. Heidmann, W. Chaibi, and F. Bondu. “Nano- and micro-optomechanical systems Optomechanical issues in the gravitational wave detector Advanced VIRGO”. In: *Compt. Rend. Phys.* 12.9 (2011), pp. 888 –897. ISSN: 1631-0705. DOI: 10 . 1016/j . crhy . 2011 . 07 . 004.
- [171] N. Vostrosablin and S. P. Vyatchanin. “Stable optical spring in the Advanced LIGO detector with unbalanced arms and in the Michelson-Sagnac interferometer”. In: *Phys. Rev. D* 89 (6 2014), p. 062005. DOI: 10 . 1103/PhysRevD . 89 . 062005.
- [172] G. Heinzl. “Advanced optical techniques for laser-interferometric gravitational-wave detectors”. PhD thesis. Max-Planck-Institut für Quantenoptik, 1999.
- [173] R. Abbott et al. *Advanced LIGO Length Sensing and Control Final Design*. Tech. rep. LIGO Laboratory, 2010.
- [174] P. Fritschel et al. “Readout and control of a power-recycled interferometric gravitational-wave antenna”. In: *Appl. Opt.* 40.28 (2001), pp. 4988–4998. DOI: 10 . 1364/AO . 40 . 004988.
- [175] K. A. Strain et al. “Sensing and control in dual-recycling laser interferometer gravitational-wave detectors”. In: *Appl. Opt.* 42.7 (2003), pp. 1244–1256. DOI: 10 . 1364/AO . 42 . 001244.
- [176] B. W. Barr et al. “Control sideband generation for dual-recycled laser interferometric gravitational wave detectors”. In: *Class. Quant. Grav.* 23.18 (2006), p. 5661. DOI: 10 . 1088/0264-9381/23/18/010.

- [177] A. Staley et al. “Achieving resonance in the Advanced LIGO gravitational-wave interferometer”. In: *Class. Quant. Grav.* 31.24 (2014), p. 245010. DOI: 10 . 1088 / 0264-9381/31/24/245010.
- [178] G. Vajente. *Note on signal recycling II: Lengths and modulations*. Tech. rep. VIR-032A-08. Virgo Collaboration, 2008.
- [179] Y. Aso et al. “Interferometer design of the KAGRA gravitational wave detector”. In: *Phys. Rev. D* 88 (4 2013), p. 043007. DOI: 10 . 1103/PhysRevD . 88 . 043007.
- [180] A. Effler. “Performance Characterization of the Dual-Recycled Michelson Subsystem in Advanced LIGO”. PhD thesis. Louisiana State University, 2014.
- [181] F. Acernese et al. “Lock acquisition of the Virgo gravitational wave detector”. In: *Astropart. Phys.* 30.1 (2008), pp. 29 –38. ISSN: 0927-6505. DOI: 10 . 1016 / j . astropartphys . 2008 . 06 . 005.
- [182] M. Mantovani and A. Freise. “Evaluating mirror alignment systems using the optical sensing matrix”. In: *JPCS* 122.1 (2008), p. 012026.
- [183] A. J. Mullavey et al. “Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers”. In: *Opt. Express* 20.1 (2012), pp. 81–89. DOI: 10 . 1364/OE . 20 . 000081.
- [184] F. Acernese et al. “A local control system for the test masses of the Virgo gravitational wave detector”. In: *Astropart. Phys.* 20.6 (2004), pp. 617 –628. ISSN: 0927-6505. DOI: 10 . 1016/j . astropartphys . 2003 . 10 . 001.
- [185] A. Stochino et al. “The Seismic Attenuation System (SAS) for the Advanced LIGO gravitational wave interferometric detectors”. In: *Nucl. Instr. Meth. Phys. Res. A* 598.3 (2009), pp. 737 –753. ISSN: 0168-9002. DOI: 10 . 1016/j . nima . 2008 . 10 . 023.
- [186] S. Goßler et al. “The AEI 10 m prototype interferometer”. In: *Class. Quant. Grav.* 27.8 (2010), p. 084023. DOI: 10 . 1088/0264-9381/27/8/084023.
- [187] K. Dahl et al. “Towards a Suspension Platform Interferometer for the AEI 10 m Prototype Interferometer”. In: *JPCS* 228.1 (2010), p. 012027.
- [188] Y. B. Band. *Light and Matter: Electromagnetism, Optics, Spectroscopy and Lasers*. Light and Matter. John Wiley & Sons, 2006. ISBN: 9780471899310.
- [189] P. Welch. “The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms”. In: *IEEE Trans. Aud. Elec.* 15.2 (1967), pp. 70–73. ISSN: 0018-9278. DOI: 10 . 1109 / TAU . 1967 . 1161901.

- [190] A. Freise et al. “Frequency-domain interferometer simulation with higher-order spatial modes”. In: *Class. Quant. Grav.* 21.5 (2004), S1067. DOI: 10 . 1088 / 0264 - 9381 / 21 / 5 / 102.
- [191] C. L. Mueller et al. “In situ characterization of the thermal state of resonant optical interferometers via tracking of their higher-order mode resonances”. In: *Class. Quant. Grav.* 32.13 (2015), p. 135018. DOI: 10 . 1088 / 0264 - 9381 / 32 / 13 / 135018.
- [192] A. Kumeta, C. Bond, and K. Somiya. “Design study of the KAGRA output mode cleaner”. In: *Opt. Rev.* 22.1 (2015), pp. 149–152. ISSN: 1349-9432. DOI: 10 . 1007 / s10043-015-0028-2.
- [193] G. Müller, T. Delker, D. B. Tanner, and D. Reitze. “Dual-recycled cavity-enhanced Michelson interferometer for gravitational-wave detection”. In: *Appl. Opt.* 42.7 (2003), pp. 1257–1268. DOI: 10 . 1364 / AO . 42 . 001257.
- [194] T. Corbitt, Y. Chen, and N. Mavalvala. “Mathematical framework for simulation of quantum fields in complex interferometers using the two-photon formalism”. In: *Phys. Rev. A* 72 (1 2005), p. 013818. DOI: 10 . 1103 / PhysRevA . 72 . 013818.
- [195] M. Evans, L. Barsotti, P. Kwee, J. Harms, and H. Miao. “Realistic filter cavities for advanced gravitational wave detectors”. In: *Phys. Rev. D* 88 (2 2013), p. 022002. DOI: 10 . 1103 / PhysRevD . 88 . 022002.
- [196] D. Brown. “Interactions of light and mirrors: advanced techniques for modelling future gravitational wave detectors”. PhD thesis. University of Birmingham, 2016.
- [197] doi: 10 . 5525 / gla . researchdata . 202.