

**School of Physics
and Astronomy**



Enhancing the QUEST experiment with a coin sized
Output Mode Cleaner for improved sensitivity to
Quantum Gravity signatures

William L. Griffiths

Submitted for the degree of Doctor of Philosophy
School of Physics and Astronomy
Cardiff University

Submission: 24/04/2023
Viva: 05/06/2023

For the imposter.

Acknowledgments

I am grateful for support from the Science and Technology Facilities Council (STFC), grants ST/T006331/1, ST/I006285/1, and ST/L000946/1, the Leverhulme Trust, grant RPG-2019-022, Kate, Hartmut and Cardiff University for making my PhD possible; thank you. Additional thank you to STFC for supporting my 4 months in LIGO; the COLA and travel funding for the long term attachment were invaluable. Thank you too to the LIGO fellowship program for the accommodation, car hire, travel arrangements and financial assistance, it would not have been possible otherwise.

A final funding related thank you to the Dr Benjamin Glyndwr Owen Scholarship, who awarded me a prize in my final year for being “A PGR student who has made an outstanding contribution to postgraduate research”. It was thoroughly unexpected, and a wonderful way to help with final year celebrations...I mean work.

Lorenzo, we met on day 1 of my PhD and you have guided my every step from that moment to this one. It is no exaggeration to say I absolutely would not be submitting this piece of work with your help and support. Your mentorship and supervision (yes, every single one of those 14 percents), and equally as importantly — your friendship — have shown me what it means to go above and beyond, setting the example of the kind of teammate I can only one day hope to be. Any successes I enjoy from this are indebted to you, and it has been the privilege of a lifetime to have worked with you.

Thank you to all the fellow students and academics I met along the way. It certainly helps to be surrounded by passionate people, and people experiencing similar joys...and tribulations. The GW lab team members were of course especially involved in my PhD. Lorenzo, Aldo, Eyal, the lab wouldn't have run without you and my progress, learning and work would have been near impossible without all of your support; I honestly cannot thank you enough. Ali, Sander, and Abhinav, thank you too for your help. I hope the experiment continues (Abhinav, that's on you now buddy, best of luck and no pressure) and we all revel in the fruits of our efforts.

Jordan and Terri, you had the misfortune of being stuck in an office with me, for which, I am very glad. You ended up organising my viva celebrations so I can't have annoyed you too much. It was great getting to know both of you and I can't wish you well enough.

Thank you to Adam Mullavey for all your help at LIGO, and the continued help with calls, emails and proof reading afterwards; I hugely appreciate the effort.

Thank you to Keiko and Ed for being my examiners and helping me put the finishing

touches on the thesis. And thank you to Emyr and Matt for your support in the final weeks.

My parents, brother and sister — I'm sure you've seen a change in me since childhood, my time in the military, a seemingly random choice to do a degree and the even more left field ambition to continue to this point. I certainly didn't see it all coming, I doubt any of you did. But I have felt comfortable pursuing these achievements knowing that you'd all support me. Weekends at the farm were a welcome break, especially when my nieces and nephew made an appearance. Thank you all, forever.

To my extended family, aunties, uncles and cousins — we don't see enough of each other but when we do your support is always warmly felt. I can always rely on you for an ego boost and that has been invaluable at times of heavy bleakness.

For those we've lost, I hope to have made you proud. To my grandad who also studied physics at Cardiff and then taught it at high school for his career; I wish we could have spoken about all of this, I'm sure you'd have been baffled by the discoveries since you studied.

Matt and Phil, I thought this thesis was tough, but keeping what I write about you two appropriate for a published piece of work is infinitely more difficult. Thank you for your support and friendship, you've both become family and I wouldn't be without you. That...that's all I've got...

...for this edition...

Eleri, my fiancé, who met me as I began year 2 of my BSc. You have known me through the most stressful periods of my life and have endured it (me) with grace, understanding, patience and love. Even though I had to accept being second to your horse, I hope you never felt second to this. Regardless of how large a feat I may consider this to be, convincing you to say yes, will forever be my greatest achievement. I struggle to find words of appropriate strength to thank you enough for being my motivation, my inspiration, my girl. For now I'll just say thank you and I'll work on showing you how grateful I am for you in the years to come, with you by my side.

A final thank you to someone who is now a memory, and from whose sowing I now reap. From beginning to end, awake, asleep, happy, sad, enthusiastic, pessimistic, ambitious and lethargic; there isn't a second you weren't there, not a moment you missed. There wasn't a sigh you didn't shrug off or a buried head you didn't support. The past is the sculptor of the future; this may not be Michelangelo's David, but you chipped away until you discovered something good enough. The imposter will likely never quieten into irrelevance but perhaps his grip will weaken and his continuous banshee cries will dissolve to an infrequent murmur. Well done, don't forget to be kind to yourself and to enjoy what follows. To the very best of times, cheers.

Summary of thesis

This PhD thesis details the work undertaken at Cardiff University on the quantum enhanced spacetime (QUEST) experiment and at laser interferometer gravitational-wave observatory (LIGO) Livingston, and it is structured as follows. Chapter one gives an overview of the theory about the three major science goals for the QUEST experiment: quantum gravity, high frequency gravitational waves and scalar field dark matter. Chapter two is a review of the foundational scientific principles and experimental methods used throughout the thesis. Chapter three is dedicated to describing the QUEST experiment, from the design to the commissioning up to the last results achieved at the time of writing. Chapter four describes the output mode cleaner (OMC), an optical cavity in the detection path of QUEST which facilitates the use of high power and is essential for QUEST's unprecedented sensitivity level goal. It was designed from first principles and has been brought to completion with a working control loop and comprehensive characterisation. Chapter five is related to a four month fellowship at LIGO Livingston. The project goal was to develop a new method of acquiring the science laser vs X-arm length error signal by cancelling the frequency noise introduced by an optical fiber. New fiber optic breadboards were designed and installed in the per-stabilised laser enclosure and the x-end station during the fellowship and the results do show promise. I am the joint first PhD student to graduate on QUEST. The other, Sander Vermeulen, worked on the data acquisition for QUEST. The Covid-19 lockdowns somewhat impacted my work. As an experimentalist, losing access to the lab for several weeks certainly delayed progress. Nonetheless, projects were sufficiently concluded to graduate.

Statement of work

The following lists specify my contributions to the research described in this thesis. Unless stated, each activity was either conducted entirely/mostly independently.

QUEST general

- Assisted with the installation of the optical benches
- Wrote the laser operation instruction manual
- Wrote the interlock instruction manual
- Decided on, bought and installed the laminar flow bench
- Assisted with deep cleaning of the lab a few times

QUEST experimental work

- Optical component characterisation: Faraday isolators, mirrors etc (alongside Dr Lorenzo Aiello)
- Laser characterisation: Beam profiling, power stabilisation, relative intensity noise assessment (alongside Dr Lorenzo Aiello)
- Prototype stages: Fabry-Perot cavity and Michelson interferometer (alongside Dr Lorenzo Aiello)
- Glued mirrors to the PZTs (alongside Dr Lorenzo Aiello)
- Soldered PZT cables to BNC connectors (alongside Dr Lorenzo Aiello)
- Redesigned a component of the “Little Input Mode Cleaner” from Caltech
- Built the “Little Input Mode Cleaner” from Caltech
- Developed Gaussian beam modelling Matlab script with focus on mode matching telescope solution generation
- Assisted in lock scheme design
- Assisted in lock commissioning
- Attempted correction of longitudinal to angular coupling
- Designed the mode matching telescope - confirming results against an independent method used by Dr Lorenzo Aiello
- Assisted with optical positioning and tuning

- Design of the CDS loop control overview screen

QUEST beam dumps

- Optical design
- Base designed, full 3D model
- Chose the glass
- Cut the glass
- Cleaned the glass (with help from Terri Pearce)
- Beam dump construction
- Attempted characterisation — limited by power meter sensitivity (no power was detected anywhere around the dumps)

Output Mode Cleaner

- Entire design from first principles, optics, actuator, geometry, everything
- Lock loop design
- Lock loop construction (in CDS)
- Lock loop commissioning (with help from Dr Eyal Schwartz, Dr Lorenzo Aiello and Dr Aldo Ejlli)
- OMC characterisation
- Breadboard design
- Breadboard construction
- Mode matching telescope(s) design, placement and tuning
- Housing design, full 3D model(s)
- Gluing of mirrors
- Soldering actuator wires
- Design of the CDS loop control overview screen

LIGO

- Design of fibre breadboard (with Dr Adam Mullavey)
- Construction of fibre breadboard
- Testing of fibre breadboard
- Characterisation of PDs
- Installation of breadboards in end chamber and main laser enclosure (with Dr Adam Mullavey)

Contents

Introduction	2
1 Fundamental physics with table-top interferometers	6
1.1 Quantum Gravity	6
1.1.1 Light cones and causal diamonds	7
1.1.2 Horizons	9
1.1.3 Reality check	12
1.1.4 Measuring quantum gravity effects at horizons and length fluctuations	12
1.1.5 Signal confirmation	13
1.2 Gravitational Waves	13
1.2.1 Wave solution to Einstein's equations in the weak field limit	14
1.2.2 Towards gravitational wave detection	15
1.2.3 Length fluctuations — Strain	17
1.2.4 Gravitational wave frequency	17
1.3 Dark Matter	18
1.3.1 Scalar fields	18
1.3.2 Fluctuations in fundamental constants	19
1.3.3 Length fluctuations	19
1.3.4 Signal searches	21
1.4 The Michelson Interferometer	22
1.4.1 Noise	24
1.4.2 Shot-noise	25
1.4.3 Radiation pressure	26
1.4.4 Standard quantum limit	26
1.4.5 Co-located Michelson interferometers	27
1.5 The quantum enhanced spacetime experiment	28
2 From theory to experiment	30
2.1 Field propagation of the Michelson interferometer	30
2.1.1 Signal sidebands by phase modulation	34
2.2 Gaussian laser beam properties	35
2.3 Optical cavities	39
2.3.1 Electric fields around a cavity	39
2.3.2 Cavity geometry	40
2.3.3 Resonance and Finesse	41
2.3.4 Frequency reference	44
2.3.5 Spatial mode filtering	44

2.3.6	Low-pass filter	45
2.3.7	Higher order mode generation	46
2.4	Modelling Gaussian beams	49
2.4.1	Mode-matching telescope solution	50
2.5	The power-recycled Michelson interferometer	52
2.6	Control loops	54
2.6.1	Transfer functions and their measurements	57
2.6.2	Noise, noise budget and noise projection	60
2.6.3	RIN suppression example	61
2.7	Signal read-out with interferometers	64
2.7.1	Heterodyne detection	66
2.7.2	Homodyne detection	67
2.7.3	For the Michelson interferometer	67
2.8	The contrast defect	69
2.9	Pound–Drever–Hall locking	71
2.10	Dither locking	73
2.10.1	Isolating noise and maximising the error	75
2.11	V shaped baffle beam dumps	76
2.11.1	The physics of beam dumps	78
2.11.2	Baffle design	83
2.11.3	Further analysis of powers	85
2.11.4	QUEST’s beam dumps	87
3	The QUEST Experiment	90
3.1	Sensitivity goal	91
3.2	Comparisons to similar experiments	92
3.2.1	Gravitational wave detectors	92
3.2.2	Fermilab Holometer	93
3.2.3	High frequency gravitational wave detectors	94
3.3	From the empty lab to QUEST	96
3.3.1	Laser beam profile characterisation	96
3.3.2	Prototype stages	98
3.4	QUEST’s design	101
3.4.1	AMY and BOB	102
3.4.2	Co-located	103
3.4.3	The power-recycled Michelson interferometers	103
3.4.4	Bench layout	104
3.4.5	The injection breadboards	105
3.4.6	Vacuum chambers	106
3.4.7	The lasers	108
3.4.8	The end test mass actuators	109
3.4.9	Interferometer mirrors	110
3.4.10	Mode matching telescopes and lenses	112
3.4.11	Contrast defect at QUEST	112
3.4.12	Readout	112
3.4.13	Squeezing	113
3.4.14	Scientific data acquisition	114
3.5	QUEST’s lock scheme	115
3.5.1	Lock acquisition	116

3.5.2	DC lock	119
3.5.3	Locking interface	120
3.5.4	Angular to longitudinal control coupling	122
3.6	Loop characterisation	123
3.7	The lock	124
3.8	QUEST Sensitivity (model)	126
3.9	Summary and discussion	128
4	The Output Mode Cleaner	130
4.1	Motivation for the OMC	132
4.2	Design considerations	135
4.2.1	Geometry	135
4.2.2	Finesse	135
4.2.3	Angle of incidence	140
4.2.4	Cavity Layout	141
4.2.5	Vacuum requirement analysis	144
4.2.6	Housing	145
4.2.7	The OMC actuator	148
4.2.8	Design specifications summary	150
4.3	Detection breadboard	150
4.4	OMC Control	152
4.4.1	Lock acquisition and auto-scan	154
4.5	Loop characterisation	155
4.5.1	Loop electronics	155
4.5.2	First servo design	157
4.5.3	Final loop transfer functions	158
4.5.4	Maximising the error signal	160
4.6	The lock	161
4.6.1	OMC lock test with a polarising beam splitter	163
4.6.2	Further ambient condition test	166
4.7	OMC characterisation	169
4.7.1	Mirror reflectivity	169
4.7.2	Length scan analysis	170
4.7.3	OMC losses	174
4.8	Commissioning lessons	176
4.9	Summary	178
5	Fibre noise cancellation method for constructing an X-arm vs science laser error signal at LIGO	180
5.1	Advanced LIGO	180
5.1.1	Challenges	184
5.2	Fibre noise cancellation method	185
5.2.1	Fibre noise cancellation at LIGO	187
5.2.2	Summary	188
5.3	The work undertaken	189
5.3.1	The breadboard(s) design	190
5.3.2	The photodetectors	191
5.3.3	Installing the breadboards	193
5.4	Results for the fibre noise cancellation scheme	194
5.5	Summary	196

6	Conclusion and outlook	198
6.1	QUEST and the OMC	198
6.1.1	Future work — QUEST	199
6.1.2	Future work — The OMC	200
6.2	LIGO	201
6.2.1	Future work — LIGO	202
	Appendix	204
A	CDS and models in Simulink	206
A.1	CDS	206
A.2	SIMULINK	208
B	Quarter waveplate argument	214
C	Mechanical dimensions	218
C.1	Beam dumps	218
C.2	Gluing jig	219
C.3	OMC housing	219

List of Figures

1.1	The light cone	7
1.2	The causal diamond	8
1.3	The light sheet/horizon of the Michelson interferometer	9
1.4	Representation of the effect a gravitational wave has on a test mass ring	16
1.5	Unequal passes through a beam splitter of the X and Y fields	20
1.6	Simple Michelson interferometer	23
1.7	Power at the anti-symmetric port as a function of differential-mode arm length	24
1.8	Overlapping horizons of independent Michelson interferometers	28
2.1	The electromagnetic fields of the interferometer	31
2.2	Intensity distributions of the Hermite-Gaussian modes	36
2.3	Amplitude distributions of the Hermite-Gaussian modes	37
2.4	Geometry of the Gaussian beam	38
2.5	Electric fields of the Fabry-Pérot cavity	39
2.6	Optical cavity geometries	41
2.7	Powers around the Fabry-Pérot cavity as a function of cavity length	42
2.8	Cavity power variations as a function of cavity couplings	43
2.9	Cavity misalignment	46
2.10	Steering mirrors and the beam walk	47
2.11	Mode-matching examples	48
2.12	q-parameters around a Fabry-Pérot cavity	50
2.13	Mode-matching telescope model	51
2.14	The power-recycled Michelson interferometer	53
2.15	Simple control loop	55
2.16	Simple control loop with noise input	56
2.17	Bode plots of high and low-pass filters	58
2.18	Bode plots of various filter combinations	59
2.19	Measuring the open loop transfer function	59
2.20	Laser power noise suppression loop	61
2.21	Bode plots for laser power stabilisation loop	62
2.22	Noise budget and result of laser power suppression loop	64
2.23	Detecting mirror displacement using local oscillator	65
2.24	Common readout techniques for the interferometer	68
2.25	Contrast defect at the anti-symmetric port	70
2.26	Phases around the Fabry-Pérot cavity as a function of cavity length	71
2.27	Pound-Drever-Hall error signal	73

2.28	Comparison between the transmission peak of a cavity and a parabola	74
2.29	Three common beam dumps	76
2.30	QUEST's glass beam dumps 3D model	77
2.31	Reflection beam cross-sectional area	79
2.32	Transmission beam cross-sectional area	80
2.33	Brewster's angle definition	82
2.34	BRDF analysis of various glass options	83
2.35	Interior angles of the V-shaped beam dump	83
2.36	Power in reflection of the dump, for both polarisations	85
2.37	Beam propagation through two pieces of glass	86
2.38	Overall power in reflection and transmission of the V-shaped beam dumps	87
2.39	The beam dump glass preparation	88
2.40	The beam dumps in use	89
3.1	QUEST design sensitivity	92
3.2	Synchronous recycling interferometer	95
3.3	Knife-edge characterisation data at two distances	97
3.4	Measured beam radii as a function of distance	98
3.5	Fabry-Pérot prototype Finesse	99
3.6	Fringe lock of prototype Michelson interferometer	100
3.7	Sensitivity curve for prototype Michelson interferometer	101
3.8	QUEST's experimental layout	102
3.9	The layout of QUEST	105
3.10	Injection breadboard layout	106
3.11	The vacuum chambers of QUEST	107
3.12	Vacuum pipe connecting vacuum chamber of QUEST	108
3.13	The ETM mounts and actuators	110
3.14	End test mass PZT locations	110
3.15	Characterisation of high reflectivity optics	111
3.16	Data acquisition signal flow	115
3.17	QUEST lock acquisition layout	117
3.18	Pound-Drever-Hall error signal offset	118
3.19	Anti-symmetric port pick-off	120
3.20	LSC overview screen	121
3.21	Signal flow for the three filter bank stages in LSC	122
3.22	Longitudinal to angular coupling	123
3.23	Open loop transfer function of QUEST's DARM loop during DC lock	124
3.24	2-minutes of DC lock	125
3.25	Weekend DC lock data from the X-arm transmission photodetector	126
3.26	Noise budget for QUEST sensitivity measurement	127
3.27	Measured displacement sensitivity curve for one of QUEST's power-recycled Michelson interferometers	128
4.1	Photos of the OMC	131
4.2	Transmission of modes from the OMC	132
4.3	Higher order mode transmission of triangular cavity as a function of cavity g-factor	138
4.4	The lengths of the linear and triangular cavities	139
4.5	BRDF vs. angle of incidence	140

4.6	OMC angles and lengths	141
4.7	Angle of incidence on the concave mirror, and how it influences spot position within the OMC	143
4.8	Index of refraction of air changing with pressure and temperature . .	145
4.9	Bandwidth as a function of round-trip-length	146
4.10	OMC original housing	147
4.11	OMC updated housing	147
4.12	The OMC actuator	148
4.13	Actuator resonance as a function of load	149
4.14	Actuator hysteresis	149
4.15	Detection breadboard layout	151
4.16	OMC control overview screen	153
4.17	OMC lock sequence	154
4.18	OMC electronic circuit	156
4.19	PZT and CDS RC circuit transfer functions	156
4.20	OMC dither lock control loop	158
4.21	Measured transfer functions of the OMC dither loop	159
4.22	Measured and modelled transfer functions of the OMC dither loop .	160
4.23	Error point spectra with tuning demodulation phase	161
4.24	Four days of OMC lock	162
4.25	Test setup for the four day lock	163
4.26	Polarisation stability test setup	164
4.27	Polarisation stability test results	165
4.28	Polarisation stability test spectra	165
4.29	Weekend ambient condition test - Control signal	166
4.30	Weekend ambient condition test - Actuator signal	167
4.31	Weekend ambient condition test - Transmission signal	167
4.32	Weekend ambient condition test - Error signal	168
4.33	Isolating frequencies in the spectra	169
4.34	OMC length scan - transmission peaks and actuator voltage	170
4.35	Actuator voltages for resonances identification	171
4.36	Transmission peak bandwidth fit	172
4.37	OMC scan with higher order modes identified	174
4.38	OMC transmission throughput	175
4.39	OMC length scan with 45° linearly polarised input light	176
4.40	Laser frequency hop regions as a function of crystal temperature . .	177
4.41	OMC length scan with dual fundamental resonance and poor input polarisation	177
5.1	The 5 longitudinal degrees of freedom for Advanced LIGO	181
5.2	Layout of the arm length stabilisation system	183
5.3	Proof of concept experimental layout	185
5.4	Proof of concept results	186
5.5	Modified layout of the arm length stabilisation system	188
5.6	Advanced LIGO lock scheme flow charts	189
5.7	The fibre breadboard layout	190
5.8	Photodetector volts to watts conversion	192
5.9	Photodetector shot noise limit test layout	192
5.10	Photodiode shot noise limit required power	193

5.11	Fibre breadboards installed in the science laser enclosure and the X-end station	194
5.12	Fibre-quarter waveplate paddle	194
5.13	Fibre cancellation method results	195
A.1	CDS with whitening and de-whitening filters	207
A.2	Nyquist frequency example	207
A.3	OMC Michelson interferometer specific Simulink model	209
A.4	Simulink model for OMC lock	211
A.5	OMC Simulink model DEMOD box	211
A.6	OMC Simulink model AUTOSCAN box	212
A.7	Filter bank screen	212
C.1	The beam dump platform dimensions	218
C.2	The jig dimensions	219
C.3	The OMC housing dimensions	219

List of Tables

3.1	Laser beam profiles	108
3.2	QUEST optical specifications	111
3.3	End test mass PZT signal distribution matrices	116
3.4	Sensitivity measurement details	126
4.1	OMC properties for QUEST and gravitational wave detectors.	131
4.2	OMC design specifications	150
4.3	OMC mirror power reflectivities	170
4.4	OMC optical properties	172

Acronyms

AC Alternating Current.

ADC Analogue to Digital Converter.

AOI Angle Of Incidence.

AS Anti-Symmetric.

ASC Angular Sensing and Control.

ASD Amplitude Spectral Density.

BRDF Bidirectional Reflectance Distribution Function.

BS Beam Splitter.

C-PIT Common Pitch.

C-YAW Common Yaw.

CARM Common Arm Length.

CD Contrast Defect.

CDS Control and Data System.

CSD Cross-Spectral Density.

D-PIT Differential Pitch.

D-YAW Differential Yaw.

DAC Digital to Analogue Converter.

DARM Differential-mode Arm Length.

DC Direct Current.

DFO Dark Fringe Offset.

EOM Electro-Optic Modulator.

ETM End Test Mass.

F-QWP Fiber Quarter Waveplate.

F-Refl Fiber Retroreflector.

FI Faraday Isolator.

Fib Fiber optic.

FSR Free Spectral Range.

FWHM Full Width at Half Maximum.

GW Gravitational Wave.

HOM Higher Order Modes.

HP High Pass.

HV High Voltage.

HWP Half Wave Plate.

IFO Michelson Interferometer.

IR Infrared.

ITM Input Test Mass.

JamMT Just another mode matching tool.

LIGO Laser Interferometer Gravitational-Wave Observatory.

LLO LIGO Livingston Observatory.

LO Local Oscillator.

LP Low Pass.

LSC Length Sensing and Control.

OMC Output Mode Cleaner.

PBS Polarising Beam Splitter.

PD Photodetector.

PDH Pound-Drever-Hall.

PLL Phase Lock Loop.

PMC Pre-Mode Cleaner.

PR-IFO Power-Recycled Michelson Interferometer.

PRM Power Recycling Mirror.

PSD Power-Spectral Density.

PSL Pre Stabilised Laser.

PZT Lead Zirconate Titanate (Pb Zr Ti) - also commonly referred to PiezoElectric Transducer.

QUEST Quantum Enhanced Space-Time.

QWP Quarter Wave Plate.

RC Resistor-Capacitor.

REFL Reflection/Reflected.

RF Radio Frequency.

RIN Relative Intensity Noise.

RoC Radius of Curvature.

SHG Second Harmonic Generator.

SNR Signal to Noise Ratio.

SRM Signal Recycling Mirror.

TEM Transverse Electromagnetic.

...I would rather have questions that can't be answered
than answers that can't be questioned.

Richard P. Feynman

Introduction

There are two theories of Physics which have managed to break free of their academic shackles and influence the wider world, making appearances in almost all modes of pop-culture from movies to poems: General Relativity (GR) and Quantum Mechanics (QM). Though not everyone could describe nor have any formal education of them, most have heard of them — *almost* everyone knows the name Einstein and *almost* everyone knows the phrase ‘sub-atomic’. The same probably cannot be said of the majority of the fundamental physical theories, even greats like the laws of thermodynamics. Something entices the mind about the universal scale of GR and the science fiction-like world of QM. Both are titans in their own right and represent the leading methods of understanding the universe; from the smallest of scales, the quantum, to the biggest of scales, the cosmological. However, they are incompatible — the crossover from where one steers the ship to the other is still not very well understood.

It is quite remarkable that the same phenomena — the interference of light — which was used to provide experimental evidence for quantum mechanical principles such as wave-particle duality [1–3], a cornerstone of quantum field theory, has also been used to observe a prediction of GR [4], the gravitational wave (GW) [5]. Modern gravitational wave detectors are modified versions of the Michelson interferometer [6] which use optical cavities to enhance the sensitivity to differential arm length change — the mechanism that generates an interference-born output signal in the Michelson interferometer. On September 14th 2015, gravitational waves generated via the inspiral and merger of two black holes were detected for the first time ever [7] by the two Laser Interferometer Gravitational-wave Observatory (LIGO) detectors, confirming once again GR’s status as the leading theory of gravity. Since then, many other gravitational wave detections have been made [8–11].

The event of the first detected gravitational waves briefly radiated more energy in the form of gravitational waves, than the combined light energy of all the stars in the observable universe for the same time period [12]. These gravitational waves then propagated for around 1.3 billion light-years [7, 12, 13], and were detected by displacing mirrors by less than a thousandth the diameter of a proton.

That is gravity at the macroscale, but what about the quantum scale? Theories of quantum gravity [14–16] reliably seem to conclude that spacetime, the fabric of the universe, is itself *quantised* and length measurements of these spacetime quanta have an inherent uncertainty. Repeated length measurements, drawing closer and closer to a Planck-like precision, will soon enough arrive at an irreducible variance [17–22] — the quantum limit of length accuracy. Observations of a quantum spacetime could provide much needed clues about how gravity manifests at the quantum level, and be the building blocks of a unifying bridge between GR and QM. It is convenient then that through the decades long push to detect gravitational waves, the technology required to be sensitive enough to length fluctuations to attempt observing quantum spacetime signatures, is now here. It seems quite fitting that it is once again the phenomenon of light interference that may be used to achieve this goal.

Building on the experimental breakthroughs of the gravitational wave detectors, it is conceivable that a quantum spacetime signal can be detected and the team behind the QUEST experiment at Cardiff University’s Gravity Exploration Group is striving to do just that. QUEST is a pair of co-located table-top Michelson interferometers with power-recycling and quantum enhancements, and the goal to reach unprecedented displacement sensitivity level in the bandwidth of 1 – 250 MHz. The main science goal is detecting/setting new upper limits for a quantum spacetime signal, with auxiliary goals of ultra-high frequency gravitational waves and scalar field dark matter research. I was fortunate enough to start my PhD only a handful of weeks before the optical tables were installed in the QUEST lab. Since their installation, I have been involved in every aspect of the optical and control system setup. I was also granted a 4 month fellowship at LIGO Livingston. My project there was to investigate a suggested modification to the lock acquisition process of the detector cavities. The modification would remove a substantial amount of auxiliary optics which each introduce loss.

My thesis is presented in six chapters:

1. This chapter introduces the physical phenomena QUEST aims to investigate and shows how each of them is expected to result in a signal at the output of a length measuring instrument. This will lead to a description of Michelson interferometers and why they are the ideal detectors for these kinds of signals. Their theoretical workings and capabilities will be introduced as well.
2. I will describe the fundamental experimental considerations and methods used throughout the PhD.
3. This chapter is where I detail QUEST’s experimental setup, the commissioning efforts, showcase the lock scheme and the latest achieved results from commissioning.

-
4. This chapter is focused on QUEST's OMC. I will discuss its function, design, lock process and characterisation. This is one of the aspects of QUEST that sets it apart from the nearest competitor experiment — without the OMC, QUEST's groundbreaking sensitivity could not be reached.
 5. In this chapter, the work at LIGO Livingston will be discussed. It begins by describing the current method used to obtain a science laser vs X-arm error signal and why that perhaps is not the best method. Then the alternative method of constructing that error signal which is less optically expensive is described. Finally, the work carried out in installing this update will be detailed.
 6. The concluding remarks, outlook and summary of the work undertaken during my PhD as well as anticipated future work.

Appendix sections include, QUEST's control and digital system, a proof that a quarter waveplate and a mirror can be used equivalently to a half waveplate and mechanical dimensions of some components I designed.

Chapter 1

Fundamental physics with table-top interferometers

This chapter introduces the theory of the science goals for the QUEST experiment. Each will be described as a physical phenomenon, with emphasis of how they are expected to produce a length fluctuation which can in principle be detected by a sufficiently sensitive experiment.

1.1 Quantum Gravity

Quantum mechanics and General relativity are irreconcilable in their current forms in describing the gravitational interaction at the quantum scale [23–25]. Otherwise successful theories in their own right, neither can describe nature in terms of the other. Attempts to reconcile the incompatibilities are ongoing [14–16] and tend to arrive at a similar conclusion: spacetime is quantum in nature and as such is subject to an uncertainty in terms of length [17–22]. Experimental evidence of this length uncertainty is yet to have been recorded, but the length noise predicted by quantum theories of gravity based on the holographic principle [26] are perhaps within reach. Irreducible variance in repeated length measurements could provide the foundations for reconciliation — a bridge between quantum mechanics and general relativity. The basis for the theory will be described as follows.

Firstly, the concepts of light cones and causal diamonds will be introduced, together with the idea of a Rindler horizon. From the description of this specific horizon, the origin of the theories of quantum spacetime will be introduced together with why an irreducible variance of repeated length measurement is expected. This, along with the entropy bound of the horizon, will highlight why (thanks to the holographic principle) a possible detection of quantum spacetime fluctuations is potentially possible.

1.1.1 Light cones and causal diamonds

Light cones define the causally connected regions from an event, X . The z axis of Fig. 1.1a is time and the $x - y$ plane is space. The perimeter of the cone is the null line, i.e. the path a photon would take were it travelling in that direction, which is at 45° with respect to both axes.

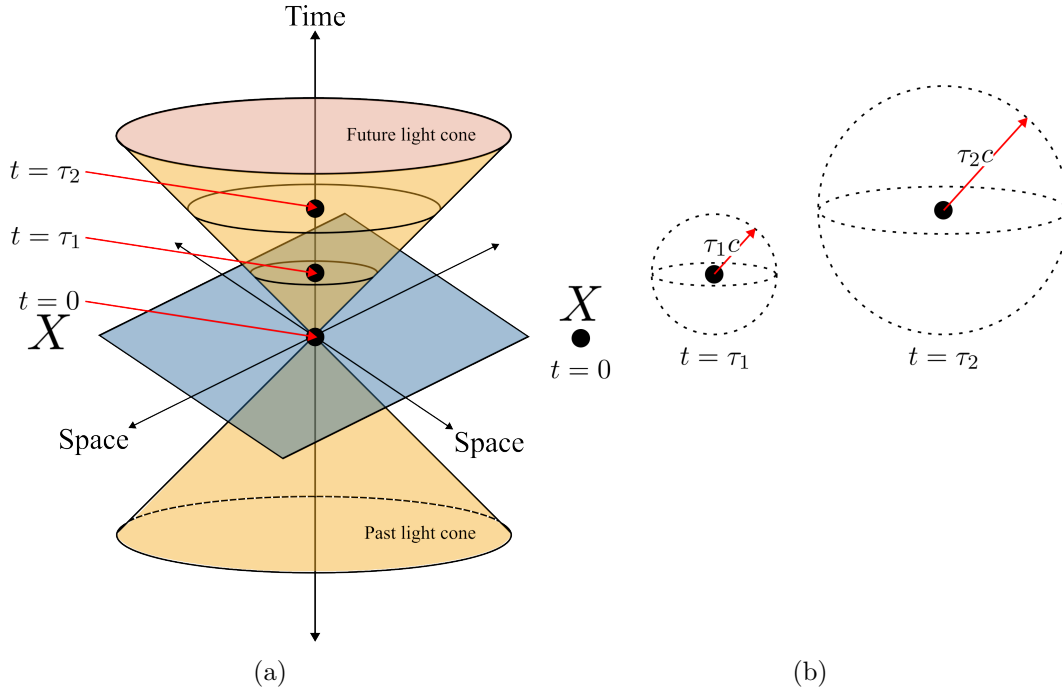


Figure 1.1: (a) A light cone, where time is the z axis and space is confined to the $x - y$ plane. The surface of the cone is the null surface, the region separating causally connected events. Within the cone events are in causal contact. (b) Sphere projections of the cone's cross section at particular times $t = 0$, $t = \tau_1$ and $t = \tau_2$. Space is 3-dimensional with time determining the sphere's radius. The surface of the sphere is the light sheet or horizon, it is the membrane between causally connected events.

The future light cone of X encompasses all events that can be affected by X while the past light cone of X contains all events that can have had an influence on it. Everything outside the cones requires faster than light speed in order to have reached and influenced X , or for it to have reached and influenced X . This is why these regions are regarded as not in causal contact.

In a 3-dimensional space, the light cone is actually a sphere; the projection of which onto the $x - y$ plane is the circle of points where a plane at constant z intersects the cone. Every time step τ forwards (or backwards), the sphere radius r grows (or is reduced) by $r = \tau c$ — see Fig. 1.1b. This sphere surrounding an event is the boundary between that which is causally connected to the event and that which is not and is defined as a light sheet.

The individual light cones of two causally connected events lead to the definition of

the causal diamond, i.e. where the future of event X_A meets the past of event X_B . Fig. 1.2a shows the respective light cones of these two events and Fig. 1.2b shows where the future of X_A and the past of X_B overlap — the causal diamond between X_A and X_B .




Figure 1.2: (a) Individual light cones of causally connected events X_A and X_B . (b) The overlap between events X_A and X_B — X_B lies within the future light cone of X_A and *ipso facto*, X_A lies within X_B 's past. The surface of the cones is the causal diamond between X_A and X_B .

The plane at the centre of the causal diamond (which in a 3-dimensional space looks like a sphere/light sheet) can be defined as a conformal Killing horizon or Rindler horizon [27].

If a Michelson interferometer (introduced in Sec. 1.4) is used for some experiment as in Fig. 1.3, then in line with the principles shown in Figs. 1.1 and 1.2, event X_A would be two photons separating at the beam splitter and heading down their respective orthogonal arm. The noteworthy limit of X_A 's future light cone is when those photons impinge on the end mirrors. This projects a light sheet with a radius equal to the arm length, Fig. 1.3a. Event X_B is then those two photons recombining and interfering back at the same beam splitter. The noteworthy limit of X_B 's past light cone is when those photons reflect from the end mirrors. This also projects a light sheet of the same radius, Fig. 1.3b.




Figure 1.3: The light sheet/horizon of the Michelson interferometer, centred on the beam splitter. (a) Shows the first event and its future light cone which in 3-dimensional space is a sphere. The event is a pair of photons separating at the beam splitter. The limit of the future light cone is where those photons impinge on their end mirror. (b) Shows the second event and its past light cone. This event is when those two photons recombine and interfere at the same beam splitter. The limit of the past light cone is where those photons reflect from the end mirrors. The dashed lines are the perimeter of a sphere which represent the light sheets of the causally overlapping events, they have a radius equal to the Michelson interferometer arm length.

This thought experiment shows how measurements taken with a Michelson interferometer trace a causal diamond. The light sheet is shown in Fig. 1.3 as the dashed sphere; it has a radius equal to the length of the Michelson interferometer arms. This light sheet surrounding the experiment could be thought of as an *experimental horizon* — the region around the experiment which separates that which is in causal contact with it and that which is not.

The importance of the experimental horizon (and horizons in general) is related to the surface area of its projection and what that surface area says about the entropy content of the volume of spacetime it encapsulates; it is related to the surface area of black holes. That connection, and why the experimental horizon is critical to the quantum gravity theories under investigation, will be highlighted in the next section.

1.1.2 Horizons

The following discusses the surface area of the black hole event horizon. In some theories of quantum gravity there is a comparison drawn between the surface area of the event horizon and the surface area of a general causal diamond — specifically how the information content of a volume of spacetime is expected to be limited. What follows is a purely geometric argument, it is not related to observers moving through various frames passing through, or observing something passing through the event horizon of a black hole.

The entropy content of a black hole is described by Bekenstein-Hawking entropy (S_{BH}) [28] and it is required for black holes to be compatible with the laws of ther-

modynamics; without this entropy, the information paradox [29] is left unaddressed.

Arguments for assuming black holes have an associated entropy include,

- The information within the black hole is not observable. Since black holes can be formed by the gravitational collapse of a star, which carries entropy, the thermodynamic description of that collapse is apparently lost behind the horizon. In order to address this, the black hole itself must have an implicit entropy, to recover the lost information.
- In thermodynamics, there are many microstates associated with a system which can be described by one observable macrostate. Some classes of black hole are completely described by only a few observable parameters [30], i.e. mass, electric charge and spin; all determined by the formation and evolutionary history of the black hole. But many formation scenarios can result in similar values for each of the observable parameters or macrostates. In this way, the information of a black hole is analogous to a thermodynamic system — many microstates (that which formed the black hole) lead to but a few observable macrostates. The black hole can therefore be described as having an associated entropy.
- Entropy can be described as a measure of missing information, or uncertainty [31]. A black hole's event horizon prevents the observation of all information which passes through. Therefore, the black hole can be considered as a place for information to become uncertain.

These are logical arguments, based largely on the reluctance to break the laws of thermodynamics. This is important because of what the entropy of a system says about the information content permitted within. If the entropy can be linked to the surface area of the horizon, not the volume that horizon encompasses, the physical ramifications are vast.

Entropy is also present in quantum systems due to the uncertainty of their microstate. If a quantum system A is entangled with another quantum system B , the entropies of A or B can not reach zero independently. Their microstates are mutually dependent, i.e. it is not possible to exactly know one without exactly knowing both. This uncertainty is known as entanglement entropy S_{ent} [28].

In leading theories of quantum gravity, the entropy of a black hole and the entanglement entropy of its horizon coincides [28, 32],

$$S_{BH} = \frac{A}{4G} \cdot \frac{k_B c^3}{h} = S_{ent} \quad (1.1)$$

where A is the surface area of the horizon and G is Newton's gravitational constant. The constants of k_B , h and c (Boltzmann's constant, Planck's constant and the speed

of light, respectively), have been included here for dimensional completeness, they are not present in [28, 32]. According to Eq. 1.1 both S_{BH} and S_{ent} are restricted to the surface area of the black hole rather than the volume, meaning that the information within the black hole is entirely encoded on the surface of the horizon. This equivalence is reached using various different approaches of fundamental physics [32–34] but, most importantly, the result of Eq. 1.1 can be shown to hold for causal diamonds as well [35–38] which is why horizons are important to describe. This link implies that the entropy content of a causal diamond (including that of empty space or some experiments) is restricted to its surface area, not volume.

It can be shown that the quantum uncertainty of horizon position is related to the entropy of a system S_{sys} [28] by

$$\delta L^2 \sim \frac{L^2}{\sqrt{S_{sys}}} \quad (1.2)$$

where L is the radius of the horizon. Since the uncertainty depends on entropy, both the macro (L) and the micro ($\ell_p = 1.6163 \times 10^{-35}$ m, the Planck length) scales are included and, and if we assume 4-dimensions [28],

$$\delta L \sim \sqrt{\ell_p L}. \quad (1.3)$$

This has the appearance of a random walk relationship [39] following root- N statistics¹ and it implies Planck length fluctuations of position. A photon traversing the horizon would follow a random walk trajectory suggesting the spacetime itself has an attributed quantum uncertainty. For the QUEST experiment as described in Chapter 3 which has arm length $L = 1.8$ m, we expect $\delta L \sim 5.4 \times 10^{-18}$ m. Initial LIGO had sufficient sensitivity to detect this level of δL [40], but the signals are expected to be in the MHz. No version of LIGO (or any of the large gravitational wave detectors) was designed for that bandwidth. Therefore, unfortunately, it is not expected that archived/future LIGO/Virgo/GEO600 data would contain these signals.

This restriction of entropy to the surface area horizon is commonly called holography [26, 34]. If sufficiently accurate repeated samples in the position of a causal diamond/horizon are measured, quantum fluctuations in the geometry of the spacetime bounded within could be observed [28].

The conclusion is that the surface area of the casual diamond restricts the entropy within and that fluctuations in spacetime, i.e. length measurements, occur. These fluctuations are not expected to be wholly independent but to be in some way

¹If particles travelling at velocity v have an average inter-particle separation R , then the time between collisions is $\tau = R/v$. The number of steps taken in measurement time T is $N = T/\tau$ and the expectation value of position goes as $\Delta x \sim \sqrt{N}\tau \sim \sqrt{T\tau}$.

entangled. These factors combined point towards a discontinuous magnitude of the spacetime fluctuations (which limits precision), potentially bringing them within reach of measurement.

1.1.3 Reality check

So far, what has been discussed is quite abstract. The key point is that measurements of spacetime, achieved by some detector repeatedly measuring length, will have some fundamental precision limit; analogous to how a particle's measured momentum is limited in precision if its position is arbitrarily well known. If this detector could measure the distance between two objects at a Planck-like accuracy while that distance was free of any perturbation, the distance measured would not be the same each time, there would be some uncertainty which cannot be surpassed — a result of the ‘size noise’ present in the spacetime quanta which together form the distance separating the objects. This uncertainty is not arbitrarily small though, the entropy is limited by the size of the detector. The entropy contains the information, and so if it is limited, the information is too; including the length variance precision. Furthermore, if two detectors are placed closely such that they overlap and share a large portion of the same spacetime volume, an entanglement between the spacetime quanta would produce coherent fluctuations in both detectors. This entanglement is expected, again, due to the entropy limit. In this case, the detector(s) is assumed to be a Michelson interferometer.

1.1.4 Measuring quantum gravity effects at horizons and length fluctuations

The causal diamond/light sheets shown in Fig. 1.2 encapsulates the whole experiment. We can break it down into a series of nested diamonds representing each quantum step the photons take along the arms. Given the uncertainty in the position of a horizon, and the random walk nature of the fluctuations, we can assume each sequential horizon is statistically uncorrelated with a scale of $\tilde{\ell}_p$. Since the position of each horizon is independent of the others, this uncertainty, in 4-dimensional space, sums to an experimentally observable fluctuation in length given by [28]

$$\delta L^2 = \frac{\ell_p L}{4\pi} \tag{1.4}$$

which scales with \sqrt{L} showing a random walk style uncertainty². For an instrument making repeated length measurements of the horizon geometry, this will appear as an irreducible noise at the output. Two identical instruments, like Michelson interferometers occupying the same or strongly overlapping causal diamonds, will

²For QUEST, with arm length 1.8 m, this gives $\delta L \sim 1.5 \times 10^{-18}$ m

measure correlated fluctuations. The entanglement of the spacetime fluctuations at the overall horizon requires fluctuations to be present in both instruments. Time integration of their cross-correlated data will allow the attenuation of uncorrelated noise and the strengthening of correlated noise.

1.1.5 Signal confirmation

A problem in experimental work is providing evidence that a given signal is indeed born of the phenomena under investigation. The question is, how might a correlated noise be confirmed as having an origin in entanglement due to overlapping detectors?

A suggested reasonable starting method for gaining corroborating evidence would be to incrementally separate the detectors until they are no longer overlapping. The correlated noise or signal should diminish with distance and be non-existent where the detectors are entirely separate.

1.2 Gravitational Waves

Before Einstein, the theory provided by Newtonian gravitation predicted gravity to be a force acting between two masses. The magnitude of the force is given by Gm_1m_2/r^2 where G is Newton's gravitational constant, $m_{1,2}$ are the masses of the bodies between which the force is exerted and r is their separation. This view poses some problems, for example, it suggests the force acts instantaneously if changes in the variables occur. For instance, if the Sun were to disappear, we would be relieved of our orbital elliptical path around it the instant it happened. On the other hand, it would take ~ 8 minutes to see the light go out raising the contradiction: the gravitational information of the Sun's disappearance arrives instantly, while the light, travelling at the maximum speed limit permitted by the laws of physics, does not.

In the Einstein theory of gravitation, gravity is not conceived as an instantaneous force felt between two objects, but rather a consequence of the effect those masses have on the spacetime around them [30]. The concept of a spacetime first originated in Einstein's 1905 work [41], though it was not then formulated exactly as spacetime. In 1908, Minkowski [42] took seriously the idea that space and time were in fact inseparable phenomena and that the fabric of the universe itself could be described as a spacetime. In 1915 the idea was made famous by Einstein's theory of general relativity [4].

In this theory, it is postulated that the presence of mass warps spacetime itself, changing the geometry of space: previously straight paths, the shortest distance between two points in a space where Euclidean geometry applies, become curved.

This curvature describes why two masses appear to be attracted to one another. The Earth travelling around the Sun is following the shortest path in the curved geometry, the so-called ‘geodesic’.

1.2.1 Wave solution to Einstein’s equations in the weak field limit

The core of general relativity is represented by Einstein’s field equations, which describe the relationship between the geometry of spacetime and the energy/matter distribution [30]:

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \quad (1.5)$$

In Eq. 1.5, $G_{\mu\nu}$ is the Einstein tensor which describes the geometry of spacetime. $T_{\mu\nu}$ is the stress-energy tensor that characterises the density of energy and momentum in spacetime and c is the speed of light. The indices μ and ν represent the temporal and spatial component of the tensor, in which the notation 0 – 3 has been adopted, with the index 0 referring to the time component and the indices 1 – 3 to the spatial ones, respectively. Later on (Sec. 1.2.2) the subscripts $\mu\nu$ will be switched to ij — in this form, only the spatial coordinates are present. The physical meaning of Einstein’s equation is perfectly summarised by the sentence of physicist John Archibald Wheeler: “*Spacetime tells matter how to move; matter tells spacetime how to curve*” [43].

Using the definition of Einstein’s tensor, Eq. 1.5 can be rewritten as

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{8\pi G}{c^4} T_{\mu\nu} \quad (1.6)$$

where $R_{\mu\nu}$ is the Ricci tensor, R the Ricci scalar and $g_{\mu\nu}$ is the metric tensor, which is the generalised gravitational potential describing the geometric and causal structure of a local spacetime. Equation 1.6 is a set of ten nonlinear partial differential equations in four independent variables, that is, it is not possible to get an analytical solution. Some assumptions about $g_{\mu\nu}$ can be made to simplify the calculations. For example, far from a region of mass sources, spacetime is flat (Minkowskian) and it is described by the metric tensor $\eta_{\mu\nu}$ given by

$$g_{\mu\nu} \equiv \eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1.7)$$

where the $(-, +, +, +)$ notation has been adopted. Considering a Minkowskian spacetime as a background with a little perturbation $|h_{\mu\nu}| \ll 1$, we can rewrite the metric tensor as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (1.8)$$

This is the so-called weak field limit and by introducing³ the trace-reversed metric $\bar{h}_{\mu\nu}$ [5, 44] where

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h \quad (1.9)$$

and now Einstein's field equations can take the form of a wave equation [5, 44]:

$$\square\bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \quad (1.10)$$

where $\square = \nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}$ is the D'Alembertian operator. Far from matter sources, $T_{\mu\nu} = 0$ and the wave equation becomes:

$$\square\bar{h}_{\mu\nu} = 0. \quad (1.11)$$

With this operator, $h_{\mu\nu}$ becomes $h(2\pi ft - \mathbf{k} \cdot \mathbf{x})$ where $f = |\mathbf{k}|c/2\pi$ represents a wave propagating at c , in the $\hat{k} = \mathbf{k}/|\mathbf{k}|$ direction [5, 44]. f is the wave frequency and $k = 2\pi/\lambda$ is the wavenumber with λ being the wavelength. This is a light speed wavelike solution to Einstein's field equations. We see that gravitational radiation can exist. This gravitational radiation is the gravitational wave.

1.2.2 Towards gravitational wave detection

Considering a wave travelling in the \hat{z} direction; in the transverse traceless (TT) gauge⁴, $h_{\mu\nu}^{TT}$ is given by [5]

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & a & b & 0 \\ 0 & b & -a & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1.12)$$

This wave can be described as a combination of two orthogonal polarisations, the so-called plus and cross polarisations

$$h_{\mu\nu}^{TT} = \hat{h}_+ + \hat{h}_\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & 0 & 0 \\ 0 & 0 & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & h_\times & 0 \\ 0 & h_\times & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.13)$$

³This is a way to simplify the calculation to arrive at the final solution of the wave equation [5, 44].

⁴The TT gauge is used because gravitational waves have an especially simplified form in this gauge [5].

A key step is to now quantify the effect a gravitational wave would have on some kind of detector, that we can imagine as a ring of test masses at rest with the ring defining the detector's proper frame. The coordinate distance of each test mass from the centre of the ring is given by ξ_i as shown in Fig. 1.4.

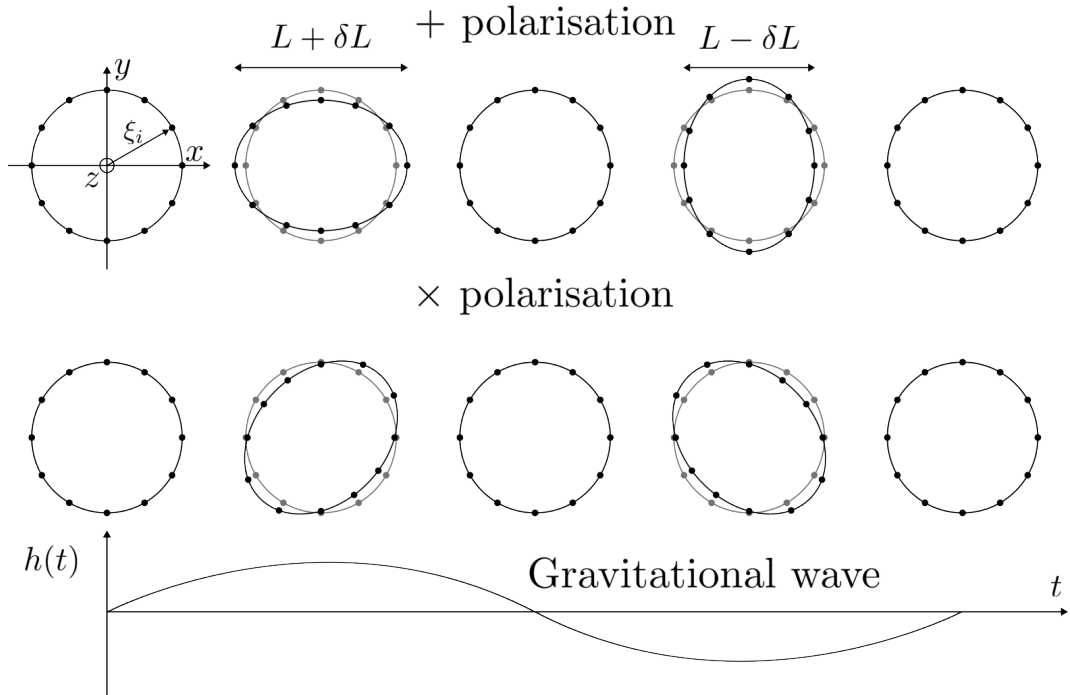


Figure 1.4: Representation of the effect a gravitational wave has on a test mass ring, in the $x - y$ plane.

A gravitational wave propagating in the \hat{z} direction has tensor components $h_{33} = 0$, leaving displacement (along with our the test mass ring) confined to the $x - y$ plane. For the plus polarisation, assuming $h_{ij}^{TT} = 0$ at $t = z = 0$ we get

$$h_{ab}^{TT} = h_+ \sin(\omega t) \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (1.14)$$

with $a, b = 1, 2$, the indices in the transverse plane [5, 44]. Since ξ_i is the distance to a test mass from the origin, with no perturbation we have $\xi_i = x_0 = y_0$. But if a perturbation is displacing the test masses by a distance $\delta x(t)$ and $\delta y(t)$, we have

$$\xi_i(t) = (x_0 + \delta x(t), y_0 + \delta y(t)), \quad (1.15)$$

and then for the plus polarisation

$$\delta \ddot{x} = -\frac{h_+}{2}(x_0 + \delta x)\omega^2 \sin(\omega t), \quad \delta \ddot{y} = +\frac{h_+}{2}(y_0 + \delta y)\omega^2 \sin(\omega t). \quad (1.16)$$

At the first order, the terms δx and δy can be neglected allowing immediate inte-

gration of the equations:

$$\delta x(t) = \frac{h_+}{2} x_0 \sin(\omega t), \quad \delta y(t) = -\frac{h_+}{2} y_0 \sin(\omega t). \quad (1.17)$$

The same approach for the cross polarisation yields

$$\delta x(t) = \frac{h_\times}{2} y_0 \sin(\omega t), \quad \delta y(t) = \frac{h_\times}{2} x_0 \sin(\omega t). \quad (1.18)$$

That is, the gravitational wave introduces a time dependent strain on the test mass ring/detector. Provided that a perpendicular wave direction relative to the detector plane, a quadrupole pattern emerges — shown in Fig. 1.4.

1.2.3 Length fluctuations — Strain

In general relativity, a geodesic describes the trajectory of a particle in spacetime that is freely falling in a gravitational field and a geodesic deviation describes how a that trajectory changes in a curved spacetime [5, 44]. The equation of the geodesic's deviation from flatness in the detector's proper frame takes the form [5, 44],

$$\ddot{\xi}^i = \frac{1}{2} \ddot{h}_{ij}^{TT} \xi^j. \quad (1.19)$$

This shows explicitly how a gravitational wave warps spacetime which has a local effect of causing distance changes between objects, ΔL (for example, test masses of a real detector⁵). Assuming we have a plus polarised gravitational wave travelling perpendicularly to the detector like a Michelson interferometer, the change in its arm lengths due to the gravitational wave passage is given by [5, 44]

$$\Delta L \sim \frac{1}{2} h L \quad (1.20)$$

where h is the amplitude of the gravitational wave. Measuring ΔL requires an instrument capable of detecting length fluctuations of order $h/2$. For the first gravitational wave detection, 'GW150914', h was of the order of 10^{-21} [7].

1.2.4 Gravitational wave frequency

In terms of binary inspiral and merger events, the frequency of a gravitational wave is a function of the system mass M , the speed of light c and Newton's gravitational constant G . Therefore, there is then a range of frequencies for the mass of the system as [45–47]

$$f_{\text{inspiral}} \lesssim (0.02 \rightarrow 0.05) \frac{c^3}{GM} \quad ; \quad f_{\text{ringdown}} \approx (0.06 \rightarrow 0.15) \frac{c^3}{GM} \quad (1.21)$$

⁵ ξ^i is the distance between two geodesics — with two test particles a fixed distance apart, this is constant in flat spacetime. It can only vary in a warped spacetime. L is *just* the distance separating the test particles, i.e. it is the length of the detector.

For QUEST the bandwidth is 1 – 250 MHz [48]. This suggests that were black hole mergers to be detected, it would require a system mass of (16μ to 10 m) solar masses. Black holes of this mass could be primordial black holes [49, 50].

Other sources of gravitational waves in the 1 – 250 MHz region include stochastic gravitational wave background and black hole superradiance [48].

1.3 Dark Matter

Since as early as the 1930s astronomers and cosmologists have suspected issues with our understanding of gravity and particle physics [51, 52]. It did not take very long for there to be a general scientific consensus [53] that there must be something out there that they could not see yet. Different observations (such as galactic rotation velocities, gravitational lensing and the angular anisotropy of the cosmic microwave background) [54–58] all arrive at the same conclusion: baryonic matter is estimated to account for less than 5% of the whole mass content of the universe. That which we cannot see was termed dark matter and dark energy. Only dark matter candidates will be described hereon.

The search for dark matter is ongoing. Low-mass scalar field dark matter models predict interesting effects that can in principle be detected by sensitive enough detectors. Here, why scalar field models are of unique interest, particularly to interferometer based experiments, will be described.

1.3.1 Scalar fields

Scalar field dark matter models predict early universe vacuum misalignment mechanisms producing low-mass ($m_\phi \ll 1\text{eV}$) particles, which would be distributed throughout the universe as a coherently oscillating field with angular Compton frequency $\omega_\phi = m_\phi c^2 / \hbar$ [59].

In these models, the dark matter field ϕ couples with standard model field suggesting a potential detection mechanism. The interaction term is an addition to the standard model Lagrangian which, considering only interactions linear in ϕ , results to be [59]

$$\mathcal{L}_{\text{int}}^{\text{lin}} = \frac{\phi}{\Lambda_\gamma} \frac{F_{\mu\nu} F^{\mu\nu}}{4} - \frac{\phi}{\Lambda_e} m_e \bar{\Psi}_e \Psi_e \quad (1.22)$$

where $F^{\mu\nu}$ is the electromagnetic field tensor, m_e the rest mass of the electron, $\bar{\Psi}_e$ and Ψ_e is the standard model electron field and its Dirac conjugate, respectively. Λ_γ and Λ_e parameterise the coupling of the dark matter field with photons and electrons, respectively.

1.3.2 Fluctuations in fundamental constants

The physical effect of the additional term in Eq. 1.22 is to induce changes to the fine structure constant α and electron rest mass m_e values. Fluctuations in these fundamental constants imply (oscillatory) changes in the physical size l and refractive index n of solids, which are parameterised as [59]

$$\frac{\delta l}{l} = \left(-\frac{\delta\alpha}{\alpha} - \frac{\delta m_e}{m_e} \right) \left(1 - \frac{\omega^2}{\omega_0^2} \right)^{-1} \quad (1.23)$$

and

$$\frac{\delta n}{n} = -5 \times 10^{-3} \left(2 \frac{\delta\alpha}{\alpha} + \frac{\delta m_e}{m_e} \right). \quad (1.24)$$

respectively.

1.3.3 Length fluctuations

An instrument sensitive to length and/or refractive index changes could in principle be exploited for searching for scalar field dark matter-induced effects. Michelson interferometers serve as ideal candidates since nowadays their sensitivity to length changes of the arms are second to none.

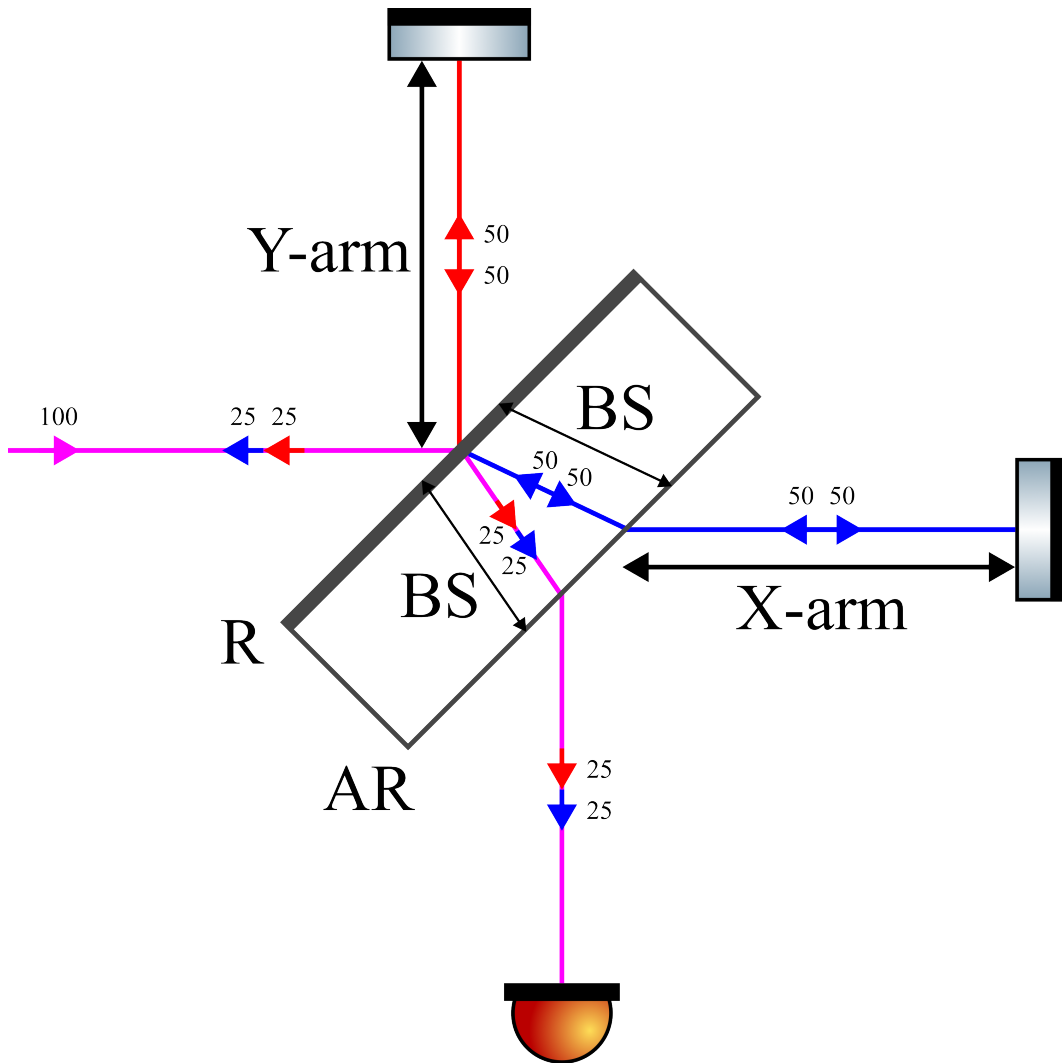


Figure 1.5: Beam splitter paths. The purple beam coming from the left is the injection beam, it originally inputs 100% of the power. The red beam is the Y-arm path after prompt reflection from the 50/50 reflective (R) coating of the beam splitter. The end mirrors are assumed 100% reflective. The red beam (shown now as a red arrow head) then travels through the beam splitter once more towards the photodetector. The blue beam is the X-arm path. It is transmitted through the 50/50 beam splitter surface and travels through the beam splitter before reaching the X-arm. Upon returning to the beam splitter, it travels through the beam splitter once more before recombining with the red beam, and a final time through the beam splitter now with the blue beam, towards the photodetector.

An additional benefit of the interferometer specifically, is that the central beam splitter interacts asymmetrically with the beam path of each arm. This will enhance a differential signal, i.e. the output to the interferometer. This is a result of the 50/50 reflective coating on the incident edge and of the purely anti-reflective coating on the rear edge of the beam splitter. The beam promptly reflected travels in the Y-arm and passes through the beam splitter only once — see Fig. 1.5 — while the transmitted beam travels in the X-arm and passes through the beam splitter three times.

The oscillation of the dark matter field would also change the size of the end mirrors, which will affect the path length of each arm as well. But since the rest size of the mirrors is intentionally equal, the effect will be largely symmetric, resulting in a negligible signal in an interferometer's output. Also, for the frequencies of interest, the wavelength of the scalar field (λ_ϕ) is expected to be at least three orders of magnitude larger than detector arm lengths [60], further minimising an end mirror's contribution to any possible dark matter signal.

Changes to the beam splitter size, δL , and refractive index, δn , will result in a differential optical path length of the arms as [60]

$$\delta(L_x - L_y) \approx \sqrt{2} \left[\left(n - \frac{1}{2} \right) \delta l + l \delta n \right]. \quad (1.25)$$

The refractive index is expected to provide a negligible contribution, i.e. have an affect \sim three orders of magnitude lower than that of size fluctuations [60], so it will not be considered in the following calculations.

1.3.4 Signal searches

The fluctuations in beam splitter size due to the oscillatory density distribution of the dark matter field we are moving through are expected to produce Doppler-shifted and Doppler-broadened signals in interferometers as [60]

$$\delta(L_x - L_y) \approx \left(\frac{1}{\Lambda_\gamma} + \frac{1}{\Lambda_e} \right) \left(\frac{n l \hbar \sqrt{2 \rho_{\text{local}}}}{m_\phi c} \right) \cos(\omega_{\text{obs}} t) \quad (1.26)$$

where $\rho_{\text{local}} = 0.4 \text{ GeV/cm}^3$ [59] is the local dark matter density and ω_{obs} the field's angular frequency, relative to the observer. Standard galactic dark matter halo models give an expected frequency linewidth of $\Delta\omega_{\text{obs}}/\omega_{\text{obs}} \sim 10^{-6}$ [60].

With two identical and co-located interferometers, the magnitude of a signal obtained from cross-correlated data⁶ is [59]

$$\delta(L_x - L_y) \approx \left(\frac{2cl\sqrt{\rho_{\text{local}}}}{\omega_\phi} \right) \text{sinc}^{-1} \left(\frac{\omega_\phi L}{c} \right) \left\{ \left(n - \frac{1}{2} \right) \left(\frac{\omega_{\text{BS}}^2}{\omega_{\text{BS}}^2 - \omega_\phi^2} \right) \left(\frac{1}{\Lambda_\gamma} + \frac{1}{\Lambda_e} \right) + n \left(\frac{10^{-2}}{\Lambda_\gamma} + \frac{5 \times 10^{-3}}{\Lambda_e} \right) \right\} \quad (1.27)$$

where the sinc function is related to the frequency response of an interferometer with arm length L , resulting in a periodic signal modulation and ω_{BS} is the (angular) mechanical resonance frequency of the beam splitter. ω_ϕ is again the field's angular Compton frequency. The $\left(\frac{\omega_{\text{BS}}^2}{\omega_{\text{BS}}^2 - \omega_\phi^2} \right)$ feature appears like a har-

⁶Cross-correlation is a signal processing technique which serves to assess how similar multiple data inputs are. In this case, the two data paths are Fast Fourier Transformed and then multiplied.

monic oscillator which suggests any driving frequency higher than the beam splitter resonance will be attenuated.

Searches for these kinds of dark matter signals have been recently carried out using data from the GEO600 gravitational wave detector [60] and Fermilab’s Holometer [59]. GEO600 is more sensitive than the other operational gravitational wave detectors (LIGO and Virgo) to signals originating at the beam splitter. This is largely due to the fact they do not use Fabry-Pérot cavities in the arms [60], and their use of world record levels of squeezing (6 dB [61]). Fermilab’s Holometer [62] consists of two co-located interferometers, ideally suited for cross-spectrum analysis which enables further noise suppression. The Holometer will be further described in Chapter 3.

1.4 The Michelson Interferometer

From the theories discussed so far, it turns out that an instrument sensitive to fluctuations in length, like an interferometer, is ideally suited to potentially detect the above phenomena: quantum gravity signatures, high frequency gravitational waves and scalar field dark matter. There are different types of interferometer configurations, yet each operates with the same general principle — interfering light to detect variations in the resulting pattern due to changes in length [63–66]. The specific type we will focus on is the Michelson interferometer [6], exploiting a Gaussian laser beam (detailed in Sec. 2.2) as a light source.

Here, the relationship between the arm lengths and the output signal generated by a Michelson interferometer will be shown.

Optical components are described by their field amplitude reflection and transmission coefficients, r and t , respectively. In terms of reflected and transmitted power, the coefficients become $R = r^2$ and $T = t^2$, respectively. We use the convention that

$$E_{refl} = rE_c \tag{1.28}$$

and

$$E_{trans} = itE_c \tag{1.29}$$

where E_c is the input field, commonly called the carrier⁷. The origin of this convention is in Sec. 2.1. The Michelson interferometer depicted in Fig. 1.6 uses a 50:50 ($R = T = 0.5$) beam splitter to split the carrier field E_c into two orthogonal arms of length L_x and L_y . At the end of each arm there is a highly reflective mirror ($R \approx 1$) which redirects the respective field back towards the beam splitter where

⁷The carrier refers to the *main* field at any point, in contrast to any other field frequencies present, such as those used for control purposes, or signals generated within the Michelson interferometer — this is discussed in more detail in Sec. 2.1.1.

interference between the two fields occurs. The two outputs of the beam splitter are the symmetric and anti-symmetric ports, with fields E_S and E_{AS} , respectively. These ports are also commonly referred to as the bright and dark ports, respectively.




Figure 1.6: The Michelson Interferometer. E_c is the input (carrier) field which is split with a 50:50 ratio at the beam splitter into the two arms of length L_x and L_y . The fields propagate along each arm, get reflected from the end mirrors and return to the beam splitter where their fields interfere. The two outputs of the beam splitter are shown as the symmetric port E_S which is directed back towards E_c and the anti-symmetric port E_{AS} which is directed *downwards* and out of the instrument.

Both outputs can be used for signal detection, but conventionally the anti-symmetric (AS) port is used. Using the convention detailed in [67] (abridged in Sec 2.1), the field at the anti-symmetric port is given by

$$E_{AS} = i \frac{E_c}{2} [e^{i\Phi_x} + e^{i\Phi_y}], \quad (1.30)$$

where $\Phi_x = 2kL_x$ and $\Phi_y = 2kL_y$ are the total phases accumulated in the crossings of X- and Y-arm, respectively, $k = 2\pi/\lambda$ is the wavenumber and λ is the wavelength of the carrier laser. Therefore,

$$E_{AS} = i \frac{E_c}{2} [e^{2ikL_x} + e^{2ikL_y}]. \quad (1.31)$$

Defining the common-mode arm length (CARM) \bar{L} as the average of the two arms and the differential-mode arm (DARM) length ΔL as the difference between the two, we have

$$\bar{L} = \frac{L_x + L_y}{2}, \quad \Delta L = L_y - L_x, \quad (1.32)$$

from which

$$L_x = \frac{2\bar{L} - \Delta L}{2}, \quad L_y = \frac{2\bar{L} + \Delta L}{2} \quad (1.33)$$

and thus we can rewrite Eq. 1.31 as

$$E_{AS} = iE_c e^{2ik\bar{L}} \cos(k\Delta L). \quad (1.34)$$

The physical quantity we are interested in measuring is the power $P_{AS} = E_{AS}E_{AS}^*$, which we can write from Eq. 1.34 as

$$P_{AS} = P_c \cos^2(k\Delta L) \quad (1.35)$$

where $P_c = E_c E_c^*$ is the power of the carrier field. Fig. 1.7 is a plot of Eq. 1.35 showing the output of the Michelson interferometer at the anti-symmetric port as a function of ΔL . At the maximums, there is perfect constructive interference and they are called the bright fringe. Minimums, where there is perfect destructive interference, are the dark fringe.




Figure 1.7: Michelson interferometer output at the anti-symmetric port. The peaks are commonly referred to as the bright fringe, here all the light is constructively interfering and exiting via the anti-symmetric port. The troughs are called the dark fringe, where all the light is destructively interfering and no light exits at the anti-symmetric port. The transmission at the symmetric port is $P_S/P_c = 1 - P_{AS}/P_c$ (neglecting losses).

1.4.1 Noise

The noises in the interferometer we are interesting in are sensing and displacement noise.

The sensors in question here are photodetectors which convert the power detected by the sensor to an output voltage. They do this by virtue of their quantum efficiency, which converts power to a photocurrent, and their transimpedance gain (provided by a resistor) which converts that photocurrent to volts. There are two main types of sensor noise. The resulting Johnson noise (which is a thermal effect due to the temperature of the resistor [68]) resulting in an output voltage variance. This is not a function of the power at the sensor; it exists when the sensor is covered and could be considered as the sensor's noise floor. This noise must be exceeded in order for the power fluctuations at the sensor to be detected. The other sensor noise is the photocurrent noise driven by photon shot-noise (described in Sec. 1.4.2)

which is a quantum photon counting error — this also results in a variance in the sensor’s output voltage. Provided there is sufficient power in the beam (which is a requirement that must be ensured for experimental work), the photocurrent or shot-noise exceeds the Johnson noise and is then the limiting sensor noise.

From Fig. 1.7, sensor noise could be understood as a noise in the y-axis. If the sensor responsible for measuring P_{AS}/P_c is outputting a voltage with variance, there is no way of determining whether that variance is due to interesting scientific phenomena driving DARM displacement, or simply some artifact of instability in the voltage output of the sensor.

Displacement noises are phenomena which move the optics and therefore lead to erroneous signals at the output since they are indistinguishable to the scientifically interesting signals. From Fig. 1.7 this could be understood as a noise in the x-axis. Seismic motion is an example of this category, but radiation pressure is an example of displacement noise as well, as will be described in Sec. 1.4.3. Improvements to displacement noises require displacement isolation systems, whether that be passive or active. For example, in LIGO the optics are suspended as a series of pendula, each providing isolation to displacement noise from the suspension point above; this is a passive measure. They also use active feedback systems to counteract seismic noise coupling to the optical tables from which the optics are suspended [69].

1.4.2 Shot-noise

As described above, differential changes in Michelson interferometer arm lengths are measured by power fluctuations at the anti-symmetric port. The maximum sensitivity to DARM length changes that an experiment can achieve depends on the stability of the power measured where no length change occurs, i.e. if the optics were perfectly still then ideally the output of the detector would be constant. The high-frequency sensitivity is often limited by photon shot-noise which has its roots in the Heisenberg uncertainty relationship between the number N_γ of and phase ϕ of photons, i.e. $\Delta N_\gamma \Delta \phi \geq 1$ [48].

The uncertainty in the number of photons in a laser beam ΔN_γ follows Poisson statistics [70], i.e. $\Delta N_\gamma = \sqrt{N_\gamma} \equiv \sqrt{\tau P / \hbar \omega}$ where τ is the measurement time, P the beam power, and ω the angular frequency of the photons.

The uncertainty in the phase can be related to the uncertainty in propagation distance (i.e. uncertainty in Michelson interferometer arm length ΔL) by $\Delta \phi = k \Delta L = 2\pi \Delta L / \lambda$, where λ is the wavelength of the laser.

The shot-noise-limited displacement amplitude spectral density (ASD) in a Michelson interferometer is given by [48, 71]

$$S_{\Delta L}^{SN}(f) = \sqrt{\frac{c \hbar \lambda}{4\pi P_{BS}}} \quad (1.36)$$

where P_{BS} is the power on the beam splitter. Eq. 1.36 shows that to increase the signal-to-shot-noise-ratio the options are to decrease λ or increase P_{BS} . Changing the laser's wavelength is not so simple, because, for example, optics have coatings designed for a specific wavelength, with narrow bandwidths, and the same is true for photo detectors — they operate within a narrow bandwidth as well. An experiment could be designed with this in mind and set out to use a much smaller wavelength laser. However, a three order of magnitude decrease in $\lambda = 1064\text{ nm}$ puts the light at the far end of UV if not x-ray light, which is not ideal for optical experiments where people are constantly nearby.

The most practical method of improving $S_{\Delta L}^{SN}$ is to increase power on the beam splitter. In order to do this, besides input power increase, a particular optical device — a resonant cavity — is exploited. A cavity can provide a three order of magnitude increase in power on the beam splitter quite readily, as described in Sec. 2.3.

1.4.3 Radiation pressure

The uncertainty in the number of photons per unit time has a second noise effect. Each photon imparts momentum to the optics they impinge [72] which means a displacement in the mirrors proportional to square root of power in the beam [73] — or equivalently, to the number of photons arriving. If the number of photons was constant, this displacement effect is cancelled by virtue of the differential arms each receiving an equal momentum transfer. However, the uncertainty in the number of photons at each optic leads to a stochastic displacement noise — the radiation pressure noise, ΔL_{RP} .

Radiation pressure has an inverse proportionality to power compared to shot-noise [73] i.e.

$$\Delta L_{RP} \propto \sqrt{P} \tag{1.37}$$

and so there is a balance to be struck between the two. Increasing the laser power will reduce the shot-noise floor, but will increase the radiation pressure noise. However, radiation pressure is a low frequency noise; for experiments like LIGO and Virgo, it is the dominant quantum noise at frequencies below $\sim 50\text{ Hz}$, above which shot-noise becomes dominant [74, 75], which means it will not be the main limiting noise for experiments looking in the kHz-MHz region.

1.4.4 Standard quantum limit

The standard quantum limit is the lowest noise floor a classical experiment can achieve without some novel quantum mechanical based intervention. Once all other noise sources are addressed, the quantum noises of shot-noise at the sensor and radiation pressure at the optics are limiting. They are a feature of quantum mechanics,

they cannot be suppressed by feedback loops. A novel method fast becoming standard practice at advanced high sensitivity interferometer experiments is to inject squeezed states of light (a method referred to as ‘squeezing’), the process of reducing one quadrature of the quantum noise of photons (number or phase) at the expense of the other, resulting in a higher signal-to-noise (SNR) than would otherwise be possible by bringing noise below the standard quantum limit.

When discussing noise, it is common to do so in the context of the SNR rather than noise alone. For example, when aiming to reduce the shot-noise limit the power can be increased. This works because of the different relationships the signal and the noise have to that increased power, as described earlier. But for something like radiation pressure or seismic noise, the improvement is likely only going to reduce noise while having no impact on the signal. So, when discussing improvements to noise it may be a direct reduction of noise, or as a byproduct of increasing both the signal and the noise amplitudes by different amounts. In both approaches, the SNR is ultimately increased.

1.4.5 Co-located Michelson interferometers

If two nearby Michelson interferometers were operated simultaneously and independently, their time averaged cross-spectrum could be used to strengthen any correlated signals while mitigating the uncorrelated noise sources. This principle of the data from two separate detectors being cross-correlated is of benefit to all phenomena which generate an output signal in Michelson interferometers — assuming non-interesting correlated noises, such as seismic noise, are sufficiently suppressed. Further, if the two independent Michelson interferometers were so nearby that they actually spatially overlapped or were ‘co-located’, this dual Michelson interferometer configuration would be particularly constructive for the quantum gravity search.

Two co-located Michelson interferometers which share a strongly overlapped region of spacetime define a common horizon. The light sheet diagram of Fig. 1.3 is expanded on to include these two co-located Michelson interferometers in Fig. 1.8. Otherwise independent entangled fluctuations are present and correlated to both instruments due to this spatial overlap, and the time integrated cross-spectrum of their separate data will establish a convenient method of enhancing the signals generated within their common horizon. As described in Sec. 1.1.5, a signal believed to have an entanglement origin could be scrutinised by varying the overlap of the detectors. There should be a relationship between separation and signal strength, with an expected signal extinction when they are no longer overlapping.

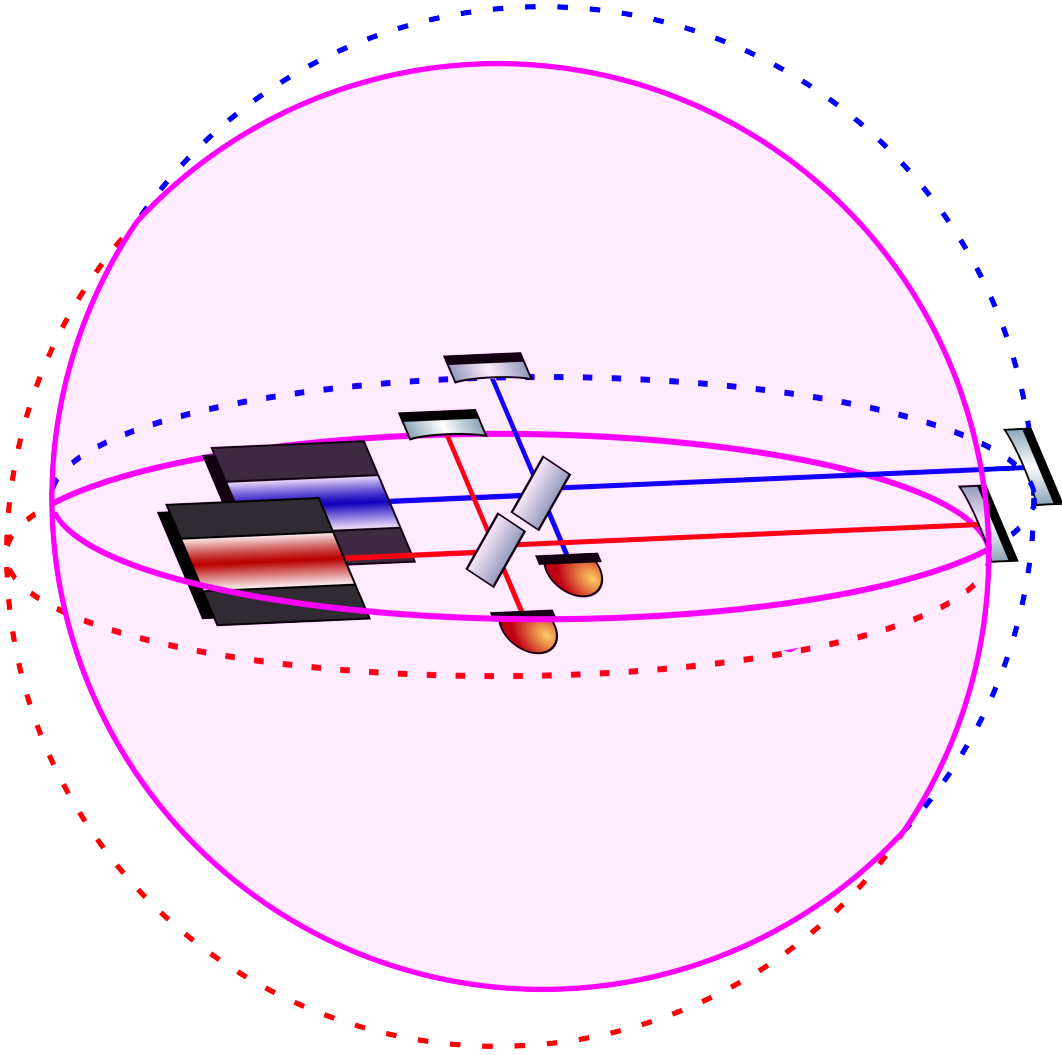


Figure 1.8: The overlapping independent Michelson interferometers generate a common horizon and causal diamond. Entangled length fluctuations within each of the separate causal diamonds are now correlated between the two instruments and will appear as a common signal in their cross-spectrum.

1.5 The quantum enhanced spacetime experiment

The *QU*antum *E*nanced *S*pace-*T*ime (QUEST) experiment [48] exploits the principles of the co-located Michelson interferometer with the express purpose of establishing new upper limits for quantum spacetime signals. It is a table-top experiment, of the order of 2 m in length with a data-acquisition rate capable of detecting individual photon cross times, making it sensitive to differential-mode length fluctuations in the 1 – 250 MHz region. Chapter 3 will describe QUEST’s experimental layout, control systems and commissioning efforts to date. Chapter 4 details the Output Mode Cleaner, an optical cavity in the detection path which plays a key role in facilitating QUEST’s groundbreaking sensitivity.

Chapter 2

From theory to experiment

This chapter aims to give an overview of concepts which are important throughout the thesis. Those familiar to interferometry, Gaussian beams, optical cavities and control schemes may not need the revision. The majority of the topics here are commonly accessible in such works as Bond [67], Saulson [76], Hecht [77], Siegman [78] and Abramovici [79].

The topics of this chapter are included given their foundational importance to the overarching themes of this thesis: the design and uses of high sensitivity optical devices, control schemes and commissioning. Where derivations are included, they are intended to provide those new to the topics a somewhat all encompassing overview, alleviating the need to refer to citations for a general understanding.

Where the work here offers a perhaps unique or original approach is in Sec. 2.4.1 where the mode-matching telescope solution generator is detailed, and Sec. 2.11 where the design of glass beam dumps, which hugely reduce scattering, is described.

2.1 Field propagation of the Michelson interferometer

This derivation is an abridged version of that in Bond et al. [67]. The major results are the anti-symmetric (AS) port field and the origin of the complex coefficient added to a field with transmission through an optic.




Figure 2.1: The electromagnetic fields of the interferometer.

The electric field \vec{E} is described by [67]

$$\vec{E}(x, y, z, t) = E_c \vec{e}_p \cos(\omega t - \vec{k} \cdot \vec{r} + \phi) \quad (2.1)$$

where E_c is the carrier field amplitude, \vec{e}_p the polarisation direction, $\omega = 2\pi\nu$ the angular frequency with ν the frequency, $\vec{k} = \omega/c$ the wavenumber, \vec{r} the direction of propagation and ϕ is a phase offset.

In the case of a linearly polarised field propagating along the z axis, Eq. 2.1 becomes

$$E(z, t) = E_c \cos(\omega t - kz + \phi). \quad (2.2)$$

When the electric field inside a Michelson interferometer arrives at the beam splitter, the field is split in the two arms according to the reflection and transmission amplitude coefficients r and t , respectively, with the phase changing ϕ_r and ϕ_t , respectively. Following the electromagnetic fields propagation shown in Fig. 2.1, we have

$$E_{X1} = tE_c e^{i\phi_t} \quad , \quad E_{Y1} = rE_c e^{i\phi_{r1}} \quad (2.3)$$

where ϕ_{r1} is the phase reflection coefficient from the *front* side of the beam splitter. We cannot assume the reflection from either side is equivalent, but because of symmetry we can assume equivalency of the transmission ϕ_t , in both directions [67].

The fields E_{X1} and E_{Y1} both propagate down their respective arms and return. Assuming the end mirrors (referred to in the field of gravitational wave detection as the end test masses, ETMs) to be fully reflective ($r_{ETM} = 1$) and to have no influence on the phase ($\phi_{r,ETM} = 0$), the fields coming from the arms and approaching the

beam splitter are

$$E_{X2} = tE_c e^{i(\phi_t + \Phi_X)} \quad ; \quad E_{Y2} = rE_c e^{i(\phi_{r1} + \Phi_Y)} \quad (2.4)$$

where Φ_X and Φ_Y are the total accumulated phase in the horizontal and vertical arms, respectively.

Following the second interaction with the beam splitter, the fields go as

$$E_S = E_{X2,trans} + E_{Y2,refl} = E_{X2}t e^{i\phi_t} + E_{Y2}r e^{i\phi_{r1}} = E_c \left[T e^{i(2\phi_t + \Phi_X)} + R e^{i(2\phi_{r1} + \Phi_Y)} \right] \quad (2.5)$$

$$E_{AS} = E_{X2,refl} + E_{Y2,trans} = E_{X2}r e^{i\phi_{r2}} + E_{Y2}t e^{i\phi_t} = E_c r t \left[e^{i(\phi_t + \phi_{r2} + \Phi_X)} + e^{i(\phi_t + \phi_{r1} + \Phi_Y)} \right] \quad (2.6)$$

where $R = r^2$, $T = t^2$ and ϕ_{r2} is the phase reflection coefficient from the *back* side of the beam splitter. For convenience, the phase factors are separated into the common and differential ones

$$\alpha_+ = \phi_{r1} + \phi_t + \frac{1}{2}(\Phi_X + \Phi_Y) \quad (2.7a)$$

$$\alpha_- = \phi_{r1} - \phi_t + \frac{1}{2}(\Phi_X - \Phi_Y) \quad (2.7b)$$

$$\beta_+ = \phi_t + \frac{1}{2}(\phi_{r1} + \phi_{r2} + \Phi_X + \Phi_Y) \quad (2.8a)$$

$$\beta_- = \frac{1}{2}(\phi_{r1} - \phi_{r2} + \Phi_X - \Phi_Y). \quad (2.8b)$$

This reduces Eqs. 2.5 and 2.6 to

$$E_S = E_c e^{i\alpha_+} [R e^{i\alpha_-} + T e^{-i\alpha_-}] \quad (2.9)$$

and

$$E_{AS} = E_c r t e^{i\beta_+} 2 \cos(\beta_-) \quad (2.10)$$

respectively.

For an optic, the conservation of energy requires that $R + T = 1$ (neglecting losses¹). With a 50/50 beam splitter, $R = T = 0.5$, which means $r = t = 1/\sqrt{2}$, giving

$$E_S = E_c e^{i\alpha_+} \cos(\alpha_-) \quad (2.11)$$

and

$$E_{AS} = E_c e^{i\beta_+} \cos(\beta_-). \quad (2.12)$$

¹The full relationship is $R + T + P = 1$ where P is the loss coefficient.

Conservation of energy also requires that the power out of a system must equal the power into the system, i.e. $P_c = P_S + P_{AS}$, with $P_c = E_c E_c^* = E_S E_S^* + E_{AS} E_{AS}^*$. After substitution, we have

$$E_c^2 = E_c^2 \cos^2(\alpha_-) + E_c^2 \cos^2(\beta_-) \quad (2.13)$$

$$1 = \cos^2(\alpha_-) + \cos^2(\beta_-). \quad (2.14)$$

Eq. 2.14 is true if and only if

$$\alpha_- - \beta_- = (2N + 1) \frac{\pi}{2} \quad (2.15)$$

where $N \in \mathbb{Z}$. Combining this with Eqs. 2.7b and 2.8b, this gives

$$\frac{1}{2}(\phi_{r_2} + \phi_{r_1}) - \phi_t = (2N + 1) \frac{\pi}{2}. \quad (2.16)$$

The convention most commonly used is to assume $N = 0$ and $\phi_{r_1} = \phi_{r_2} = 0$ [67], which gives $\phi_t = -\pi/2$. For keeping track of phase for power calculations, this is equal to $\phi_t = \pi/2$, or equivalently, $N = -1$.

With this convention in mind, the amplitude coefficients r and t for an optic are recast as r and it , respectively, to account for this phase change in the transmitted field, i.e. $E_r = rE_c$ and $E_t = itE_c$. Substituting β_+ and β_- back in

$$E_{AS} = E_c e^{i(\phi_t + \frac{1}{2}(\phi_{r_1} + \phi_{r_2} + \Phi_X + \Phi_Y))} \cos\left(\frac{1}{2}(\phi_{r_1} - \phi_{r_2} + \Phi_X - \Phi_Y)\right) \quad (2.17)$$

and including the particular values for each, remembering $\Phi_{X,Y} = 2kL_{X,Y}$

$$E_{AS} = E_c e^{i(\pi/2 + \frac{1}{2}(0+0+2kL_X+2kL_Y))} \cos\left(\frac{1}{2}(0-0+2kL_X-2kL_Y)\right) \quad (2.18)$$

$$E_{AS} = E_c e^{i(\pi/2 + \frac{1}{2}(2kL_X+2kL_Y))} \cos\left(\frac{1}{2}(2kL_X-2kL_Y)\right) \quad (2.19)$$

$$E_{AS} = E_c e^{i(\pi/2)} e^{ik(L_X+L_Y)} \cos(kL_X - kL_Y) \quad (2.20)$$

remembering $\bar{L} = (L_X + L_Y)/2$ and $\Delta L = L_Y - L_X$

$$E_{AS} = iE_c e^{ik(2\bar{L})} \cos(-k\Delta L) \quad (2.21)$$

\cos is an even function and so we arrive at the anti-symmetric field as defined in Eq. 1.34.

2.1.1 Signal sidebands by phase modulation

Phenomena that drive differential-mode arm (DARM) length changes to the Michelson interferometer couple power from the carrier into new fields with an offset frequency called sidebands; from the super-position principle, they are separate fields from the carrier [80]. To derive these sidebands, take from Eq. 2.4 (though the same is true for both arms [67]), including the above-mentioned convention

$$E_{Y2} = rE_c e^{i2kL_Y}. \quad (2.22)$$

If something lengthens the arm at a certain frequency, this introduces an additional phase component in the field

$$E_{Y2} = rE_c e^{i(2kL_Y + kL_Y A_p \cos(\Omega_p t))} \quad (2.23)$$

where A_p is the amplitude of the effect and $\Omega_p = 2\pi\nu_p$ is the angular frequency ν_p of the length modulation. The factor $e^{i(kL_Y A_p \cos(\Omega_p t))}$ is the phase modulation to the carrier with modulation index $m_p = kL_Y A_p$, which can be generally expanded using the Bessel Functions $J_n(m_p)$ [67] as

$$e^{im_p \cos(\phi)} = \sum_{n=-\infty}^{\infty} = i^n J_n(m_p) e^{in\phi}. \quad (2.24)$$

For small A_p , $m_p \ll 1$ and this expansion gives us

$$E_{Y2} = rE_c e^{i(2kL_Y)} \left[1 + \frac{ikL_Y A_p}{2} e^{-i\Omega_p t} + \frac{ikL_Y A_p}{2} e^{i\Omega_p t} \right] + \text{higher order terms} \quad (2.25)$$

where the terms inside the square brackets are for the leading order ($n = 0$), and the next-to-leading-orders ($n = -1$ and $n = 1$), respectively. In the approximation $m_p \ll 1$ the higher order terms contribution is so small that they can be neglected (though if m_p is significant, many non-negligible sidebands are generated), which reduces Eq. 2.25 to

$$E_{Y2} = a_c + a_{-\Omega}(t) + a_{\Omega}(t) \quad (2.26)$$

where a_c is the carrier field and $a_{-\Omega}(t)$ and $a_{\Omega}(t)$ are the lower and upper first-order sidebands, respectively, shifted from the carrier frequency by $-\nu_p$ and $+\nu_p$, respectively. This is the result for the Y-arm of a Michelson interferometer. In the case where some phenomenon drives a balanced DARM fluctuation — i.e. of equal magnitude in both arms — the sidebands generated in the X-arm are identical except they are 180° out of phase to those in the Y-arm. In this instance with the Bessel expansion and simplification, $E_{X2} = a_c - a_{-\Omega}(t) - a_{\Omega}(t)$ [67].

This coupling of light from the carrier to sidebands by differential-mode length

modulation has the effect of generating a signal at the anti-symmetric port even when the carrier is maintained at the dark fringe, since those sidebands constructively interfere at the beam splitter. Phenomena which generate CARM fluctuations also produce sidebands in the same manner, but these do destructively interfere and do not exit at the anti-symmetric port.

Other types of modulation to a field, amplitude and frequency modulation are not considered here. QUEST, which is described in Chapter 3, is interested in phase modulation sidebands generated by differential-mode fluctuations in the interferometer arm lengths. Phase modulation effects, both differential- and common-mode, are also useful for control schemes. Manually driving a length change generates sidebands in the same way and those sidebands can be used for length sensing of an optical device. This will be more apparent in Secs. 2.9 and 2.10 where specific control schemes which generate and use phase modulated sidebands are described.

2.2 Gaussian laser beam properties

The laser beams used in optical interferometry are Gaussian beams, which are solutions to the paraxial wave equation [78]. In the paraxial wave equation, the assumption is that a wave function $u(x, y, z)$ which has electric field

$$E(x, y, z) = u(x, y, z)e^{-ikz} \quad (2.27)$$

varies slowly in z (the propagation direction) compared to x and y [67]. The assumption extends to say that the second derivative of z is also very small; this is equivalent to saying that the wave vector is very close to parallel to the z axis. By substituting the field into the standard wave equation we end up with the differential equation for u

$$(\partial_x^2 + \partial_y^2 + \partial_z^2) u(x, y, z) - 2ik\partial_z u(x, y, z) = 0. \quad (2.28)$$

Now we make the paraxial wave assumptions [67]:

$$|\partial_z^2 u(x, y, z)| \ll |2k\partial_z u(x, y, z)| \quad , \quad |\partial_x^2 u(x, y, z)| \quad , \quad |\partial_y^2 u(x, y, z)|. \quad (2.29)$$

and Eq. 2.28 simplifies to the paraxial wave equation,

$$(\partial_x^2 + \partial_y^2) u(x, y, z) - 2ik\partial_z u(x, y, z) = 0. \quad (2.30)$$

A field which solves Eq. 2.30 has a paraxial beam shape when used in the form of Eq. 2.27.

There is an infinite set of solutions to Eq. 2.30, each describing a unique electromagnetic distribution of the field transverse to the beam propagation direction — which here is assumed to be z . The solutions are called modes. The modes de-

scribed in this thesis are the Hermite-Gaussian modes [81]; there are other bases the beam can be decomposed into, Laguerre–Gaussian [82], Ince-Gaussian [83] and Hypergeometric-Gaussian [84], but these will not be discussed here.

The number of nodes splitting the mode in the x, z (tangential) and y, z (sagittal) planes are given by the subscripts n and m , respectively with $n, m \in \mathbb{Z}^+$. Modes are then labelled as TEM_{nm} , where TEM stands for transverse electromagnetic. The zeroth order solution TEM_{00} is called the fundamental mode or the carrier; these three designations are used interchangeably. This is the purely Gaussian mode associated with signal generation and instrumentation control and it is the mode usually made resonant inside a cavity as will be described in the next section, Sec. 2.3. Modes where $n+m > 0$ are referred to as the higher order modes. The intensity distributions of the fundamental mode and some low order higher order modes are shown in Fig. 2.2. The field amplitude for the modes in the tangential plane is shown in Fig. 2.3. This highlights how each node is actually a phase flip in the field. This will be important when discussing the design of a triangular cavity in Sec. 4.2.2.




Figure 2.2: Intensity distributions of the Hermite-Gaussian modes [81]. Numbers refer to the amount of nodes creating separate lobes; the first number for the $n(x, z)$ plane and second for the $m(y, z)$ plane.




Figure 2.3: Higher order Hermite-Gaussian mode field amplitude for the $n(x, z)$ order mode [78].

A Gaussian laser beam is a superposition of all the modes where usually the fundamental mode is the overwhelmingly dominant mode. The Gaussian name is because of the intensity distribution observed through the beam's cross section — it is maximised at the propagation axis and *falls off* with a Gaussian distribution from that centre, see Fig. 2.4a. The size of the beam is defined from the radius ω taken where $I = I_{max}/e^2$, where I is the beam intensity and e is Euler's constant. The Gaussian profile is maintained throughout propagation, but the variance increases because the beam diverges with distance propagated. The Gaussian intensity distribution $I(z)$ at propagation distance z is [78]

$$I(z, r) = \frac{2P_c}{\pi\omega(z)^2} e^{-2r^2/\omega(z)^2} \quad (2.31)$$

where P_c is the total power of the beam, ω is the beam radius and r is the distance from the beam's propagation axis.

One of the main properties of a Gaussian beam is the waist — where the beam radius is at its smallest when converging/being focused. At the waist position z_0 the beam has radius ω_0 ; see Fig. 2.4b. Propagation distance z (which can be negative) is measured from the waist and the beam properties at any z can be calculated provided both ω_0 and z_0 are known.

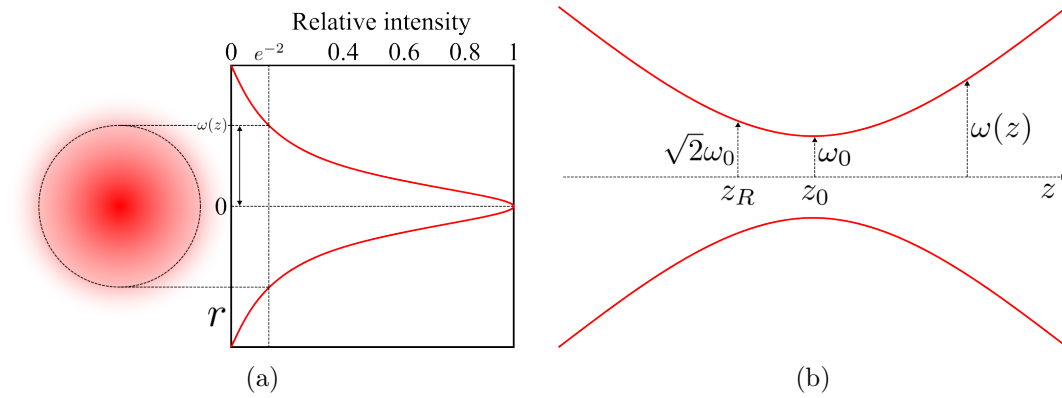


Figure 2.4: (a): Cross sectional profile of a Gaussian beam. The beam radius is defined as the distance from the axis to the point where the intensity falls to $1/e^2$ the peak. (b): Side profile of Gaussian beam through a focus point. ω_0 is the waist radius located at z_0 , z_R the Rayleigh length measured from z_0 . $\omega(z)$ is the radius of the beam at propagation distance z from z_0 .

Another important property is the Rayleigh length (or Rayleigh range) z_R , which is the distance required for the cross sectional area of the beam to double that of the cross sectional area at the waist, it is described by [78]

$$z_R = \frac{\pi\omega_0^2 n}{\lambda} \quad (2.32)$$

where λ is the wavelength of the laser and n the refractive index of the medium the beam is propagating through. The beam radius $\omega(z)$ at a given z is [78]

$$\omega(z) = \omega_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}. \quad (2.33)$$

Unlike a plane wave where the wavefront is always planar by definition, a Gaussian beam exhibits a curved wavefront as it propagates. The radius of curvature of the beam wavefront is $R(z)$ [78]

$$R(z) = z + \frac{z_R^2}{z}. \quad (2.34)$$

These properties are important for ensuring the beam is matched to the geometry of an optical device it is to impinge on, as will be more apparent in the following sections.

If a Gaussian beam of power P_{in} propagates through an iris of radius a , then the power that propagates through the iris P_{out} is given by [78]

$$P_{out} = P_{in} \left[1 - \exp\left(\frac{-2a^2}{\omega(z)^2}\right) \right]. \quad (2.35)$$

Any loss in power is due to clipping — the amount of the beam blocked due to being wider than the aperture.

A more thorough description of the Gaussian beam includes the Gouy phase [85]. When the beam is focused there is a phase shift from one side of the focus to the other with respect to that of a plane wave experiencing the same focusing. Each mode, including the fundamental, accumulates a specific amount of this Gouy phase. The Gouy phase of the fundamental and that which separates successive order modes from it is [78, 85]

$$\eta_G(z) = -\arctan\left(\frac{z}{z_R}\right) \quad (2.36)$$

Since the order of a mode is defined as the sum total of nodes, $n + m$, the shift of a particular mode is [78, 85],

$$\eta_{G, nm}(z) = (n + m + 1)\eta_G(z) \quad (2.37)$$

2.3 Optical cavities

Optical cavities are arrangements of mirrors which, under the right conditions which are described in this section, trap the electric field and make it resonate. They can be used for many different purposes either as the main component of an experiment, or as a supplementary component. Some of the main experimental uses will be described in this section.

2.3.1 Electric fields around a cavity

An exploited feature of an optical cavity is to trap the electric field and make it resonate before allowing it to pass through or be reflected. Where the field ends up is a matter of the interference conditions, which is determined by the relationship between the cavity length and the field frequency. In order to appreciate the properties of the cavity fields, we can break them down into separate infinite series — the reflected field, the field circulating inside the cavity, and the transmitted field.

Starting with the incident field, E_c , we describe all others in terms of the reflectivities (r_i) and transmissivities (t_i) of the mirrors where subscript i identifies the mirror with $i = 1$ and $i = 2$ for the input and output mirrors, respectively.




Figure 2.5: Electric fields in and around a cavity of length L . The left mirror is the input mirror with reflectivity r_1 and transmissivity it_1 . The right mirror is the output or end mirror, r_2 and it_2 .

As Fig. 2.5 shows, the reflected field E_r consists of a promptly reflected field (E_{r0}), and of the fields which have entered the cavity and completed n number of laps ($E_{r1,2,\dots,n}$), for $n = 1 \rightarrow \infty$, until transmitting back out through the first mirror in the direction of the source.

The circulating field E_{circ} begins at the immediately transmitted incoming field ($E_{circ,0}$), followed by the fields again completing n laps.

The transmitted field E_t begins with the field that enters the cavity and is transmitted through the second mirror without completing any laps. Followed by each field having completed n laps, and then transmitting through the second mirror.

The electric fields around a cavity of length L are mathematically expressed as [67, 76, 86],

$$E_r = E_c \frac{r_1 - r_2(r_1^2 + t_1^2)e^{-2ikL}}{1 - r_1r_2e^{-2ikL}} \quad (2.38)$$

$$E_{circ} = E_c \frac{it_1}{1 - r_1r_2e^{-2ikL}} \quad (2.39)$$

$$E_t = E_c \frac{-t_1t_2e^{-ikL}}{1 - r_1r_2e^{-2ikL}} \quad (2.40)$$

In terms of measured power, Eqs. 2.38 - 2.40 become,

$$P_r = P_c \frac{R_1 + R_1^2R_2 + 2R_1R_2T_1 + R_2T_1^2 - 2r_1r_2R_1 \cos(2kL) - 2r_1r_2T_1 \cos(2kL)}{R_1R_2 - 2r_1r_2 \cos(2kL) + 1} \quad (2.41)$$

$$P_{circ} = P_c \frac{T_1}{R_1R_2 - 2r_1r_2 \cos(2kL) + 1} \quad (2.42)$$

$$P_t = P_c \frac{T_1 T_2}{R_1R_2 - 2r_1r_2 \cos(2kL) + 1} \quad (2.43)$$

It is possible to see that the powers are all dependent on the relationship between the separation of the mirrors L and the wavelength of the laser (absorbed in the wavenumber k).

2.3.2 Cavity geometry

Three common layouts for the optical cavity are the linear Fabry-Pérot cavity such as those in the arms of gravitational wave detectors [87], triangular such as the LIGO input mode cleaner [88] and the bow-tie, such as that used for the LIGO Output Mode Cleaner [89]. These common configurations are shown in Fig. 2.6.




Figure 2.6: Cavity geometries. (a) is the linear Fabry-Pérot, (b) the triangular and (c) the bow-tie configurations.

For whichever geometry the cavity is designed, there is a stability requirement that exists between the cavity length and the curvature of the mirrors of which it is comprised. Each mirror of the cavity has a so called *g*-factor [44, 78]

$$g_i = 1 - \frac{L}{RoC_i} \quad (2.44)$$

where L is the cavity length and RoC_i is the radius of curvature of the i^{th} mirror. For a cavity to be stable the condition of

$$0 < g_1 g_2 g_i < 1 \quad (2.45)$$

must be met [44]. If it is met, this means that for any point with the cavity, the transverse structure and phase of the beam is repeated at that point after each round-trip of the cavity [44]. The selection of which geometry is appropriate depends on the intended use of the cavity. Though a linear cavity is perhaps the simplest to construct, it does direct all of the reflected field back towards the source while the other two direct the reflected field away. However, the reflection towards the source may actually be why the cavity is used — the arm cavities of the large gravitational wave detectors for example. There may be a constraint on the physical size of the cavity. If a large round-trip is desirable but space is limited, the bow-tie may be the best option. The question of cavity geometry was one of the first tackled in the design of QUEST’s Output Mode Cleaner, the subject of chapter 4.

2.3.3 Resonance and Finesse

Plotting Eqs. 2.41, 2.42 and 2.43, reveals the dependence on L of the powers — see Fig. 2.7.




Figure 2.7: Cavity resonance condition. Red is reflected, blue is transmitted and black is circulating power. (a): A *low quality* cavity with $r_{1,2} = 0.8$, $t_{1,2} = \sqrt{1 - r_{1,2}^2} = 0.6$. (b): A *higher quality* cavity with $r_{1,2} = 0.96$, $t_{1,2} = \sqrt{1 - r_{1,2}^2} = 0.28$. Note the y axis, more reflective mirrors lead to increased internal power. Subscripts 1 and 2 denote the input and output mirrors, respectively.

From the two examples of Fig. 2.7, (a) with low reflectivity mirrors and (b) with high reflectivity mirrors, it is clear that $R_{1,2}$ and $T_{1,2}$ play a huge role in determining how fine the resonance is and how much power is built internally, but for resonance itself, it is the relationship between L and λ that dictates when it can happen.

Eqs. 2.42 and 2.43 are maximised when the denominator is minimised, which occurs where $\cos(2kL) = 1$. This condition of maximum circulating and transmitted power (and minimised reflected power) is the resonance condition.

From the $\cos(2kL) = \cos(2 \cdot 2\pi/\lambda \cdot L) = 1$ requirement for resonance, it is clear that cavity resonance is achieved when the length of the cavity is an integer number of half wavelengths,

$$L = n \frac{\lambda}{2} \text{ for } n = 1, 2, 3, \dots \quad (2.46)$$

which is visible in Fig. 2.7. This is the condition to allow a standing wave to form within the cavity. The reflectivity of the mirrors determines the internal amplification, and the proportion of the field reflected and transmitted. Either the length of the cavity is controlled to suit the wavelength of the laser, or the wavelength of the laser is controlled to suit the length of the cavity.

The cavities in Fig. 2.7 are set up to be what's called *Impedance matched*, where the

reflectivity/transmissivity of both mirrors are the same i.e. $R_1 = R_2$ and $T_1 = T_2$ (assuming the ideal case, where there is no loss); subscripts 1 and 2 denote the input and output mirrors, respectively. This configuration is preferred for a cavity where the transmitted beam is used since, on resonance transmitted power is equal to input carrier power; $P_t = P_c$ (ignoring loss). Sometimes though, the reflected beam is used instead (such as the cavities within the arms of the large gravitational wave detectors). If $T_1 > T_2$, the cavity is called *Overcoupled* while for $T_1 < T_2$, it is called *Undercoupled*. The different configurations and the effect they have on the cavity powers are shown in Fig. 2.8.




Figure 2.8: Display of different cavity couplings. T_1 is shown on the x axis, and $T_2 = 1 - T_1$. Left shows undercoupled since $T_1 < T_2$. Right shows overcoupled since $T_1 > T_2$. Where $T_1 = T_2$ this is critically coupled or impedance matched. Red is reflected, blue is transmitted and black is circulating power.

For the impedance matched case (0.5 on the x axis of Fig. 2.8), transmitted power is greatest which could be of benefit to certain setups, while with overcoupled higher circulating power can be achieved. This is very beneficial if, for example, the objective of the cavity is to amplify the power within, or keep photons bouncing back and forth for as long as possible to sample a length change imperceptible to a single crossing.

An important parameter is the cavity Finesse \mathcal{F} which is defined as [78],

$$\mathcal{F} = \frac{\nu_{FSR}}{\nu_{cav}} \quad (2.47)$$

where ν_{FSR} is the cavity Free Spectral Range, i.e. the frequency distance between the successive transmission peaks [78]

$$\Delta\nu_{FSR} = \frac{c}{2L} = \frac{c}{L_{rt}} \quad (2.48)$$

with L as the cavity length and c the speed of light. L_{rt} is the cavity round-trip-length which for a linear cavity is $2L$.

ν_{cav} is the cavity bandwidth [78],

$$\Delta\nu_{cav} = \frac{2\Delta\nu_{FSR}}{\pi} \arcsin\left(\frac{1 - r_1 r_2}{2\sqrt{r_1 r_2}}\right) \quad (2.49)$$

where r_1 , r_2 are the mirror reflection coefficients. For highly reflective mirrors, $r_1 \approx r_2 \approx 1$ and the arcsin function can be reduced to its argument, i.e. [78],

$$\mathcal{F} \approx \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} \quad (2.50)$$

For a fixed length cavity, improvements to the mirror reflectivities will reduce the bandwidth².

2.3.4 Frequency reference

Fig. 2.7 and the equations of Sec. 2.3.1 highlight that if a cavity's length can be assumed as stable, then the cavity can be used as a witness for fluctuations in the laser's frequency — optical cavities are commonly used for this purpose. If the length is fixed in the sense that it has no actuators (but is still subject to length noise due to vibrations, thermal expansion etc), then the resonance condition is maintained by acting on the wavelength (frequency) of the laser. This is not a drastic wavelength variation, it is a minute change well within the bandwidth of the components, i.e. $\Delta\nu \ll$ the full-width-half-maximum. With a high Finesse cavity, the bandwidth is narrow which means the range of frequencies which will result in a single resonance gets smaller. Therefore, higher Finesse cavities are more challenging to control and hold on resonance, but if it can be achieved, the laser's frequency will be more stable.

2.3.5 Spatial mode filtering

In addition to power build up and frequency stabilisation, another exploited feature of optical cavities is spacial mode filtering. Assuming the fundamental mode is the signal carrier, the higher order modes in the beam do not contribute to the signal, they can therefore be considered as a pure noise source which needs to be suppressed — via spacial mode filtering. To utilise this feature, the impedance matched condition is ideal since the transmitted beam is always separate from the input and reflected beams and so can be used downstream of the cavity. This description will therefore assume the cavity is impedance matched. A consequence of the Gouy phase shift is that when entering a cavity, higher order modes have a

²For cavities with more than two mirrors, it is simply a case of including each mirror, i.e. $\mathcal{F} \approx \pi\sqrt{r_1 r_2 r_3} / (1 - r_1 r_2 r_3)$ [90].

frequency shift $\Delta\nu_{nm}$ relative to the fundamental mode [90]

$$\Delta\nu_{nm} = \frac{c}{L_{rt}}(n+m)\frac{1}{\pi}\arccos(\sqrt{g_1g_2}) \quad (2.51)$$

where g_i is the g-factor of the i^{th} cavity mirror as introduced in Sec. 2.3.2, Eq. 2.44. When the fundamental mode³ is in resonance within the cavity, and as such is transmitted through, the higher order modes present in the beam are not resonant because of the $\Delta\nu_{nm}$ frequency shift. The higher order modes are therefore reflected from the cavity providing then a purely fundamental laser mode downstream of the cavity. This is why cavities act as spacial mode filters and are said to ‘clean’ the beam (of higher order modes).

The transmission of higher order modes with respect to the fundamental mode from a linear cavity is given by [90]

$$\frac{T_{nm}}{T_{00}} = \frac{1}{\left[1 + \left(\frac{2}{\pi}\mathcal{F} \sin\left(\frac{2\pi}{L}\Delta\nu_{nm}\right)\right)^2\right]^{1/2}}. \quad (2.52)$$

2.3.6 Low-pass filter

The final interesting feature of cavities is their frequency response transfer function $H(\nu)$ given by [91]

$$H(\nu_m) = \left(\frac{t_1t_2}{1-r_1r_2}\right)^2 \frac{1}{\sqrt{1+i(\nu_m/\nu_p)}} \quad (2.53)$$

where $t_{1,2}$ and $r_{1,2}$ are the transmission and reflection coefficients for the cavity mirrors, respectively, i is the imaginary number, ν_m is some modulation frequency and ν_p the cavity pole frequency $\nu_p = \nu_{cav}/2$.

Normalised to the maximum transmission, the gain of a cavity is [91]

$$|H(\nu_m)| = \frac{1}{\sqrt{1+(\nu_m/\nu_p)^2}} \quad (2.54)$$

and the phase delay is [91]

$$\arg[H(\nu_m)] = \tan^{-1}\left(\frac{\nu_m}{\nu_p}\right). \quad (2.55)$$

This is a classic low-pass filter transfer function [79, 91] which means that if an input parameter is fluctuating, then fluctuations at frequency $\nu_m > \nu_p$ will be suppressed by the cavity which can be useful for noise suppression. However, it also attenuates

³A cavity can be made resonant with any mode, but in this case it is assumed the fundamental mode is the resonant one.

potentially interesting signals in exactly the same manner. The cavity pole is therefore a property that requires careful design so not to preemptively begin suppressing signals.

2.3.7 Higher order mode generation

Cavities have states of resonance where the Gaussian beam geometry is repeated each round-trip, these states are called eigenmodes. The eigenmode is defined by the cavity axis and the geometry of the mirrors, including separation and radius of curvature. Where the beam wavefront matches each mirror's radius of curvature, the cavity is said to be mode-matched. Imperfect alignment and/or mode-matching generate higher order modes within the cavity, with power from the fundamental mode coupling to those higher order modes. Their coupling mechanism is described well in [67] — each factor and the mitigation strategy will be described in the following sections.

Misalignment

If the input beam is misaligned to the cavity axis, either by tilt α or translation x_0 (Fig. 2.9a and Fig. 2.9b, respectively), power is coupled to first order higher order modes [92].




Figure 2.9: Cavity misalignment. (a): Angular misalignment, the beam has angle α to the cavity axis. (b): Translation misalignment, the beam is distance x_0 from the cavity axis.

Two dedicated mirrors, ideally with no other optics in between them or the cavity, are required to steer the beam into the cavity — see Fig. 2.10. The greater the separation of the mirrors, the more angular control they provide by simple virtue of trigonometry.




Figure 2.10: The beam walk. (a): the near mirror only is acted on and a maximum transmission is found, this could be assumed to be the best alignment. (b): the far mirror is now acted on and the transmission is reduced, it is a worse alignment than (a). (c): The near mirror is acted on again which brings the beam closer to the cavity axis and provides a higher transmission than in (a).

Alignment can be carried out in the form of a *beam walk*; the process of iterative improvements to the transmitted power when a length scan (or frequency shift) is occurring. Tuning one mirror followed by the other to maximise transmitted power can (almost certainly will) lead to erroneous maximums — it may appear like a maximum has been reached but it is a local maximum, not the global maximum. With the beam walk, once a maximum is reached, one mirror is intentionally misaligned and the other is used to recover and ideally improve the previous maximum.

The idea of a beam walk is shown in Fig. 2.10. Say while acting on the pitch control of the closest mirror as in Fig. 2.10a, the transmission reached a maximum. The beam is not very well aligned to the cavity axis but it is the best that can be achieved with the current position of the furthest mirror. Now consider the case where the furthest mirror is acted on as in Fig. 2.10b, the beam is worse than in Fig. 2.10a and so the transmission is reduced. However, now acting on pitch of the closest mirror again, the beam can be brought to very close to the cavity axis as in Fig. 2.10c. This will have a higher transmission than the previously achieved value of Fig. 2.10a. The beam walk is carried out for both pitch and yaw degree of freedom, it is an iterative (and somewhat laborious) process. This could be automated with some algorithm and motorised mirror mounts, which would also present the opportunity for an alignment feedback loop to maximise transmission.

The length scan will show the transmission spectrum provided it covers at least one ν_{FSR} . This can be used to simultaneously maximise the fundamental transmission while minimising the higher order modes visible in the spectrum.

Mode-matching

Mode-matching is related to the beam wavefront's radius of curvature matching that of the cavity's concave mirror (there could be multiple curved mirrors, but for the current description it will be assumed only one is present). For this to happen, the correctly sized and positioned waist must be generated within the cavity. Once the beam reaches the curved mirror from the waist, through symmetry the mirror will

focus the beam back to the waist — in this condition, the Gaussian beam retraces itself geometrically. Deviation in the size or position of the waist will result in a mismatch of the radius of curvatures and the field will not be correctly repeated each round-trip, causing a coupling of power to second order higher order modes.

Examples of mode-matching are in Fig. 2.11.




Figure 2.11: Mode-matching examples; green lines depict a matching radius of curvature where blue are incorrect. (a) — (d) show a linear cavity with a plane input and a concave end mirror. The waist must be on the plane mirror since only the curved end mirror can focus the beam back to the waist. Also a plane mirror is equivalent to a curved mirror with infinite radius of curvature and at the waist ($z = 0$), $R(0) = \infty$. (a) has too small a waist, the beam diverges too much and becomes larger than the mirror, and $R(z) > RoC$. (b) has too large a waist, $R(z) < RoC$. (c) has a correctly sized but incorrectly positioned waist. $R(z)$ is wrong for both mirrors. (d) is ideally mode-matched, $R(z) = RoC$. (e) shows a symmetrical cavity with equally concave mirrors either end, requiring the waist be positioned at the center of the cavity. (f) is a triangular cavity with two plane mirrors and one concave mirror. The plane mirrors do not focus the beam in any way and so the symmetry of the cavity puts the waist in between the plane mirrors.

Mode-matching is generally controlled using a two lens telescope, one for each degree of freedom (position and size). As with alignment, mode-matching can be tuned while observing the transmission spectrum of a cavity length/frequency scan. Second order higher order modes will reduce in amplitude while the fundamental mode will grow.

To design the telescope, usually lenses of fixed and standard focal length are used — so the position of the available lenses is what needs to be determined. It is not ideal to allow these lenses to be positioned downstream of the steering mirrors because moving a beam using the steering mirrors with a lens in the path is likely to cause deformations and astigmatism in the beam. A method which was developed during this PhD to solve the mode-matching telescope is described in the next section.

2.4 Modelling Gaussian beams

When modelling a Gaussian laser beam propagation, the required information is the waist size and position ω_0 and z_0 , respectively, and the refractive index(es) of the medium(s) it will propagate through. With those, the beam radius $\omega(z)$ and radius of curvature $R(z)$ at any propagation distance z (measured from the waist) can be calculated. The complex beam parameters (also known as the q-parameters) are defined as [78]

$$q(z) = z + iz_R \quad (2.56)$$

$$\frac{1}{q(z)} = \frac{1}{R(z)} - i \frac{\lambda}{\pi n \omega(z)^2}. \quad (2.57)$$

To calculate the outgoing (subscript 2) light ray from the input (subscript 1) light ray the ABCD matrices (which are optical component operator matrices) are used; they are defined as [67, 78]

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} x_1 \\ \theta_1 \end{pmatrix} \quad (2.58)$$

where x_i is the distance of the ray from the optical axis and θ_i the divergence angle to the optical axis. Among the many component specific matrices, the most exploited in this thesis are:

— propagation through medium of refractive index n a distance d

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 & \frac{d}{n} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ \theta_1 \end{pmatrix}; \quad (2.59)$$

— propagation through a thin lens⁴ of focal length f

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ \theta_1 \end{pmatrix}; \quad (2.60)$$

— reflection from a mirror of radius of curvature R_{oC}

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -\frac{2}{R_{oC}} & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ \theta_1 \end{pmatrix}; \quad (2.61)$$

— propagation from one medium to another, with refractive indices n_1 and n_2 , respectively

$$\begin{pmatrix} x_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & \frac{n_1}{n_2} \end{pmatrix} \begin{pmatrix} x_1 \\ \theta_1 \end{pmatrix}. \quad (2.62)$$

⁴The thin lens approximation means there is no need to account for the change in refractive index; Eq. 2.60 is only valid for this approximation. This will be the assumption for all cases in this thesis.

These ABCD matrices can be used to tune the Gaussian beam q-parameters [78]

$$q_2 = \frac{Aq_1 + B}{Cq_1 + D}, \quad (2.63)$$

which facilitates modelling the entire optical system. It is possible therefore to rewrite Eqs. 2.59 - 2.62 as

$$q_{2prop} = q_1 + \frac{d}{n}, \quad (2.64)$$

$$q_{2lens} = \frac{q_1 f}{f - q_1}, \quad (2.65)$$

$$q_{2mirror} = \frac{q_1 RoC}{RoC - 2q_1}, \quad (2.66)$$

$$q_{2med} = \frac{q_1 n_2}{n_1}. \quad (2.67)$$

2.4.1 Mode-matching telescope solution

The principle behind the method developed in this thesis to solve mode-matching telescopes will be described here. It relies on propagating the unknown lens positions through the otherwise known optical layout of the system.

The first step is to identify the cavity waist ω_{cav} — with this, the cavity can then be considered as a *black box*. All we need to know is the required waist and position for that black box. It is worth saving this information if the rest of the layout changes, for example. There are four q-parameters to be computed when tracing the beam around the cavity of Fig. 2.12: q_{cav} describes the beam at the input; q_1 at the concave (in this case) end mirror after propagating the cavity length, L ; q_2 upon reflection from the end mirror and finally q_3 after propagating L back to the input.




Figure 2.12: q-parameters around a linear cavity. Mode-matching is the condition that $q_3 = q_{cav}$. RoC_{PRM} and RoC_{ETM} are the radius of curvatures of the power recycling mirror and the ETM, respectively.

In this example the input mirror (called the power recycling mirror, PRM) is flat i.e. $RoC_{PRM} = \infty$. Because of this, we know this is where the waist must be located (see Fig. 2.11) and it can be attributed to the starting q-parameter.

The q-parameters follow from one another as such⁵

$$\frac{1}{q_{cav}} = \frac{1}{RoC_{PRM}} - i \frac{\lambda}{\pi \omega_{cav}^2} \quad (2.68)$$

$$q_1 = q_{cav} + L \quad (2.69)$$

$$q_2 = \frac{q_1 RoC_{ETM}}{RoC_{ETM} - 2q_1} \quad (2.70)$$

$$q_3 = q_2 + L \quad (2.71)$$

$$q_3 = q_{cav} \quad (2.72)$$

With the values for RoC_{PRM} , L and RoC_{ETM} this can be solved to provide the cavity waist q-parameter.

Alternatively, from Eqs. 2.32 and 2.34 (where we substitute $R(z)$ for RoC_{ETM} and z for L), the cavity waist can be calculated as

$$\omega_0 = \sqrt{\frac{\lambda \sqrt{L(RoC_{ETM} - L)}}{\pi}} = \sqrt{\frac{\lambda \sqrt{g_{cav} RoC_{ETM} L}}{\pi}} \quad (2.73)$$

from which we generate the cavity z_R and substitute this into Eq. 2.56 where $z = 0$.

From here, the telescope solutions can be calculated provided a waist upstream of the cavity is known. If a waist is not known then a beam characterisation, such as the knife-edge method [93] must be carried out — the process used to characterise lasers as part of this thesis is given in Appendix 3.3.1. This example assumes the laser is output directly from the emitter with a known waist ω_0 (q-parameter q_0) at a known position inside the laser head, y — see Fig. 2.13.




Figure 2.13: Depiction of the telescope model. The laser on the left is providing the beam from a known waist within the laser. The beam propagates through two lenses and the input mirror of the cavity where it must be a waist of the correct size for the cavity.

Fig. 2.13 shows the telescope model being solved. Red spaces are optional lens-free zones; they ensure space for other optics which it may not be preferable for the lenses to be amongst. n_1 is assumed to be 1 for the entire model, the only refractive index mentioned is that of the mirror, which has refractive index n_2 and thickness

⁵Assuming $n = 1$.

m. Following the q -parameters of Fig. 2.13,

$$q_0 = i \frac{\pi \omega_0^2}{\lambda} \quad ; \quad q_1 = q_0 + y + s_1 + x_1 \quad (2.74)$$

$$q_2 = \frac{q_1 f_1}{f_1 - q_1} \quad ; \quad q_3 = q_2 + x_2 \quad (2.75)$$

$$q_4 = \frac{q_3 f_2}{f_2 - q_3} \quad ; \quad q_5 = q_4 + (d - x_1 - x_2) + s_2 \quad (2.76)$$

$$q_6 = q_5 n_2 \quad ; \quad q_7 = q_6 + \frac{m}{n_2} \quad (2.77)$$

$$q_8 = \frac{q_7}{n_2} \quad (2.78)$$

Now we can solve the problem using the previously calculated q_{cav} as the desired output. Since there are two unknowns x_1 and x_2 , it is necessary to separate the real and imaginary components of the final condition in order to have a pair of simultaneous equations:

$$re[q_8] = re[q_{cav}] \quad (2.79)$$

$$im[q_8] = im[q_{cav}] \quad (2.80)$$

This process is carried out for each combination of the available lenses.

Additional features can be added, such as calculating the beam radius at each point in order to estimate the clipping loss with each optic.

2.5 The power-recycled Michelson interferometer

Combining the interferometer with a cavity is not a new concept. Gravitational wave detectors like LIGO and Virgo use several cavities in conjunction with a Michelson interferometer. The configuration discussed here is the power-recycled Michelson interferometer, where an additional mirror (the power-recycling mirror, PRM) is placed at the symmetric port, between the source and the beam splitter.




Figure 2.14: The power-recycled Michelson interferometer. The additional power-recycling mirror directs the symmetric port light back towards the beam splitter. In the dark fringe condition, the Michelson interferometer forms a cavity with the power-recycling mirror, enhancing the power circulating within.

This additional mirror transforms the interferometer into a cavity of sorts. The fringe condition at the anti-symmetric port determines the *Finesse* of the cavity since the interferometer now behaves like a mirror with a variable reflectivity. At a bright fringe, the interferometer directs all the light away from the power-recycling mirror and out of the anti-symmetric port — this is analogous to having fully transmissive mirrors in the cavity, the light will pass right through (when resonant with the common arm length). In the dark fringe condition, the interferometer directs all the light back towards the power-recycling mirror and none leaves via the anti-symmetric port — this is then analogous to having a fully reflective end mirror in the cavity (assuming no losses). Thus, at the dark fringe, the power-recycled Michelson interferometer can be treated as a linear cavity with cavity length equal to the common arm length (\bar{L} in Eq. 1.32) plus the distance from the power-recycling mirror to the beam splitter. The ‘reflectivity’ of the Michelson interferometer is determined by how dark the dark fringe can be made⁶, with loss to the anti-symmetric port accounting for its transmissivity. To fully exploit the power-recycled Michelson interferometer configuration, the darker the dark fringe, the better. If excess light cannot be stopped leaking out of the anti-symmetric port, it is equivalent to a high loss cavity.

The overcoupled configuration (power-recycling mirror of lower reflectivity than the dark fringe Michelson interferometer ‘reflectivity’) will generate the largest available power build up of the possible configurations — see Secs 2.3.1 and 2.3.3. With high quality mirrors and appropriate reflectivity choices, internal power can reach orders of magnitude higher than the injection power.

The inclusion of the power-recycling mirror coupled with holding the Michelson interferometer at the dark fringe is a very efficient method of increasing the power

⁶The arm end mirrors are assumed to be almost fully reflective.

at the beam splitter, which is the strategy for improving the shot-noise-limited ASD (see Eq. 1.36). This is a primary reason for operating at the dark fringe. Another reason is to limit the power falling on the photodetectors — low noise, high bandwidth photodetectors typically have small power thresholds and risk being saturated when the dark fringe is not well maintained.

2.6 Control loops

Given the sensitivity of the devices described so far, it is no surprise that feedback control schemes are required to hold them at their operating points. If the mirrors of a Michelson interferometer or a cavity were free to move, or the laser free to drift in frequency or output power, then nothing much could be clearly determined by the output of the system — any scientific signals would be buried by noise. Different noise sources manifest at different frequencies, shot noise (high frequency) and radiation pressure (low frequency) were introduced in Secs 1.4.2 and 1.4.3, respectively. A more intuitive noise source for a table-top optical experiment is seismic activity which covers a broadband of very low frequencies. Higher frequency noise sources include things like the electrical outlet AC frequency (50 Hz for the UK), and mechanical resonances of components which can be several hundreds of Hz up to many kHz. This is all to point out that it is not as simple as holding a system at a single DC operational point, the operating point is the desired output for all frequencies. The loop must behave appropriately to noises at different frequencies across its operational bandwidth.

A feedback loop [76, 79] as shown in Fig. 2.15, is a system where the real output $o(\omega)$ is observed by some sensor $S(\omega)$ which provides a feedback signal $y(\omega)$ which is compared to the desired output, the set-point $z(\omega)$. If $y(\omega) \neq z(\omega)$, an error signal $e(\omega)$ is generated and used to provide corrective measures. In the case of a negative feedback loop⁷, the error is $e(\omega) = z(\omega) - y(\omega)$. This error is then passed through a controller, $H(\omega)$, where appropriate amplifications are applied by a servo to transform $e(\omega)$ into $c(\omega)$, the control signal. The actuator $A(\omega)$ takes $c(\omega)$ and acts on the plant $P(\omega)$, the component ultimately under control.

⁷So called because the output is subtracted from the set-point.




Figure 2.15: Simple negative feedback loop. $z(\omega)$ is the set-point, the desired output. $o(\omega)$ is the actual output. The sensor $S(\omega)$ measures the output and provides the feedback signal $y(\omega)$. $e(\omega)$ is the error signal, the difference between $z(\omega)$ and $y(\omega)$, it shows how far from the desired output the system is. The controller $H(\omega)$ takes in the error and amplifies it appropriately, outputting the control signal $c(\omega)$ to the actuator $A(\omega)$ which drives the plant $P(\omega)$ is the device under control.

The loop is divided into parts, each with their own contribution to the whole:

- $P(\omega)$ — The Plant is what is under control/being monitored, it outputs $o(\omega)$.
- $S(\omega)$ — The Sensor is the component used to monitor $o(\omega)$ and provide the feedback signal $y(\omega)$.
- $H(\omega)$ — The Controller (or servo) transforms the error signal, $e(\omega)$ into the control signal $c(\omega)$.
- $A(\omega)$ — The Actuator which is acting on the Plant, it is driven by c .

Each component has a transfer function defined as the output over the input, which has a frequency dependency. The main properties of the transfer function are the gain (some multiple) and the phase (some delay) of the input signal, which are visually displayed by a Bode diagram [76, 79, 94].

Given the definition of the transfer function, signals around the loop are multiplied by the device they pass through, i.e. $c(\omega) = e(\omega)H(\omega)$. The *open loop* transfer function, $G(\omega) = H(\omega)A(\omega)P(\omega)S(\omega)$ is what the loop components collectively do to the input signal. The *closed loop* transfer function, when feedback is included, can be defined by tracing signals around the loop:

$$e(\omega) = z(\omega) - y(\omega), \quad (2.81)$$

$$c(\omega) = e(\omega)H(\omega), \quad (2.82)$$

$$y(\omega) = c(\omega)A(\omega)P(\omega)S(\omega). \quad (2.83)$$

Combining all of these,

$$y(\omega) = e(\omega)G(\omega) \quad (2.84)$$

$$y(\omega) = [z(\omega) - y(\omega)]G(\omega) \quad (2.85)$$

$$y(\omega) = \frac{z(\omega)G(\omega)}{1 + G(\omega)} \quad (2.86)$$

and so the closed loop transfer function is,

$$\frac{y(\omega)}{z(\omega)} = \frac{G(\omega)}{1 + G(\omega)}. \quad (2.87)$$

Where open loop gain $|G(\omega)| \gg 1$, then $y(\omega) \approx z(\omega)$ i.e. the measured output follows the input — this condition is assumed for the bandwidth of the loop.

A loop is unstable if an input creates a ‘diverging’ output. In Eq. 2.87, this will happen when $|G(\omega)| \rightarrow -1$. In terms of a sine wave, multiplying the amplitude at each point by -1 is mathematically equivalent to *delaying* it by $\pm 180^\circ$. Input signals with this phase delay will cause a runoff towards instability if their amplitude is not properly attenuated. Loops remain stable by ensuring the frequency at which $\arg[G(\omega_{180})] = \pm 180^\circ$ is first approached, the gain for that ω_{180} is much less than 1 i.e. $|G(\omega_{180})| \ll 1$. That way, each time it passes through the loop its amplitude is reduced and it cannot build to make the loop unstable.

The unity gain frequency is the frequency at which $|G(\omega)| = 1$ — meaning signals will pass through unaltered in amplitude. $G(\omega)$ must be designed to ensure this frequency is a safe distance from $\arg[G(\omega)] = \pm 180^\circ$ — this is called the phase margin⁸.

The ideal case of Fig. 2.15 is modified in Fig. 2.16 to a more realistic case. Here the plant’s output is corrupted by noise $N(\omega)$, the noise of the plant itself. Depending on what the plant is and/or which property of the output is under control determines the type of noise under consideration. A laser for example has noise in the power, frequency and polarisation of its outputted field. They may not be strongly coupled, they could all have a different actuator, and thus would require a dedicated feedback loop each for suppressing the respective noises. In the design of the loop the other components are tested and proven to have negligible noise compared to the plant — as described in Sec. 2.6.2.

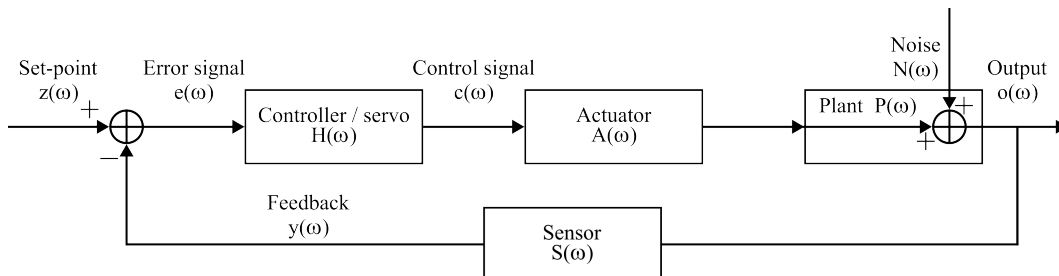


Figure 2.16: Simple negative feedback loop as in Fig. 2.15 but with a noise input $N(\omega)$. It is placed within plant because it is assumed that the components of the loop have negligible noise compared to that of the plant.

⁸There is no ‘right’ value for the phase margin, but in general it should be kept as far from $\pm 180^\circ$ as possible. Finding a working value could come down to trial and error for a specific loop.

Now using the same approach as the loop equations above, the feedback signal looks like:

$$y(\omega) = \frac{z(\omega)G(\omega)}{1 + G(\omega)} + \frac{N(\omega)S(\omega)}{1 + G(\omega)} \quad (2.88)$$

and if we assume the open loop gain is very large $|G(\omega)| \gg 1$ and $|G(\omega)| \gg |S(\omega)|$ then this simplifies to

$$y(\omega) = z(\omega) + \frac{N(\omega)}{1 + G(\omega)} \quad (2.89)$$

where we can see how the closed loop suppresses the noise of the plant.

2.6.1 Transfer functions and their measurements

In this section frequency is described as $\nu = \omega/2\pi$ because of the way filter shapes are typically defined, by their gain change with frequency, not angular frequency.

One way to describe the general trend of a transfer function is by using a DC gain, poles and zeros [76, 79, 94]. A single pole creates a *low-pass* filter which has the characteristic transfer function

$$T_{LP}(\nu) = \frac{k}{1 + i\nu/\nu_{\text{cutoff}}} \quad (2.90)$$

where k is the DC gain⁹ and ν_{cutoff} is the pole or *cut-off* frequency. High-pass filters tend to be drawn with this characteristic transfer function:

$$T_{HP}(\nu) = \frac{ki\nu}{i\nu + \nu_{\text{cutoff}}} \quad (2.91)$$

which contains a zero (where $\nu = 0$) and a pole (where $\nu = \nu_{\text{cutoff}}$).

Low-/high-pass refers to the shape of the filter. In a low-pass filter signals of $\nu < \nu_{\text{cutoff}}$ are passed while $\nu > \nu_{\text{cutoff}}$ are increasingly attenuated — see Fig. 2.17a. On the other side, high-pass filters increasingly attenuate signals where $\nu < \nu_{\text{cutoff}}$ while passing signals where $\nu > \nu_{\text{cutoff}}$ — see Fig. 2.17b. .

Eqs. 2.90 and 2.91 have a phase relationship — their gain and phase are calculated as is custom with complex numbers i.e. gain is $|T(\nu)|$ and phase is $\arg[T(\nu)]$. The filter's frequency dependent gain G is generally shown in dB, where

$$G_{\text{dB}}(\nu) = 20 \log_{10}(G(\nu)).$$

Filters are described by their frequency dependency — for example, a single pole low-pass is described as having a $1/\nu$ shape¹⁰. This is because for every order of magnitude of frequency increase (called a decade), gain is reduced by an order of magnitude, or, equivalently, by 20 dB. The opposite is true for the high-pass: gain

⁹The DC gain is a 'flat' gain factor across all frequencies.

¹⁰This would be pronounced "one by f" or "one over f".

increases by an order of magnitude or 20 dB for every decade in frequency — see Figs. 2.17a and 2.17b.

The low-pass (LP) filter introduces a 90° phase delay, where phase goes through -45° at ν_{cutoff} , see Fig. 2.17c. For the high-pass (HP), phase begins at $+90^\circ$ and is delayed to 0° going through $+45^\circ$ at ν_{cutoff} , see Fig. 2.17d.

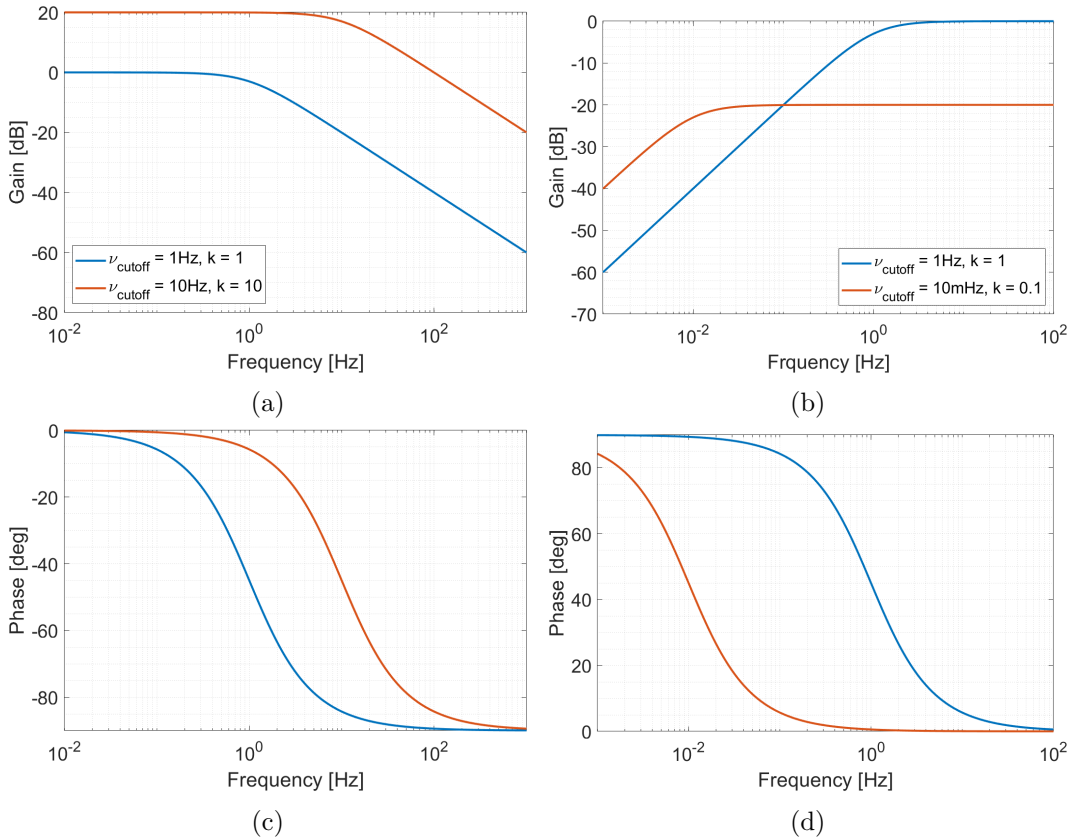


Figure 2.17: Bode plots showing filters with different cut-off frequencies and DC gains. In a traditional Bode plot, the top plot shows gain against frequency while the lower one shows the phase against frequency. (a) and (c) represent the gain and phase of a pole, respectively. (b) and (d) represent the gain and phase of a zero, respectively.

When filters are cascaded, their transfer functions are multiplied — in the log space of dB, this is equivalent to summing them. Fig. 2.18 shows examples of some low-pass and high-pass filter combinations.

A combination of two low-pass filters creates two poles which results in a $1/\nu^2$ gain trend¹¹ and an additional 90° phase delay — bringing it to a total of -180° . Conversely, two zeros will create a ν^2 gain trend, and the phase will begin at $+180^\circ$.

These kinds of basic transfer functions can be generated within electronic servos or dedicated analogue filters. For a plant of unknown transfer function, it is typically measured by injecting an excitation signal as an input to a component and measuring

¹¹This would be pronounced “one by f squared” or “one over f squared”.

the output. Dividing the output by the input provides the information about what the component does to signals of that frequency.

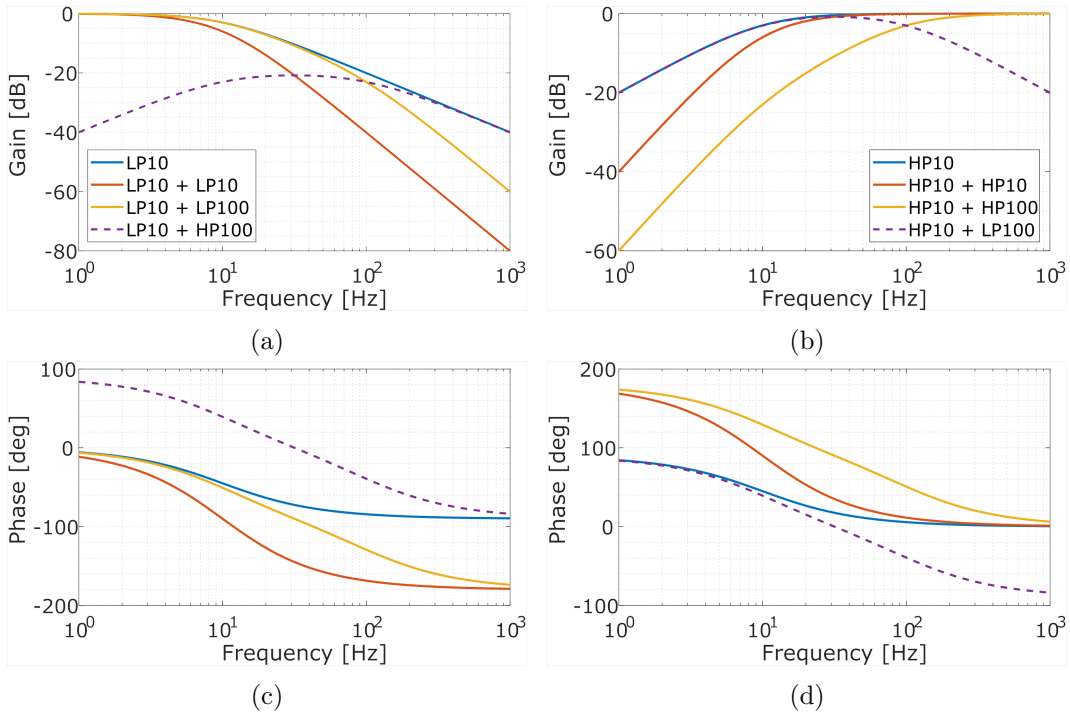


Figure 2.18: Bode plots of cascaded filters. LP10 is a low-pass where $\nu_{\text{cutoff}} = 10$ Hz. LP100 is a low-pass where $\nu_{\text{cutoff}} = 100$ Hz. HP10 is a high-pass where $\nu_{\text{cutoff}} = 10$ Hz. HP100 is a high-pass where $\nu_{\text{cutoff}} = 100$ Hz. (a) and (b) show the gains of the listed combinations of filters, (c) and (d) show the phase of those combinations.

Fig. 2.19 shows a method of measuring the open loop transfer function.

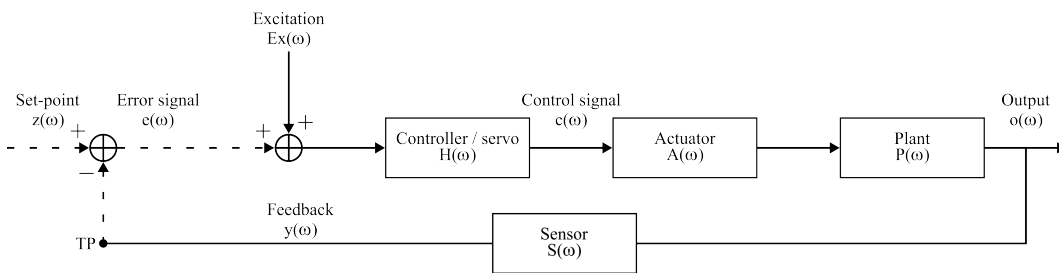


Figure 2.19: Measuring the open loop transfer function of a simple negative feedback loop; the dashed lines represent signal paths which are no longer in use/connected, i.e. the loop is open. An excitation signal $Ex(\omega)$, usually in the form of a ‘swept sine’ signal, is injected and the test point TP, which is just the sensor output $y(\omega)$ are used to measure the open loop transfer function.

From Fig. 2.19, we can see that the signal measured at TP is the output of the sensor $y(\omega)$ i.e.

$$TP = y(\omega) = Ex(\omega)G(\omega) \quad (2.92)$$

and so $G(\omega)$ is measured as

$$G(\omega) = \frac{TP}{Ex(\omega)}. \quad (2.93)$$

Where $Ex(\omega)$ is some well defined swept sine signal. With Eq. 2.93 the closed loop transfer function can be calculated according to Eq. 2.87.

2.6.2 Noise, noise budget and noise projection

The word noise has been used quite a lot so far — shot noise and radiation pressure noise, for example. But the general idea is that noise is an unwanted input to a system/detector which must be suppressed to prevent it compromising the detector's sensitivity. The worst case is that noise disables the detector entirely; pushing an optical cavity off resonance for example.

Noise in a loop can be added as a separate input or as has been done in this thesis so far, by including it as part of one of the main inputs — the set-point for example. Or if wishing to specify noise at the output of the plant, it could be added at the output of the plant (Fig. 2.20b shows this). Noise is important to be aware of since it may limit the loop efficiency, bandwidth or prevent stability. If there is a noise spike past the unity gain frequency which occasionally peaks above $|G(\omega)| = 1$, this can drive the loop to instability. A power spectrum of various points of the loop can assist in finding these features and specific filters can be added to target the problems.

A noise budget is a display of the major loop component noises in order to assess their respective contributions. An important measurement to include is the input noise of whatever analyser is being used for the measurement. The analyser input noise is the 'noise floor' of the measurement. If a component is quieter than the analyser then it should be amplified first¹². Measuring the components individually enables the loop designer to understand where the noise is coming from and which component is likely to be responsible for any problems. All the relevant component noises can then be summed in quadrature to get an idea of the total expected noise of the loop. For example, $N_{tot} = \sqrt{N_A^2 + N_{PDJ}^2}$ where N_{tot} is the total noise estimate, N_A is the analyser input noise and N_{PDJ} is the photodetector Johnson or thermal noise ($N_{PDJ} = \sqrt{4k_B T R}$ where k_B is Boltzmann's constant, T is the temperature in Kelvin and R the resistance [68]). If attempting to measure the power noise of a laser, N_{RIN} and at some frequencies $N_{RIN} < N_{tot}$, the power noise cannot be accurately measured at those frequencies.

A noise projection uses the measured open loop gain, $G(\omega)$, to model what the loop *should* do to the noise of the component in the loop which is under control, the plant¹³. Multiplying the plant's noise spectrum by $1/(1 + G(\omega))$ gives the projected

¹²And of course the noise of the amplifier itself therefore needs to be measured as well.

¹³Up until now, 'the plant' has been used in reference to the transfer function of the device, but

plant noise when the loop is closed. This can also be done to specific signals, the error signal for example.

2.6.3 RIN suppression example

Lasers have an inherent intensity noise, commonly referred to as the RIN — relative intensity noise. It is the intensity noise spectrum normalised to the DC output. The suppression of this RIN is a straightforward example of a feedback loop. A test setup undertaken during this PhD work is in Fig. 2.20. Fig. 2.20a shows the optical layout — how it was arranged on the bench, Fig. 2.20b shows how the feedback loop can be drawn. Fig. 2.20a does not have a noise input, it is inherent in the laser but for the control loop approach, the noise is considered as a separate input, added at the plant output.

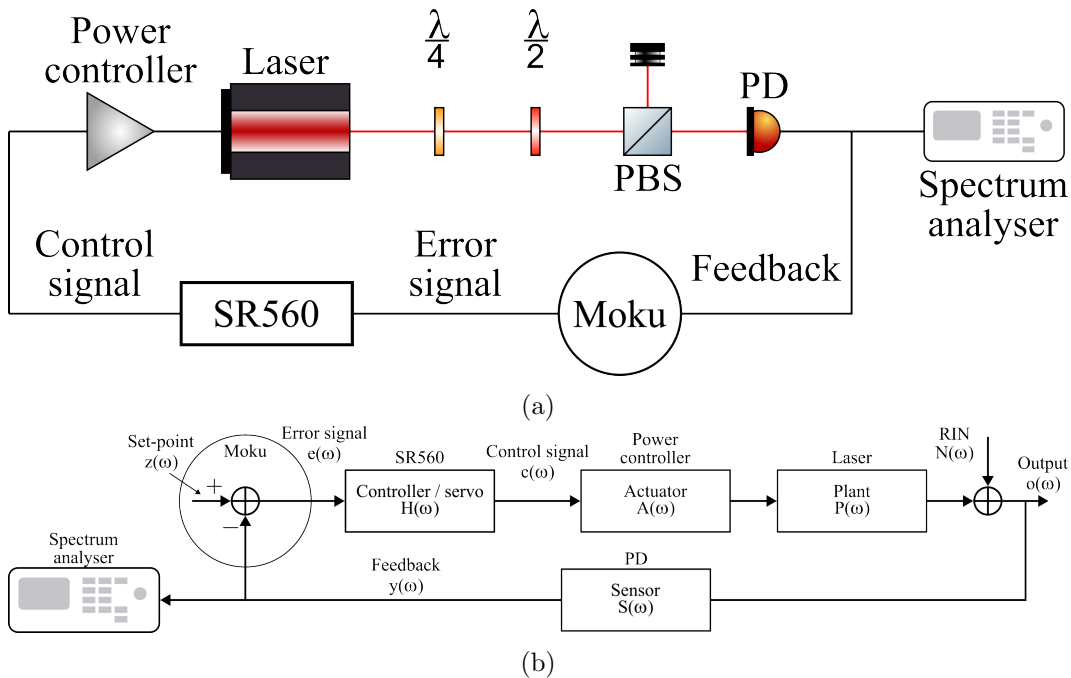


Figure 2.20: Layout for RIN example. (a): Optical representation, this shows the components as they were organised on the bench. The Moku [95] was there to add an offset/set-point and flip the sign of the sensor signal. The SR560 is an analogue filter servo. (b): Control loop representation. This is how the feedback loop operated. The servo is the filter and gain applied by the SR560.

In this test setup, the Plant is a Coherent Mephisto laser [96]. The output power of the laser is set by the drive current, which can be controlled via a voltage input. The controller/servo in this example is an analogue filter servo, the SR560 [97] and the sensor a Thorlabs PDA20CS2 photo detector [98]. The laser was set to output 0.5 W as its optimum working point, and the power falling on the photodetector was

it can also simply refer to the device itself. If a laser's frequency is being controlled by a loop, then the laser is 'the plant' and so this plant has an associated noise spectrum — frequency noise in this example.

reduced using waveplates and a polarising beam splitter (PBS). The actual power on the photodetector was not measured, instead the voltage from the photodetector was set (using the waveplates) to 500 mV. The sensor output was then fed through a Moku:Lab [95] which applied a -1 gain and added a 500 mV set-point. This is where the error signal is generated, subtracting the sensor output from the set-point.

A Rohde & Schwarz spectrum analyzer was used to take a transfer function of the plant in the frequency range 1 Hz - 10 kHz.

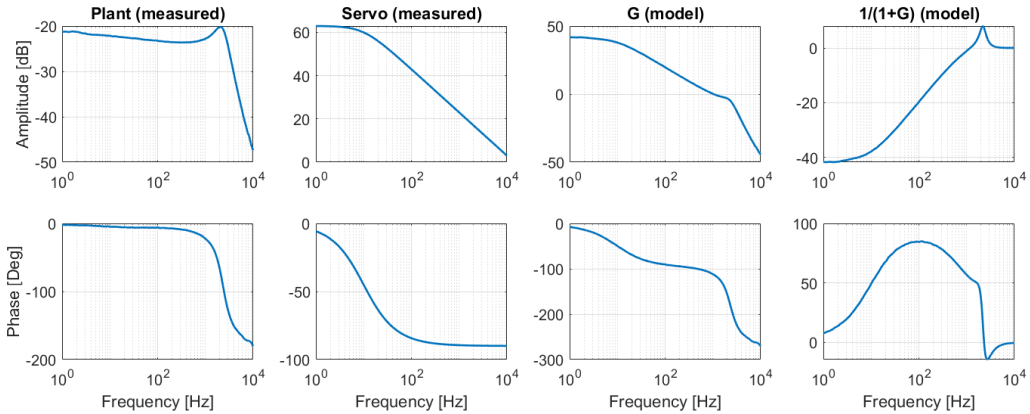


Figure 2.21: Measured transfer functions of the plant and servo, modelled open and closed loop transfer functions.

The Plant of Fig. 2.21 has the transfer function shape of a double pole low-pass filter — the gain is reasonably flat out to ~ 2 kHz after which point it sharply reduces the amplitude at a rate of 40 dB/decade, and there is a -180° phase delay. A single pole low-pass falls at a rate of 20 dB/decade and introduces a phase delay of -90° . The shape is due to the laser components' inability to respond to current fluctuations above around 2 kHz. The small 'gain peak' at 2 kHz is likely due to a circuit resonance.

The servo filter was designed to establish a unity gain frequency of 1 kHz and a safe phase margin of 70° . This is achieved using a single pole low-pass filter with a DC gain of 63 dB and cut-off frequency 10 Hz. After 10 Hz, this will lose 20 dB per decade, so by 1 kHz the remaining gain is 23 dB. It will also have lost 90° phase. Combined with the plant (G (model) of Fig. 2.21), we achieve our 0 dB gain at 1 kHz and phase of -110° , giving us a 70° margin. The loop suppression can then be modelled — shown in the $1/(1+G)$ (model) plots, Fig. 2.21. This effectively shows what signals will be multiplied by when the loop is closed. Signals below 1 kHz will be suppressed since they are multiplied by negative dB gain (< 1 linear gain).

RIN noise budget and result

Fig. 2.22 shows the noise budget for the laser intensity noise suppression project.

Measured spectra are solid lines, dashed lines are calculated models.

1. ‘Analyser input noise’ was measured with a BNC terminator attached to the input.
2. ‘Photocurrent noise’ is equivalent to photon shot noise; it is the expected voltage noise of the photodetector as a result of the shot noise in the laser power. It was calculated as $\sqrt{2eIG_{PD}}$ where e is the charge of the electron, I is the current at the photodetector, calculated by V/G_{PD} . V is the voltage output by the photodetector and G_{PD} is the transimpedance gain of the photodetector (at the 0 dB gain setting, $G_{PD} = 1.51 \text{ k}\Omega$).
3. ‘Johnson noise’ is the thermal noise of the resistor, calculated as $\sqrt{4k_B T G_{PD}}$ where k_B is Boltzmann’s constant, T is the temperature in Kelvin [68].
4. ‘SR560 + PD’ is the combined noise of the SR560 at 40 dB DC gain and the photodetector when the laser was off, this is the photodetector powered with the sensor uncovered in the ambient light of the lab. These components were combined because their respective noises were lower than the analyser input noise; only with the 40 dB gain on the SR560 could either be seen and even then the photodetector noise was only marginally higher than the SR560 by itself. They were left as a combined noise source, the data was divided by 100 (40 dB) to recover the appropriate level.
5. The ‘Total loop noise’ is the quadrature sum of these listed noise sources.
6. The ‘Free running laser’ is the laser noise when the loop is open, i.e. there is no feedback.
7. ‘Projection’ is then this free running laser data multiplied by $1/(1 + G(\omega))$ (shown in Fig. 2.21), with the total loop noise added. This presents the projected outcome of the loop’s suppression of the plant noise, the RIN.
8. ‘Loop closed laser’ is where feedback is engaged, the suppression is clear and even just surpasses the projection. Between $\sim 3 - 60 \text{ Hz}$, at least an order of magnitude suppression was achieved.

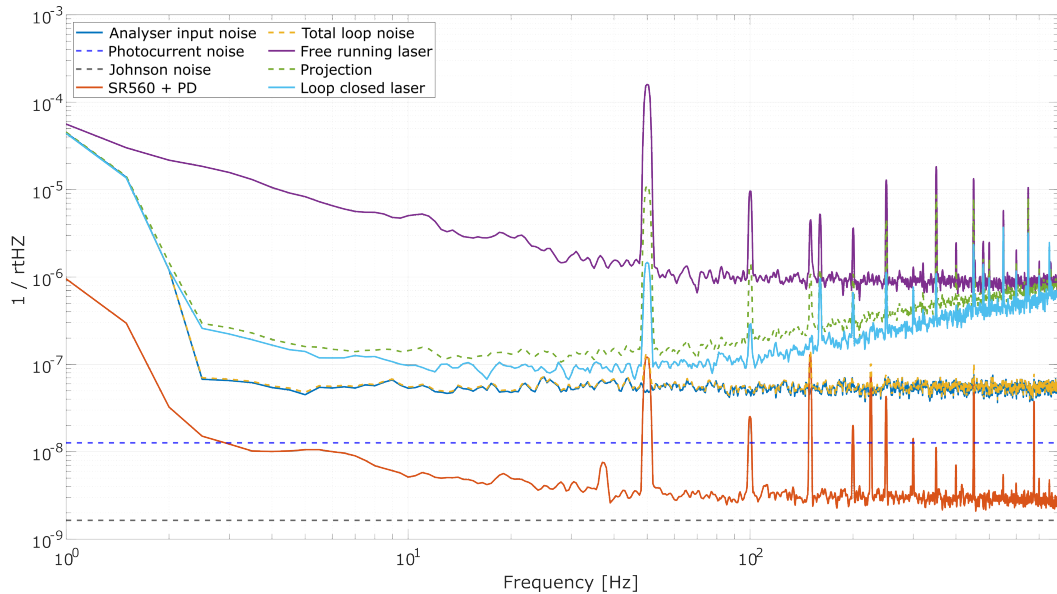


Figure 2.22: RIN Noise budget and suppression. The noises present in the system are the analyser input, photocurrent or shot noise, Johnson thermal noise and the photodetector measured noise with no laser. These were summed in quadrature to calculate the noise projection. The free running RIN is the measured laser intensity noise with no feedback, the loop closed RIN is the laser intensity noise with feedback. Between $\sim 3 - 60$ Hz, at least an order of magnitude suppression was achieved.

2.7 Signal read-out with interferometers

When some phenomenon drives a DARM length change in the Michelson interferometer, phase modulated sidebands are generated on the carrier as described in 2.1.1 [67, 99]. The carrier is at frequency ν_c and the phenomenon has frequency ν_p , leading to signal sidebands of frequency $\nu_{sig} = \nu_c \pm \nu_p$. These sidebands constructively interfere at the beam splitter and make it to the photodetector at the anti-symmetric port. However, for the 1064 nm IR lasers at QUEST (and with the large gravitational wave detectors), the carrier field frequency is $\nu_c \approx 3 \cdot 10^{14}$ Hz. QUEST’s bandwidth extends to 250 MHz, this would put the lower signal sidebands at $\nu_{sig,low} = 3 \cdot 10^{14}$ Hz $- 250 \cdot 10^6$ Hz $\approx 3 \cdot 10^{14}$ Hz. There are no devices capable of measuring the changes to individual cycles at anything close to that frequency, yet. Further, the phase of an individual field is also not something which can currently be detected or monitored; it must be measured against some reference. This change in phase may actually be the origin of the signal we are looking for in the first place, see Sec. 2.1.1; so it is important to find a way to measure it.

Conveniently, a solution for measuring both frequency and phase relies on the same technique, signal mixing and demodulation. The process uses a stable field to down-convert the high frequency signal carrying field to a detectable low frequency, and by doing so, also provides a reference to detect changes to phase. This section gives a description of how mixing and demodulation can be used to find phase information of high frequency fields and *read out* the signal. To simplify things, let’s first consider

how the displacement of a single mirror could be detected.




Figure 2.23: Mirror displacement detection. (a): Probe beam with mirror displacement. The beam reflection is routed to a photodetector; there is no way to detect the mirror movement with this method. (b): Probe beam mixed with local oscillator beam. The beams are superimposed at the lower beam splitter and together shine on the photodetector for analysis. If these lasers are of a different frequency, the difference frequency *beat* is used and this would be heterodyne detection. If the lasers have the same frequency (possibly even originating from the same laser), then this would be homodyne detection.

A probe beam has electric field

$$E_{pro} = A_{pro}e^{i2\pi\nu_{pro}t}, \quad (2.94)$$

where A_{pro} and ν_{pro} are the amplitude and frequency of the beam. We can use this beam to measure the displacement of a mirror, by identifying the phase changes induced by the mirror moving.

The mirror displacement Δx_m will generate a phase change in the field as

$$\Delta\phi_m = \frac{2\pi \cdot 2\Delta x_m}{\lambda}. \quad (2.95)$$

If we can measure $\Delta\phi_m$, we can calculate Δx_m . The probe field is modified and becomes

$$E_{pro} = A_{pro}e^{i(2\pi\nu_{pro}t + \Delta\phi_m)}. \quad (2.96)$$

This is the field shining on the photodetector in Fig. 2.23a. To extract $\Delta\phi_m$, we use signal mixing via homodyne/heterodyne detection. They work in an identical way, the difference being that homodyne means the fields being mixed have the same frequency and heterodyne is where the fields have a different frequency.

The non-probe field is assumed to be a stable reference against which fluctuations in the probe field can be measured. This non-probe field is commonly referred to as the local oscillator (LO),

$$E_{LO} = A_{LO}e^{i(2\pi\nu_{LO}t)}, \quad (2.97)$$

where A_{LO} and ν_{LO} are the amplitude and frequency of the local oscillator beam, re-

spectively. At the lower beam splitter of Fig. 2.23b, E_{pro} and E_{LO} are superimposed on one another. At the photodetector of Fig. 2.23b we have the field

$$E_{pd} = E_{pro} + E_{LO} = A_{pro}e^{i(2\pi\nu_{pro}t + \Delta\phi_m)} + A_{LO}e^{i(2\pi\nu_{LO}t)} \quad (2.98)$$

and then photodetectors measure the power,

$$P_{pd} = |E_{pd}|^2 = |A_{pro}|^2 + |A_{LO}|^2 + 2A_{pro}A_{LO} \cos(2\pi(\nu_{pro} - \nu_{LO})t + \Delta\phi_m). \quad (2.99)$$

The first two terms are DC, that is, they do not oscillate. The third term oscillates at the difference frequency between the probe and local oscillator beams, and it is a function of $\Delta\phi_m$, the interesting information.

2.7.1 Heterodyne detection

If using heterodyne sensing the probe beam and the local oscillator have a different frequency i.e. $\nu_{pro} \neq \nu_{LO}$. The difference frequency is the heterodyne frequency, ν_h , which is commonly called the *beat note*. This frequency we can detect by design,

$$\nu_h = \nu_{pro} - \nu_{LO}. \quad (2.100)$$

The DC components of Eq. 2.99 are not relevant because they do not contain the information we are after — so if we pass the photodetector signal through a high-pass filter, either by AC coupling the photodetector or otherwise, we can isolate the oscillating part which carries the $\Delta\phi_m$ information

$$P_{pd-ac,heterodyne} = 2A_{pro}A_{LO} \cos(2\pi\nu_h t + \Delta\phi_m). \quad (2.101)$$

Using a trigonometric identity, this becomes

$$P_{pd-ac,heterodyne} = 2A_{pro}A_{LO} [\cos(2\pi\nu_h t) \cos(\Delta\phi_m) - \sin(2\pi\nu_h t) \sin(\Delta\phi_m)] \quad (2.102)$$

and can be separated into the in-phase (with the cos function in Eq. 2.101), ϵ_I , and quadrature ϵ_Q components,

$$\epsilon_I = A_{pro}A_{LO} \cos(\Delta\phi_m), \quad (2.103)$$

$$\epsilon_Q = A_{pro}A_{LO} \sin(\Delta\phi_m), \quad (2.104)$$

such that

$$P_{pd-ac,heterodyne} = 2\epsilon_I \cos(2\pi\nu_h t) - 2\epsilon_Q \sin(2\pi\nu_h t). \quad (2.105)$$

This is the electronic signal output by the photodetector. We can electronically multiply this by another sinusoid — a process called mixing¹⁴. When signals are multiplied in the time domain, by convolution theory, that is equivalent to adding or subtracting in the frequency domain [100]. This is an important feature to note since for the majority of this thesis, signals will be defined by their frequency and when dealing with demodulation, mixing etc., adding frequencies will be the norm.

Mixing $P_{pd-ac,heterodyne}$ with a reference signal of frequency ν_h and a tunable phase θ_{demod} will allow us to extract ϵ_I and ϵ_Q separately.

$$\begin{aligned} P_{pd-ac,heterodyne} \cdot \cos(2\pi\nu_h t + \theta_{demod}) &= 2\epsilon_I \cos(2\pi\nu_h t) \cdot \cos(2\pi\nu_h t + \theta_{demod}) - \\ &\quad 2\epsilon_Q \sin(2\pi\nu_h t) \cdot \cos(2\pi\nu_h t + \theta_{demod}) \\ &= \epsilon_I [\cos(\theta_{demod}) + \cos(4\pi\nu_h t + \theta_{demod})] - \\ &\quad \epsilon_Q [\sin(4\pi\nu_h t + \theta_{demod}) - \sin(\theta_{demod})] \end{aligned} \quad (2.106)$$

The components oscillating at $4\nu_h$ can be discarded with a low-pass filter, leaving

$$P_{pd-ac,heterodyne} = \epsilon_I \cos(\theta_{demod}) + \epsilon_Q \sin(\theta_{demod}) \quad (2.107)$$

θ_{demod} is then tuned to maximise ϵ_I or ϵ_Q (while minimising the other), giving us a direct measure of $\Delta\phi_m$.

2.7.2 Homodyne detection

If using homodyne detection the probe beam and local oscillator have the same frequency. At Eq. 2.99, $\nu_{pro} = \nu_{LO}$ and so $\nu_h = 0$. i.e.

$$P_{pd-ac,homodyne} = 2A_{pro}A_{LO} \cos(\Delta\phi_m). \quad (2.108)$$

We already have a direct measure of Δx_m since power fluctuations at the photodetector are caused by changes of $\cos(\Delta\phi_m)$, with maximum sensitivity at the mid-point of the cosine cycle. The problem with this method is the inability to distinguish between power fluctuations caused by signals or the laser output itself. Also, as will be more thoroughly described in Sec. 2.7, combining the probe beam with the local oscillator (as in the case of Fig. 2.23b) is itself a technical challenge since they generally do not take the same path.

2.7.3 For the Michelson interferometer

Now we can look at how these methods are implemented for a Michelson interferometer.

¹⁴Not to be confused with optical mixing, where laser fields are superimposed — the kind of mixing described here is electronic and is done digitally.




Figure 2.24: Three common readout techniques for the interferometer [99]. Red solid line is the carrier, blue dashed lines are signal sidebands and green dashed lines are heterodyne sidebands. (a): Heterodyne readout; this technique relies on electro-optical modulator (EOM) generated sidebands to leave the interferometer due to the Schnupp asymmetry and act as the local oscillator for the signal sidebands. (b): Homodyne readout; this technique uses a pick-off of the carrier prior to entry into the interferometer. The pick-off is then routed to the output to act as the local oscillator, it is aligned to overlap with the signal sidebands. (c): DC-readout; this technique uses a small differential-mode length difference to allow some constant amount of carrier light to leak out and act as the local oscillator for the signal sidebands.

With a heterodyne detection scheme (Fig. 2.24a), radio frequency (RF) sidebands (generally in the high kHz to low-mid MHz range) are generated on the carrier prior to it entering the interferometer. These sidebands are obtained with an electro-optical modulator which is driven by a RF signal generator. The sidebands have a frequency of ν_{het} and are of exactly the same mathematical form as the signal sidebands discussed in Sec. 2.1.1. The Schnupp asymmetry is a Michelson interferometer arm length difference on the order of cm that is intended to hold the sidebands away from their dark fringe [99], allowing them to leave the anti-symmetric port, while maintaining the dark fringe for the carrier. This then provides a local oscillator at the detection port in the form of the sidebands. They optically mix with the signal sidebands at the output and demodulation using the same RF signal generator isolates ν_p leading to signal detection.

A homodyne detection scheme (Fig. 2.24b) requires some pick-off of the carrier before it reaches the interferometer. This pick-off is then routed to the detection port outside of the interferometer and is able to act as the local oscillator. The method has issues related to the stability of the local oscillator beam in terms of position and phase. Since the beams need to spatially overlap to ensure interference, and the phase needs to be stable since the point of the local oscillator is to act as a reference, the optics used to route it need dedicated feedback control. Given these challenges, this is why homodyne is not widely used¹⁵.

¹⁵Though it will be used for A+ [101].

DC-readout

DC-readout (Fig. 2.24c) is a special case of homodyne detection. A very small ($\ll \lambda/4$) asymmetry in the interferometer arms is maintained such that some constant amount of carrier light exits the anti-symmetric port and is thus present to act as the local oscillator. This method removes the instability in ensuring the local oscillator and signal beams are overlapping since they are the same beam [99]. This readout scheme is detailed more in Chapter 4, since it is used for the experiment of this thesis.

2.8 The contrast defect

The dark fringe of a Michelson interferometer is never 100% dark in reality because of fundamental mode light which imperfectly interferes on the beam splitter and leaks out at the anti-symmetric port. This is caused by optical imperfections in the beam splitter and ETMs. The non-perfect beam splitter results in an unequal power split between the two arms meaning totally destructive interference is not possible. In addition, differential reflectivity, surface integrity and cleanness differences in the ETMs generate asymmetrical scattered light and non-homogeneous phase changes across the beam wavefronts returning to the beam splitter. At the beam splitter, this light does not destructively interfere and couples to the anti-symmetric port. Also, where very high power or a power-recycled Michelson interferometer configuration (Sec. 2.5) is used, thermal lensing can become problematic. The heat from the high power effects the optics, the beam splitter in particular; this couples light to higher order modes [48]. These sources of light make up the contrast defect (CD), a combination of non-perfectly destructing fundamental mode and higher order modes; they keeping the Michelson interferometer's output from reaching a purely dark fringe.

This contrast defect is not a feature of DARM changes and so does not contribute to the signal strength — on the contrary, contrast defect is always present and increases with power at beam splitter, contributing only to noise and thus reducing the SNR. The contrast defect can be defined as,

$$\text{CD} = \frac{P_{AS,DF}}{P_{BS,DF}} \quad (2.109)$$

where $P_{AS,DF}$ is the power leaving the anti-symmetric port at the dark fringe and $P_{BS,DF}$ is the power on the beam splitter at the dark fringe. Ideally $P_{AS,DF}$ could be driven to be zero, but given the phenomena listed above, it will never be so. The leftover power due to the contrast defect is therefore a pure noise contribution. It is minimised by careful experimental construction and high-quality optics; but when in operation it is the power level against which the local oscillator must be relatively

dominant.

The higher order mode content of the contrast defect can be attenuated by way of spacial mode filtering cavities, Sec. 2.3.5; dedicated mode cleaning optical cavities are used for this purpose. When in use, the output of a Michelson interferometer is not shone directly on a photodetector but passes through the mode cleaning cavity first. This cavity is made resonant with the fundamental mode and reflects the higher order modes, leaving only the fundamental mode to reach the photodetector.

Unfortunately, the fundamental mode content of the contrast defect passes through the mode cleaner. But, given the higher order modes highly suppressed, the overall noise power at the photodetector is greatly reduced.

Where the DC-readout scheme is used a small DARM offset (δ_{DFO}) is maintained — as introduced in Sec. 2.7.3. This provides a local oscillator in the form of the dark fringe offset power. To ensure a good SNR the dark fringe offset power must be dominant over the contrast defect at the photodetector. The two cases, where a mode cleaning cavity is not used and where one is used, are shown in Figs. 2.25a and 2.25b, respectively.




Figure 2.25: Model of the anti-symmetric (AS) port power exiting the Michelson interferometer and reaching the photodetector as a function of DARM. The dashed black line is an idealised model, where perfect destructive interference occurs and 0W is output. Green is power present due to the fundamental mode contrast defect P_{Fund} and blue is the power in the higher order modes P_{HOM} . The red line is the combined output $P_{AS,DF} = P_{Fund} + P_{HOM} \neq 0W$. In order to provide a dominant signal stream and local oscillator a DARM imbalance is maintained, δ_{DFO} , which ensures there is power in the form of a dark fringe offset P_{DFO} where $P_{DFO} > P_{AS,DF}$; i.e. the contrast defect light is dominated by the dark fringe offset (signal carrying) light. (a) This is the case where no mode filtering cavity is included and the higher order modes are free to reach the photodetectors. (b) This is the case where a mode filtering cavity is included and the higher order modes are greatly suppressed. The noise power reaching the photodetector is much less meaning the dark fringe offset can also be reduced without affecting the SNR.

2.9 Pound–Drever–Hall locking

Pound–Drever–Hall locking [102] is a heterodyne type scheme used for sensing the round-trip-phase of a cavity. Sec. 2.7 shows how heterodyne and homodyne schemes can be used to sense displacement of a single mirror. But for a cavity, it is the round-trip-length that ultimately matters because changes in the cavity length induce phase changes in the reflected, circulating and transmitted field of a cavity. Fig. 2.7 shows the power of the field around a cavity, but that is only one side of the problem; the other side is the phase of the fields as shown in Fig. 2.26.




Figure 2.26: The phase of the circulating (black), transmitted (blue) and reflected (red) fields of the cavity as a function of change in length.

The reflected field is the one normally used for this scheme because it has the benefit of being insensitive to power fluctuations caused by factors other than the cavity length [102]¹⁶.

In order to sense the phase change, a local oscillator is required in the form of sidebands created by an electro-optic modulator. The frequency of the sidebands is chosen such that they fall outside of the cavity bandwidth and thus are anti-resonant when the carrier is resonant. The electro-optic modulator produces phase-modulated sidebands onto the field passing through it. With a small modulation depth¹⁷ $\delta \ll 1$, the field E_i transmitted from the electro-optic modulator¹⁸ is [103]

$$E_i = E_0 e^{i\omega_c t} \left[1 + \frac{\delta}{2} e^{i\omega_m t} - \frac{\delta}{2} e^{-i\omega_m t} \right] \quad (2.110)$$

where E_0 is the field input to the electro-optic modulator, ω_c the carrier angular frequency and ω_m the modulation angular frequency. These can now be considered separate fields (combining via super-position): the carrier $E_c = E_0 e^{i\omega_c t}$ and

¹⁶This is because the reflected field for the carrier is at (and held at) a minimum on resonance (Fig. 2.7).

¹⁷This is equivalent to the small modulation index approximation for the Bessel Function treatment of Sec. 2.1.1

¹⁸It is called E_i because it will be the input field to the cavity — it being the transmitted field of the electro-optic modulator is a minor point and not worth naming it after.

the upper and lower sidebands $E_{s+} = E_0 \frac{\delta}{2} e^{i(\omega_c + \omega_m)t}$ and $E_{s-} = -E_0 \frac{\delta}{2} e^{i(\omega_c - \omega_m)t}$, respectively. Since these fields are of a different frequency, their cavity response is not the same. In reflection on resonance, the carrier and the sidebands beat, then the demodulation of these beat signals generates an error signal.

As is visible in Fig. 2.26, the phase sensitivity to cavity length is maximised on resonance, when the reflected field looks like [103]

$$E_{ref} = E_i e^{i\omega_c t} \left[F_{ref}(\phi_c) + F_{ref}(\phi_{s+}) \frac{\delta}{2} e^{i\omega_m t} - F_{ref}(\phi_{s-}) \frac{\delta}{2} e^{-i\omega_m t} \right], \quad (2.111)$$

where F_{ref} is the complex transfer function of the reflection port — that is, the reflection beam over the incident beam. $\phi_c = 2\omega_c L/c$ is the round-trip-phase of the carrier. $\phi_{s+} = 2[\omega_c + \omega_m]L/c$ and $\phi_{s-} = 2[\omega_c - \omega_m]L/c$ are the round-trip-phases for the upper and lower sidebands, respectively.

At the reflection photodetector we see $P_{ref} = E_{ref} E_{ref}^*$ which produces DC terms and terms which oscillate at ω_m and $2\omega_m$. The DC terms are generated by the carrier and sidebands each beating against themselves, while $2\omega_m$ terms are generated by the sidebands beating against the opposite sideband. The ω_m terms are produced by the carrier beating against the sidebands which provide the strongest information about cavity round-trip-length; these are the signals of interest. The other terms can be rejected with AC coupling to the photodetector (or a high-pass filter), and a low-pass filter with cut-off larger than ω_m ¹⁹. With this in mind, the power at the photodetector which oscillates at ω_m is

$$P_{ref,\omega_m}(\phi_c) = 2\sqrt{P_c P_{sb}} \left(F_{ref}^*(\phi_c) F_{ref}(\phi_{s+}) e^{i\omega_m t} - F_{ref}^*(\phi_c) F_{ref}(\phi_{s-}) e^{-i\omega_m t} + F_{ref}(\phi_c) F_{ref}^*(\phi_{s+}) e^{-i\omega_m t} - F_{ref}(\phi_c) F_{ref}^*(\phi_{s-}) e^{i\omega_m t} \right) \quad (2.112)$$

where P_c is the power in the carrier and P_{sb} is the power in each sideband. There is a common complex factor in Eq. 2.112. This factor is designated $\epsilon_{ref}(\phi_c)$

$$\epsilon_{ref}(\phi_c) = 2\sqrt{P_c P_{sb}} \left(F_{ref}(\phi_c) F_{ref}^*(\phi_{s+}) - F_{ref}^*(\phi_c) F_{ref}(\phi_{s-}) \right) \quad (2.113)$$

and using complex analysis, it can be shown that Eq. 2.112 becomes [103]

$$P_{ref,\omega_m}(\phi_c) = \epsilon_{ref,I}(\phi_c) \cos(\omega_m t) + \epsilon_{ref,Q}(\phi_c) \sin(\omega_m t), \quad (2.114)$$

where $\epsilon_{ref,I}$ is the component which is in phase with the real part (amplitude) of ϵ_{ref} , and $\epsilon_{ref,Q}$ is the quadrature phase component (the imaginary part), of ϵ_{ref} [103].

These components can be isolated by demodulation — mixing them electronically

¹⁹See Sec. 2.7 for a description of the demodulation process in general.

with the modulation signal at ω_m and low-pass filtering. Tuning of the demodulation phase θ_{demod} maximises one of the components, which can then be exploited as an error signal for feedback to the cavity length or laser frequency

$$err_{ref}(\phi_c) = \epsilon_{ref,I}(\phi_c) \cos(\theta_{demod}) + \epsilon_{ref,Q}(\phi_c) \sin(\theta_{demod}). \quad (2.115)$$

A plot of these error signals is shown in Fig. 2.27.




Figure 2.27: The Pound–Drever–Hall error signals of Eq. 2.115. (a): $\omega_m \gg \nu_{cav}$. (b): $\omega_m < \nu_{cav}$ where ν_{cav} is the cavity bandwidth (Eq. 2.49).

The motivation for keeping $\omega_m \gg \nu_{cav}$ is because the amplitude of the error signal is greater than when $\omega_m < \nu_{cav}$. Very near to the resonance the carrier receives a large phase shift from the cavity — see Fig. 2.26. But if $\omega_m \gg \nu_{cav}$, as is desirable, the sidebands will not — they are far from their resonance and so their phase is largely unaffected. Moreover, the power in both the carrier and sidebands do not change greatly when the carrier is very near resonance — see Fig. 2.7. The biggest source of change between the fields is their relative phase near resonance and this is why $\epsilon_{refl,Q}$ is larger than $\epsilon_{refl,I}$ for the carrier error — see the central feature of Fig. 2.27a. The relative phase information is contained within the imaginary part of Eq. 2.113, which $\epsilon_{refl,Q}$ is in phase with — where $\epsilon_{refl,I}$ is instead in phase with the amplitude shift of the fields.

2.10 Dither locking

For the input cavities and the interferometer itself, sidebands can be added to the fundamental mode using an electro-optic modulator such as with Pound–Drever–Hall style locks [102]. For other cavities, particularly those in the detection path, it could

be detrimental to use an electro-optic modulator because it introduces some loss²⁰. To generate the phase modulation sidebands for length sensing in the detection path, it is preferable to instead act on the cavity length at some constant frequency. This is known as dithering the cavity. The dither is small, deviating only a few % down from peak transmission — the process of dithering and how it couples to length noise is shown below.

The blue curve of Fig. 2.28 is the modelled transmission of the a cavity (square of Eq. 2.40 with $r = \sqrt{0.925}$, $t = \sqrt{1 - 0.925}$). This curve highlights that when close to resonance the peak of the transmitted power can be approximated as a negative parabola with form $P = P_{max} - \kappa\Delta x^2$, shown as the black dashed line in Fig. 2.28, $\kappa = 2.2 \cdot 10^{16}$.




Figure 2.28: Comparison between the transmission peak of a cavity and a parabola.

P_{max} is the peak transmission power, κ , contains information about the cavity such as Finesse and can be assumed as constant²¹. Δx is the cavity length displacement. If left alone, this would simply be a noise feature, $\Delta x = x_n$, keeping the cavity from stable resonance. To control the cavity, we must observe this noise in order to suppress it.

With dither locking this is achieved by using an oscillator to sinusoidally perturb the cavity length. Now, $\Delta x = x_n + x_d \sin(\omega_d t + \theta)$, where x_d is the amplitude (in meters) and ω_d the angular frequency of the oscillator's dither signal. The unknown phase θ encompasses the delay between the oscillator output, the PZT and, since we are looking for this signal in transmission, also the cavity response delay. The

²⁰For example, the electro-optic modulator used at QUEST has a measured throughput of 97.8 % which is decently high, but a the detection path any loss is to be avoided at all costs.

²¹For this treatment it can be assumed to be constant. In reality, this is the optical gain of the cavity which can drift for reasons outside of length fluctuations.

parabola approximation of transmitted power takes the form

$$P_t = P_{max} - \kappa(x_n + x_d \sin(\omega_d t + \theta))^2 = P_{max} - \kappa \left[x_n^2 + \frac{x_d^2}{2} + 2x_n x_d \sin(\omega_d t + \theta) - \frac{x_d^2 \cos(2\omega_d t + 2\theta)}{2} \right] \quad (2.116)$$

where $2x_n x_d \sin(\omega_d t + \theta)$ is linearly proportional to x_n ; isolating it gives a of measure x_n .

2.10.1 Isolating noise and maximising the error

In order to isolate the component of the signal linearly proportional to x_n , we do some further electronic mixing and demodulating/filtering — the transmitted signal is multiplied by both in phase and quadrature (90 ° out of phase) clones of the dither signal

$$\begin{aligned} P_I = P_t \sin(\omega_d t) &= \left(P_{max} - \kappa \left[x_n^2 + \frac{x_d^2}{2} + 2x_n x_d \sin(\omega_d t + \theta) - \frac{x_d^2 \cos(2\omega_d t + 2\theta)}{2} \right] \right) \sin(\omega_d t) = \\ P_{max} \sin(\omega_d t) - \kappa x_n^2 \sin(\omega_d t) - \kappa \frac{x_d^2}{2} \sin(\omega_d t) - \kappa x_n x_d \cos(\theta) + \kappa x_n x_d \cos(2\omega_d t + \theta) + \\ &\quad \frac{\kappa x_d^2 \sin(3\omega_d t + 2\theta)}{4} + \frac{\kappa x_d^2 \sin(\omega_d t + 2\theta)}{4} \quad (2.117) \end{aligned}$$

$$\begin{aligned} P_Q = P_t \cos(\omega_d t) &= \left(P_{max} - \kappa \left[x_n^2 + \frac{x_d^2}{2} + 2x_n x_d \sin(\omega_d t + \theta) - \frac{x_d^2 \cos(2\omega_d t + 2\theta)}{2} \right] \right) \cos(\omega_d t) = \\ P_{max} \cos(\omega_d t) - \kappa x_n^2 \cos(\omega_d t) - \kappa \frac{x_d^2}{2} \cos(\omega_d t) + \kappa x_n x_d \sin(\theta) - \kappa x_n x_d \sin(2\omega_d t + \theta) + \\ &\quad \frac{\kappa x_d^2 \cos(3\omega_d t + 2\theta)}{4} - \frac{\kappa x_d^2 \cos(\omega_d t + 2\theta)}{4} \quad (2.118) \end{aligned}$$

AC signals can be rejected using low-pass filters, isolating the DC component:

$$P_I = -\kappa x_n x_d \cos(\theta) \quad (2.119)$$

$$P_Q = \kappa x_n x_d \sin(\theta). \quad (2.120)$$

As θ is still unknown, either P_I or P_Q is maximised by multiplying by the rotation matrix with controllable angle ψ , $R(\psi)$

$$R(\psi)P_{I,Q} = \begin{pmatrix} \cos(\psi) & -\sin(\psi) \\ \sin(\psi) & \cos(\psi) \end{pmatrix} \begin{pmatrix} -\kappa x_n x_d \cos(\theta) \\ \kappa x_n x_d \sin(\theta) \end{pmatrix} = \begin{pmatrix} -\kappa x_n x_d \cos(\psi - \theta) \\ -\kappa x_n x_d \sin(\psi - \theta) \end{pmatrix} \quad (2.121)$$

At this point, the signals obtained are the error signals

$$\begin{pmatrix} e_I \\ e_Q \end{pmatrix} = \begin{pmatrix} -\kappa x_n x_d \cos(\psi - \theta) \\ -\kappa x_n x_d \sin(\psi - \theta) \end{pmatrix} \quad (2.122)$$

since in tuning ψ to maximise either e_I or e_Q (while simultaneously minimising the other), we are left with a DC signal linear to x_n which has a 0 point crossing on resonance (where $x_n = x_d = 0$).

Both e_I and e_Q can be used as the error signal, but only one is required. In order to maximise the chosen one, an excitation can be sent into the system. When taking a spectrum of each of the error points after demodulation, this excitation line will be present in the spectra. ψ can then be tuned to minimise the line in the error path which will not be used. This is shown in practice in Sec. 4.5.4.

2.11 V shaped baffle beam dumps

Commonly available beam dumps such as the examples in Fig. 2.29 are perfectly serviceable for in-air experiments.




Figure 2.29: Three common beam dump options [104]. (a): Razor stack style. (b): Radiator style. (c): Trap style.

They are coated with optically absorbant material, such as graphite, which is quite coarse and results in two major issues: it makes them totally incompatible with vacuum usage and their Bidirectional Reflectance Distribution Function (BRDF) is actually quite substantial because scattering, including back reflection, is not negligible. BRDF describes how a surface scatters incident light. The higher quality/more smooth the surface, the less scatter produced, including back reflection. BRDF is measured in per solid angle, sr^{-1} because the sphere of light emitting from a point contains the integrated total of all of that light. If each infinitesimally narrow direction contained light, the integrated sphere would contain infinite light. So, one cannot describe power in a specific direction, it must be per solid angle in that direction.

For in-air, low power operation, the beam dumps of Fig. 2.29 are acceptable. But

for vacuum or higher power use where scattered light must be highly controlled, then a more advanced method is required.

One solution are V-shaped baffles, made of a very absorbant glass with ultra low BRDF owing to the well polished surface, depicted in Fig. 2.30.




Figure 2.30: QUEST’s glass beam dumps, showing the aluminium base plate and beam path. The plate has 2mm deep grooves and the glass is glued down. The small etch to the right side is the intended incident point for the beam, to assist in positioning. Dimensions for the base are in Appendix C.1.

These kinds of beam dumps are used at LIGO and other sites but a comprehensive description of how and why exactly they work was not found. So, this section serves to give what is believed to be the motivation for using these kinds of beam dumps and describes why they are designed as they are.

Summarising the analysis:

- Glass which is very absorbant to 1064 nm light and has a very low BRDF is used.
- Two pieces are arranged in a V-shape in order to generate several internal ‘bounces’.
- The angle between the pieces is such that one of them is parallel to the incident light while the other is the piece the light first impinges on.
- The initial bounce is designed to be at Brewster’s angle [105, 106] which minimises reflection for the ‘correctly’ polarised incident light, which by design and placement, will be the overwhelming majority of the light.
- Almost all of the light is therefore transmitted into the glass and absorbed, being transformed into heat. This is why very absorbant glass is required.
- Given the low BRDF, most of the light that is reflected is contained to the

specular reflection, it is not scattered in all directions — especially back towards the source.

- Subsequent bounces absorb orders of magnitude of whatever light was not transmitted at the initial bounce.
- The light which finally exits the dump is infinitesimal in power.
- Scatter is minimal. There is no out-gassing from coarse coatings — these can be used in vacuum.

2.11.1 The physics of beam dumps

The Fresnel equations [80] describe how the electromagnetic field is reflected and transmitted at an interface/upon meeting a change in refractive index. Each linear polarisation has its own set of equations. The polarisation (typically of the E-field) is defined against the plane of incidence, the plane through which the wave's k-vector is oriented. S-polarised (P-polarised) light is where the E-field is perpendicular (parallel) to the plane of incidence.

The amplitude reflectance (r) and transmittance (t) of the surface are given by [80]

$$r_{\perp} = \frac{n_i \cos(\theta_i) - n_t \cos(\theta_t)}{n_i \cos(\theta_i) + n_t \cos(\theta_t)} \quad ; \quad t_{\perp} = \frac{2n_i \cos(\theta_i)}{n_i \cos(\theta_i) + n_t \cos(\theta_t)} \quad (2.123)$$

$$r_{\parallel} = \frac{n_i \cos(\theta_t) - n_t \cos(\theta_i)}{n_i \cos(\theta_t) + n_t \cos(\theta_i)} \quad ; \quad t_{\parallel} = \frac{2n_i \cos(\theta_i)}{n_i \cos(\theta_t) + n_t \cos(\theta_i)} \quad (2.124)$$

where \perp and \parallel define the fields perpendicular and parallel polarised to the plane of incidence, respectively, n is a medium refractive index and θ is the angle of the k-vector to the surface normal. Subscripts i and t refer to incident and transmission, respectively.

For both polarisations the Snell equation holds [80]

$$n_i \sin(\theta_i) = n_t \sin(\theta_t). \quad (2.125)$$

These transformation matrices for each polarisation can be derived as a function of θ_i , n_i and n_t

$$\begin{pmatrix} r_{\perp}(\theta_i, n_i, n_t) \\ t_{\perp}(\theta_i, n_i, n_t) \end{pmatrix} = \begin{pmatrix} \frac{n_i \cos(\theta_i) - \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}}{n_i \cos(\theta_i) + \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}} \\ \frac{2n_i \cos(\theta_i)}{n_i \cos(\theta_i) + \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}} \end{pmatrix} \quad (2.126)$$

$$\begin{pmatrix} r_{\parallel}(\theta_i, n_i, n_t) \\ t_{\parallel}(\theta_i, n_i, n_t) \end{pmatrix} = \begin{pmatrix} \frac{n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} - n_t^2 \cos(\theta_i)}{n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} + n_t^2 \cos(\theta_i)} \\ \frac{2n_i n_t \cos(\theta_i)}{n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} + n_t^2 \cos(\theta_i)} \end{pmatrix}. \quad (2.127)$$

Ordinarily, these amplitude coefficients could just be squared to create the power coefficients, but for a device dealing specifically in loss as a function of θ_i it is necessary to go a little deeper. Power in a beam is defined as [80]

$$P = AI \quad (2.128)$$

where A is the beam area and I is the beam intensity

$$I = \left(n \frac{\epsilon_0 c}{2} \right) |E|^2 \quad (2.129)$$

where ϵ_0 , c and E are the vacuum permittivity constant, the speed of light and the beam's electric field amplitude, respectively. The refractive index and area can be altered during a transmission, this is why it is not necessarily as simple as squaring Eqs. 2.126 and 2.127 to calculate the transmitted power. Refractive index and area are not changed for the reflected beam, but it will be included in the derivation for completeness. The following derivation applies to both S- and P-polarised fields, so to reduce clutter the subscripts \perp and \parallel will be omitted.

$$R, T = \frac{\text{Reflected, Transmitted power}}{\text{Incident power}} = \frac{I_{r,t} A_{r,t}}{I_i A_i}, \quad (2.130)$$

giving us

$$R, T = \frac{n_{r,t} |E_{0r,t}|^2 A_{r,t}}{n_i |E_{0i}|^2 A_i}. \quad (2.131)$$

For reflection $n_i = n_r$, and symmetry demands the beam's area is unaffected by reflection, i.e. $A_i = A_r$ — see Fig. 2.31²².




Figure 2.31: Reflection beam area.

²²It may actually be unsymmetrical due to divergence/convergence, but reflection is not what causes the asymmetry.

$$R = \frac{n_r |E_{0r}|^2 A_r}{n_i |E_{0i}|^2 A_i} = \frac{|E_{0r}|^2}{|E_{0i}|^2}. \quad (2.132)$$

Since $|E_{0r}|/|E_{0i}|$ is the fundamental definition of amplitude reflectance, r , this gets

$$R = r^2. \quad (2.133)$$

Transmission is a bit more complicated since the beam area and refractive index are both changed.




Figure 2.32: Transmission beam area.

Looking at the 2D representation in Fig. 2.32, it is possible to see that the incident and transmitted beam widths $w_{i,t}$ are related via

$$\frac{w_t}{w_i} = \frac{\cos(\theta_t)}{\cos(\theta_i)}. \quad (2.134)$$

The beam is Gaussian, and so the area is proportional to the width

$$\frac{A_t}{A_i} = \frac{\cos(\theta_t)}{\cos(\theta_i)} \quad (2.135)$$

leading to

$$T = \frac{n_t |E_{0t}|^2 A_t}{n_i |E_{0i}|^2 A_i} = \frac{n_t |E_{0t}|^2 \cos(\theta_t)}{n_i |E_{0i}|^2 \cos(\theta_i)}. \quad (2.136)$$

Again, $|E_{0t}|/|E_{0i}|$ is the fundamental definition of amplitude transmittance, t , we get

$$T = \frac{n_t \cos(\theta_t)}{n_i \cos(\theta_i)} t^2. \quad (2.137)$$

Finally, for the power reflectance and transmittance matrices of a surface, for each

polarisation

$$\begin{pmatrix} R_{\perp}(\theta_i, n_i, n_t) \\ T_{\perp}(\theta_i, n_i, n_t) \end{pmatrix} = \begin{pmatrix} \left[\frac{n_i \cos(\theta_i) - \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}}{n_i \cos(\theta_i) + \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}} \right]^2 \\ \frac{4n_i \cos(\theta_i) \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}}{\left[n_i \cos(\theta_i) + \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} \right]^2} \end{pmatrix} \quad (2.138)$$

$$\begin{pmatrix} R_{\parallel}(\theta_i, n_i, n_t) \\ T_{\parallel}(\theta_i, n_i, n_t) \end{pmatrix} = \begin{pmatrix} \left[\frac{n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} - n_t^2 \cos(\theta_i)}{n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} + n_t^2 \cos(\theta_i)} \right]^2 \\ \frac{4n_i n_t^2 \cos(\theta_i) \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)}}{\left[n_i \sqrt{n_t^2 - n_i^2 \sin^2(\theta_i)} + n_t^2 \cos(\theta_i) \right]^2} \end{pmatrix} \quad (2.139)$$

These matrices can be used to trace a beam's power around a system. For example, the power remaining in the P-polarised field after 3 reflections will be

$$P_{R_{\parallel, tot}} = P_{\parallel, 0} R_{\parallel}(\theta_{i1}, n_{i1}, n_{i1}) R_{\parallel}(\theta_{i2}, n_{i2}, n_{i2}) R_{\parallel}(\theta_{i3}, n_{i3}, n_{i3}) \quad (2.140)$$

and the total power in the transmitted fields will be

$$P_{T_{\parallel, tot}} = P_{\parallel, 0} T_{\parallel}(\theta_{i1}, n_{i1}, n_{i1}) + R_{\parallel}(\theta_{i1}, n_{i1}, n_{i1}) T_{\parallel}(\theta_{i2}, n_{i2}, n_{i2}) + R_{\parallel}(\theta_{i2}, n_{i2}, n_{i2}) T_{\parallel}(\theta_{i3}, n_{i3}, n_{i3}). \quad (2.141)$$

With a complete picture of what happens to a beam when it is reflected and transmitted, attention can be turned to the components of Eq. 2.140 and 2.141, θ_i and n_t (it is assumed that $n_i = 1$).

Brewster's angle [105, 106] is the angle of incidence at which the P-polarised field (E_{\parallel}) has zero reflection as depicted in Fig. 2.33.




Figure 2.33: Brewster's angle. (a): S- and P-polarised light impinging at θ_B , S-pol is fully reflected while P-pol is transmitted. (b): P-pol light impinging so there is no reflection.

The angle itself is given by [105],

$$\theta_B = \arctan\left(\frac{n_2}{n_1}\right) \quad (2.142)$$

This can be appreciated from Eq. 2.139 — i.e. R_{\parallel} will = 0 and $T_{\parallel} = 1$.

If P-polarised light can be made incident at θ_B then reflection should be minimised. This results in almost all of the light being transmitted into the material. If it is a very absorbant material, then the light should mostly be converted to heat, before reaching the other side and transmitting back out.

The other component of Eq. 2.140 and 2.141 is n_t , or in general, which glass to use. LIGO have made similar beam dumps with some being made from Schott BG39 glass [107–111]. This glass is readily available from Schott [112] and for prototypes of these beam dumps, we decided to follow suit. For interest, results of a BRDF analysis including BG39 is shown in Fig. 2.34, though unfortunately the original source could not be found, it was shared via email.




Figure 2.34: Analysis of various glass options.

BG39 has a refractive index of 1.526 [112], giving a $\theta_B \approx 56.8^\circ$ when incident from air.

2.11.2 Baffle design

The baffles are positioned such that the P-polarised light impinges on the first face at θ_B , thereby eliminating the specular reflection almost entirely. After that, each subsequent reflection absorbs orders of magnitude of the remaining light (see Fig. 2.38), resulting in a negligible amount finally exiting the dump.

Fig. 2.35 shows an example geometry and beam path, tracing the angle of incidences (AOIs) for each reflection.




Figure 2.35: Interior angles of the V-shaped beam dump as a function of the incident angle and the plate separation angle.

The angle between the pieces of glass is Φ , the closer Φ is to 0, the more reflections are achieved. Indeed parallel pieces of glass would result in all reflections incident at θ_B , which is hypothetically ideal. But if you desire the V-shape for compactness, the opening of the baffle would quickly become too narrow for a beam to pass through unhindered.

An odd number of reflections results in the exiting beam travelling somewhat in the direction of the input beam — even though there should be negligible power, it is better to avoid this.

Having four reflections seems to be a good compromise since it is easily achieved with a suitably wide baffle opening. Having a section of parallel glass leading to a V shape, in order to get two or three θ_B reflections was considered, but the discontinuous joint between the glass would introduce edges inside the baffle, entirely negating the feature of this glass having low BRDF.

As if by design, $\Phi = 90 - \theta_B$ is possible for four reflections. For BG39, $\Phi = 33.2^\circ$ with this approach. It is beneficial to do this because the second piece of glass is then parallel to the incoming beam (double check Fig. 2.35 if you missed it) and it can therefore be used as a method of coarse alignment. Pointing that glass in the source direction will get the first reflection close to θ_B . There do not seem to be any other significant options of choosing Φ , other than perhaps if a particular exiting beam direction is desired.

The BG39 baffles at QUEST were designed to accommodate four reflections, with the glass separated by angle $\Phi \approx 33.2^\circ$. The first reflection will be at $\theta_B \approx 56.8^\circ$. The subsequent reflection AOIs can be calculated (see Fig. 2.35) and their respective reflectance calculated for each polarisation. Since perfectly polarised/aligned incident light cannot be expected, analysis for both polarisations is carried out below. It is also useful to know the sensitivity to angular accuracy since positioning will be subject to some tolerance.

Fig. 2.36a uses Eqs. 2.138 and 2.139 to show the total power in the final beam, after all reflections as a function of θ_i . Brewster's angle is quite clear. The power in the first reflection only is included for context (Fig. 2.36b), to show the value of subsequent reflections. If $\theta_1 = 58^\circ$ instead of the ideal θ_B , the power exiting the baffle is lowered by \sim five orders of magnitude by virtue of there being three further reflections.




Figure 2.36: Power in reflection. (a): total power after all internal reflections. (b): the power in only the first reflection; highlighting the benefit of multiple internal reflections.

If the light is well polarised and the position of the baffle is good, then the exiting beam will be almost extinguished. There will be some marginal ellipticity in the polarisation, but whatever power is in S should still be absorbed by at least four orders of magnitude. There is quite some forgiveness in the AOI of the first reflection, as long as we are within 3° or so of θ_B , the total exiting power is expected to be below seven orders of magnitude of the input power. There should also be very little scatter due to the high quality and clean surface — most power should remain in the spectral reflection.

2.11.3 Further analysis of powers

In order to get an idea of how good these beam dumps are, it can be useful to calculate the expected transmission and reflection for each point. BG39 has an internal transmittance of $\alpha = 1.452 \times 10^{-4} \text{mm}^{-1}$ [112]. For a distance L propagated through the glass, the transmitted power is then $P_t = P_i \alpha^{L/1e-3}$ [113] where P_i is the input power.




Figure 2.37: Transmitted and reflected fields, with distances travelled through the glass and angles for each part.

Following the beam in Fig. 2.37 through the glass and out the other side,

$$L_1 = \frac{w}{\cos(\theta_b)}, \quad (2.143)$$

$$P_{1[\perp,||]} = P_c T_{1[\perp,||]} \alpha^{L_1} T_{2[\perp,||]}, \quad (2.144)$$

$$L_2 = \frac{w}{\cos(\theta_d)} \quad (2.145)$$

$$P_{2[\perp,||]} = P_c R T_{3[\perp,||]} \alpha^{L_2} T_{4[\perp,||]}. \quad (2.146)$$

The total transmitted power is then

$$P_T = P_{1\perp} + P_{1||} + P_{2\perp} + P_{2||}. \quad (2.147)$$

Plotting for the sum of transmissions around the baffle, and from the first reflection/transmission shows how the majority of the total comes from the first one — see Fig. 2.38.




Figure 2.38: Power %s around the baffle. T is calculated using Fresnel's equations for the amount of the beam transmitted through the inner face, multiplied by the absorption distance travelled through the glass and the transmission through the outer face. R_{\perp} and R_{\parallel}^{1st} are in Fig. 2.36.

2.11.4 QUEST's beam dumps

The beam dumps at QUEST were made from Schott BG39 glass bought from Newport Optics. The construction process is shown in Fig. 2.39. Dimensions for the base are in Appendix C.1.




Figure 2.39: The beam dump glass preparation. Individual pieces of 2mm thick, 50x50mm were bought. The pieces were protected using clear, waterproof Gorilla tape shown in (a); this was necessary because the surface integrity is a key component to their function and the process of cutting could be destructive. They were first clamped between two metal brackets and then the saw was mechanically driven at around 0.5mm per minute — (b). After cutting, the tape was removed using acetone in an ultrasonic cleaner; they were left to soak while being agitated by small vibrations for a weekend — (c). This acetone bath was followed by wiping the glass using alcohol and lens cloths to remove any residual adhesive. The final cleaning step was to use First Contact. The pieces were then fixed to the platform using a standard superglue — (d) (dimensions given in Appendix C.1).

There was an attempt made to measure the transmitted and reflected beams with the intention of characterising the extinction. It was not possible due to the lack of sensitivity of the power meters. Even in complete darkness and the highest power available at the time, no measurable power was recorded. Some photos taken during an attempt are in Fig. 2.40.




Figure 2.40: The beam dump in use; (a) and (b) show the non-optimal and optimal incident angle, respectively. These are photos taken through an infrared viewer (hence the poor quality). They are not intended to give a scientific result, they are here merely for interest. When attempting to measure any power around the beam dump the sensors did not detect anything. The extinction of the light is quite visually impressive when the optimal incident angle (Brewster's angle) is achieved.

Chapter 3

The QUEST Experiment

The QUEST experiment [48] aims to be the most displacement sensitive table-top laser interferometer ever commissioned in the 1 – 250 MHz bandwidth. The principal science goal is the quantum gravity research detailed in Sec. 1.1, with auxiliary goals of searches for ultra high-frequency gravitational waves (Sec. 1.2) and scalar field dark matter (Sec. 1.3). It uses science, technologies and methods spearheaded by the current gravitational wave detectors with an aim to surpass all previous displacement sensitivity benchmarks in the bandwidth of interest.

The previous experiment investigating this quantum gravity model was the Fermilab Holometer [62], which set upper limits on the holographic noise expected by quantum gravity theories based on the holographic principle. It is described in more detail in Sec. 3.2.2. It consisted of a pair of overlapping power-recycled Michelson interferometers which were independent of each other, ensuring they did not share any components. Data was cross-correlated up to 25MHz in order to observe any correlated differential displacements while attenuating uncorrelated noise.

Comparatively, QUEST is a smaller scale table-top experiment operating in the 1 – 250 MHz bandwidth. There are two identical, co-located, shot-noise-limited power-recycled Michelson interferometers of $\cong 1.8$ m arm length placed within independent vacuum systems, mounted on a suspended optical bench for seismic isolation. Signal extraction will be carried out with a DC-readout scheme (Sec. 2.7). A key method of surpassing the upper limits set by the Holometer is to operate at higher power than it did¹, which will improve the shot-noise-limited displacement sensitivity (Eq. 1.36). Given QUEST’s bandwidth, shot-noise is the limiting noise floor and as such, higher power is always of benefit (i.e. we do not expect radiation pressure to be a concern). This higher power brings with it new technical complications and challenges, not least of which is thermal lensing of the optical components which couples fundamental light to higher order modes. These higher order modes lower

¹The Holometer $P_{BS} = 2$ kW vs QUEST’s $P_{BS} = 10$ kW.

the SNR of the detector because they remove power from the signal carrying field while increasing the power in the other (noisy) fields. The higher noise power also risks premature saturation of the photodetector sensors because with higher noise on the detector, a larger DC-offset is required to ensure it is the dominant field. One crucial component which will facilitate the use of high power for groundbreaking sensitivity is the Output Mode Cleaner; a dedicated optical cavity, described in great detail in Chapter 4, placed in the detection path designed to reduce noise prior to light reaching the photodetectors. QUEST also uses squeezed states of light [114], a technology first proven at GEO600 [115] and now used by the LIGO [116] and Virgo [117] collaborations as a method of reducing quantum noise. It will not be described in much detail in this thesis. Given the compact, table-top, re-configurable design, QUEST will be the first detector with the capacity to investigate theories of quantum spacetime which predict angular correlations [27, 118, 119] by exploiting several different optical layouts. The initial configuration of QUEST is the familiar 90° Michelson interferometer arm separation but future configurations include various $\neq 90^\circ$ arm separations, arms with additional bends, and even 3-D designs [48].

This chapter is dedicated to the QUEST experiment in which the performed commissioning and the current status will be detailed. It is worth pointing out from the outset that the word ‘data’ is used almost exclusively in the control and analysis context. It will be made clear when discussing the scientific data of QUEST; otherwise, it should be assumed that data collection is referring to data regarding control systems.

3.1 Sensitivity goal

QUEST’s projected sensitivity to displacement noise for one of its power-recycled Michelson interferometers, assuming 10 kW of power at the beam splitter, is shown in Fig. 3.1 [48]. The data from the pair of power-recycled Michelson interferometer’s will be cross-correlated and averaged — it is expected to take $\sim 5 \times 10^5$ s to reach 5σ confidence level in detecting or constraining quantum gravity signals [48].




Figure 3.1: QUEST design sensitivity [48]. The Fermilab Holometer sensitivity (based on their P_{BS}) for a single of their power-recycled Michelson interferometers is included for comparison. ‘Cavity dynamic response’ takes into account the loss of sensitivity to phenomena which modulate DARM at integer multiples of the power-recycling cavity FSR — this is because DARM will appear unchanged where oscillatory phase changes of $n\pi$ ($n \in \mathbb{Z}$) occur within a single light crossing time: $2L/c = 1/\nu_{FSR}$ per 2.48 (in this case, L is the length of the power-recycling cavity, described in Sec. 3.4.9). The low-pass filtering of the cavity is also clear with the gradual reduction in the sensitivity baseline.

3.2 Comparisons to similar experiments

A brief description of similar experiments is given here. Most are based on the same technologies and principles, but the last couple in Sec. 3.2.3 are very different; they are included because of their bandwidths.

3.2.1 Gravitational wave detectors

Current gravitational wave detectors are configured as power- and signal-recycled Michelson interferometers, with Fabry-Pérot cavities in each arm. Power-recycling is done in order to improve the shot-noise related SNR. Light returning from the Michelson interferometer to the symmetric port is reflected back towards the beam splitter and kept coherent with the injection light. The number of round-trips of the power-recycling cavity, n has the effect of improving the SNR for the given input power by \sqrt{n} [76]. The power-recycling cavity is resonant with the carrier frequency. A signal-recycling is a similar concept, though it is instead resonant with the signal sidebands. The signal-recycling cavity facilitates a tuning of the Michelson interferometer’s frequency response, it can widen the detector’s general bandwidth

[120, 121] or conversely greatly improve sensitivity to narrower bandwidths [76]. QUEST has no current plans to implement signal-recycling.

The arm Fabry-Pérot cavities² serve to contain each photon for an average number of $\approx \mathcal{F}/2\pi$ [78] round-trips. In the case of Advanced LIGO $\mathcal{F} \approx 450$ [87, 122, 123], meaning the phase shift produced by the gravitational wave altering the arm length is sampled by the average photon ≈ 72 times, compounding the effect of the length change. This provides a proportionate gain to the phase sensitivity introduced due to differential-mode arm (DARM) length fluctuations of the Michelson interferometer. However, there is some debate whether these arm cavities actually reduce sensitivity to quantum spacetime signals — this is an ongoing area of research [48, 124, 125].

Readout of these detectors has a bandwidth up to only a few kHz [126]. This is ideal for compact binary mergers, with frequencies low enough to cause an effectively stationary strain for each photon completing its many round-trips of the arms. But for phenomena causing strain/length noise at higher frequencies, frequencies comparable to the single photon cross time of the detector (~ 75 kHz for a 4 km arm), these detectors are not sensitive.

Additionally, phenomena thought to generate coherent signals in detectors occupying highly overlapped volumes of space-time (i.e. co-located detectors) cannot be investigated due to there being only one detector at each site. There is also no possibility to re-configure these large ground based gravitational wave detectors to any new geometry — they form fixed Michelson interferometers with 90° separated arms and given their size, it is not feasible that the geometry will change. Any phenomena with angular correlations which could be better probed by arms separated by $\neq 90^\circ$, and even some 3-D configuration, are out of reach to these detectors.

3.2.2 Fermilab Holometer

The Fermilab Holometer [62] was the first detector of its kind. A pair of twin, co-located 40 m power-recycled Michelson interferometers, with independent injection, detection, vacuum, electronics and control schemes. The outputs were cross-correlated with a bandwidth of 25 MHz. The Holometer constrained the magnitude of quantum space-time fluctuations in the cross-spectrum of co-located power-recycled Michelson interferometers [127]. However, given the size of this experiment reconfiguring it to any other geometry was not possible. Though this experiment was designed specifically for quantum space-time searches, those theories predicting angular correlations [27, 118, 119] were not investigated.

²These arm cavities are an example of an overcoupled cavity. The circulating and reflected power is increased while the transmitted power is reduced, compared to that of a critically coupled cavity with similarly reflective mirrors. This is because the goal of the cavities is not to transmit the beam but rather to contain the photons for many round-trips and direct them back towards the Michelson interferometer, which in terms of the fields of the cavity, is the reflected field.

It was also restricted in making improvements to its shot-noise limited ASD. With no inclusion of higher order mode filtering optics in the detection path, the full power of the higher order mode light (as introduced in Sec. 4.1) reached the low-power threshold photodetectors. In order to resolve faint signals the power in the fundamental mode DC-offset must be dominant over the contrast defect (Sec. 2.8) which includes higher order mode light. The lack of mode filtering required a larger DC-offset which approached the photodetector power threshold — preventing Fermilab from increasing their circulating power any further.

Additionally, the Holometer did not use vacuum squeezing, the process of driving one uncertainty quadrant (phase or number/field amplitude) of the photon down at the expense of increasing the other, in order to improve the quantum-noise-limit for the frequency of interest [73, 114].

3.2.3 High frequency gravitational wave detectors

There was a laser interferometry based experiment built to detect 100 MHz gravitational waves before Fermilab’s Holometer. They used two synchronous recycling interferometers whose outputs were cross-correlated, with each achieving a strain sensitivity of $\sim 10^{-16} \text{ Hz}^{-1/2}$ at 100 MHz [128]. This style of interferometer was chosen over the Fabry-Pérot Michelson interferometer and L-shaped cavity Michelson interferometer because of the relative simplicity [129]. With the synchronous interferometer the two paths are identical, they are just propagated in opposite directions. This means there is only one degree of freedom to control. It also means that any differential imperfections of the mirrors are no longer troublesome since each path is constructed of the same mirrors. The experimental layout of one of the interferometers is in Fig. 3.2.




Figure 3.2: The synchronous recycling interferometer [128]. This is the experimental layout of one of the interferometers used. The input laser is split at the beam splitter and made resonant with the recycling cavity which is formed with a recycling mirror, two end mirrors and the transfer mirror. The two inputs to the recycling cavity cycle around the cavity in opposite directions until leaking back to the beam splitter and towards the detection photodetector.

There are two interesting gravitational wave detectors which searched in the MHz region [130, 131]. Neither relied on laser based interferometry but both used optical cavities of different sorts.

In 2001 a group used a pair of cylindrical microwave cavities which were coupled via a small iris on the cavity axes [130]. The experiment established a symmetrical resonating field in the cavities of frequency ω_s . When a gravitational wave passes through at frequency ω_{GW} and harmonically varies the length of the cavities, a power transfer to an anti-symmetric resonant field at frequency ω_a is expected. In order to generate this field, the harmonic length variation must have frequency $\omega_{GW} = |\omega_a - \omega_s|$. The difference frequency, i.e. the required frequency of the gravitational wave, is a function of the coupling iris diameter. This experiment achieved a displacement sensitivity of $\delta L/L = 3.3 \times 10^{-20} \text{ Hz}^{-1/2}$ at coupling frequency 1.38 MHz.

In 2006 another group used a pair of torus shaped microwave wave-guides with conducting probes to detect fluctuations in the polarisation of the resonating field [131]. A gravitational wave with a propagation direction perpendicular to the wave-guide axis is expected to warp the space-time such that the polarisation angle is directly altered, with this effect being compounded each round-trip of the resonating field. The probe only conducts if the polarisation axis is perturbed. This experiment used cross-correlation between two detectors to improve sensitivity by at least an order of magnitude. They achieved a noise spectral density of $\sim 10 \times 10^{-14} \text{ Hz}^{-1/2}$ at 100 MHz.

3.3 From the empty lab to QUEST

The following subsections, 3.3.1 and 3.3.2 briefly describe some of the preliminary work and prototype stages for the QUEST experiment. At the beginning of this PhD the lab was empty, not even the optical tables were installed. As QUESTS's first PhD student (who was focused on the optical experimental work), these were some of the activities which took over a year and a half to complete; so it is worth giving them a space in this chapter.

3.3.1 Laser beam profile characterisation

Characterising a laser's beam profile is the process of measuring the beam size at several positions in order to model the waist. This can be done at any point along the layout if optics exist and are changing the beam shape. For example, it could be necessary to characterise the beam following a series of lenses in order to assess lens positions or plan for additional optics downstream. Characterisation can, maybe should, also be done as a first step in using a laser. The laser itself generates a waist and in order to model the experimental layout, that waist must be understood.

The process of characterisation is to measure the beam size at different locations along the propagation. This is done by incrementally blocking the beam using something with a straight edge, ideally it should be thin in order to not cause orientation issues. If something a few cm thick were used, ensuring the edge was exactly parallel to the beam propagation direction would be important. We used a single sided razor blade. The beam is slowly blocked by translating the razor perpendicularly to the beam propagation direction. We used a micrometer translation stage. The power on the other side of the razor is measured for each step.

The Gaussian intensity of the beam defined in Eq. 2.31 is a function of r , the distance from the propagation axis which at this point is unknown. So r becomes $(x - r)$ where x is the distance the razor is translated. Also, given we are using a power meter, which naturally integrates the intensity over the sensor surface area and so intensity become power [93],

$$P(z, r) = \frac{2P}{\pi\omega(z)^2} e^{-2(x-r)^2/\omega(z)^2} \quad (3.1)$$

Data at the two extreme points of a measurement is in Figs. 3.3a and 3.3b. This P vs x is then transformed into a Gaussian distribution by way of an averaging derivative method [93]

$$\frac{dP}{dx} = -\frac{1}{2} \left(\frac{y_{i+1} - y_i}{x_{i+1} - x_i} + \frac{y_i - y_{i-1}}{x_i - x_{i-1}} \right) \quad (3.2)$$

A negative appears here where it is not in [93] because they begin with a fully

blocked beam and steadily reveal it rather than our method of the reverse. Eq 3.2 is the average of the derivatives between each point and its two immediately adjacent points. The result produces an array of length one less than the data array and so we added a zero to end of the array. The result is in Figs. 3.3c and 3.3d where the Gaussian distribution, Eq. 3.1 has also been fit.

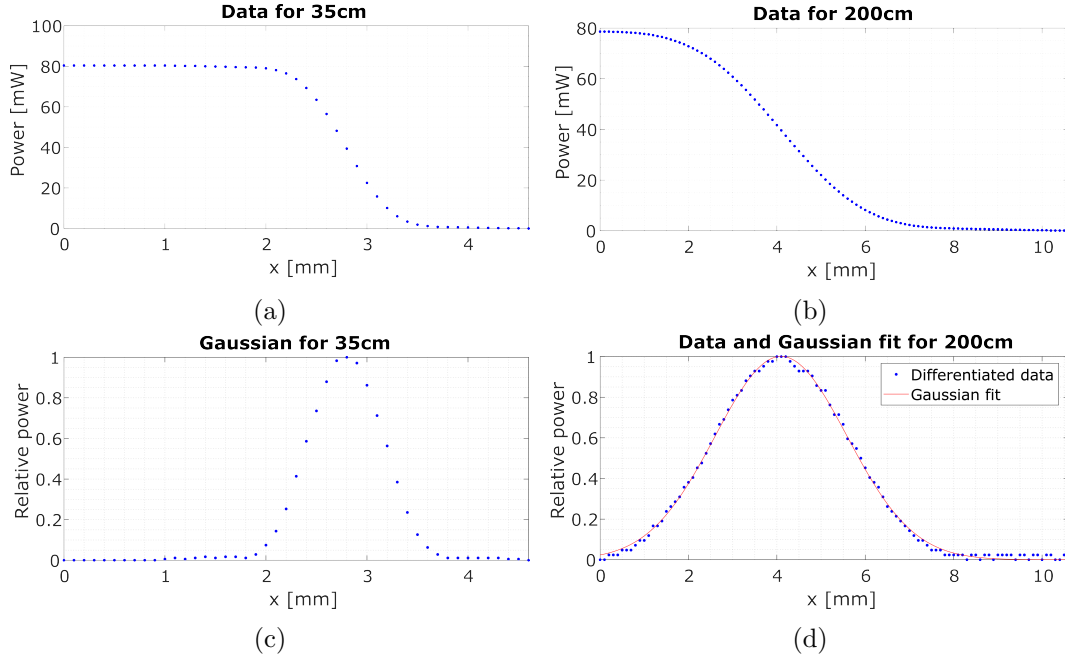


Figure 3.3: Knife-edge data at two distances, 35 cm and 200 cm from the laser output. (a) and (b) show the raw measured data where the beam is incrementally revealed, (c) and (d) are the averaged derivative data and a Gaussian fit.

With the many Gaussian fits providing a beam radius at the propagation distances, the beam divergence can be modelled. This is done using a fit which includes a ‘quality factor’ M where for an ideal Gaussian $M^2 = 1.0$ [78, 132]. The beam radius $\omega(z)$ is

$$\omega(z) = \omega_0 \sqrt{1 + \left[\frac{(z - z_0)\lambda M^2}{\pi\omega_0^2} \right]^2} \quad (3.3)$$

where λ is the laser wavelength, z is the measured distance from some reference (the edge of the laser head in this case) and z_0 is the position of the waist. Data for the measured beam radii and the fit (Eq. 3.3) are in Fig. 3.4




Figure 3.4: Measured beam radii plotted against distance from the laser output. Two fit lines are shown, one with the ideal Gaussian beam where $M^2 = 1.0$ and one with a slightly less pure Gaussian where $M^2 = 1.1$.

For the final measurements, whose data are in Table 3.1, we used a Thorlabs beam profiler [133]. The process is the same in principle except the profiler measures the beam dimensions in less than a second. This of course hugely sped up the process — each measured position took seconds to gather dimension data for, rather than the almost half an hour of the ‘old school’ knife edge method we initially carried out.

During each characterisation, the lasers were outputting approximately 0.5 W since that is the optimum operating power.

3.3.2 Prototype stages

There were three prototype stages in preparation for QUEST. They were all completed prior to the installation and commissioning of the control and data system now in use. A brief description of each will be given here since they all contributed to the foundations of the current workings of QUEST.

The 1.5 m Fabry-Pérot cavity

On October 14th 2020, we achieved our first milestone of locking the in air Fabry-Pérot cavity. It was a 1.5 m linear cavity with a flat input mirror and a 6 m radius of curvature end mirror (the same ones used for the final QUEST layout).

The cavity was setup by first placing the input optics up to and including the steering mirrors (but not the mode-matching lenses), and then using fixed irises to ensure the beam was level and directed along the bench screw holes. The end mirror was

placed and aligned by making sure the reflected beam followed the same path back through the irises. The input mirror was then positioned using the same process. After this, the mode-matching telescope lenses were placed in approximately the correct positions to create the ~ 0.9 mm cavity waist; the positions were found with the script as described in Sec. 2.4.1. They were mounted on micrometer translation stages for fine tuning. Once all the optics were on the table, fine alignment and mode-matching was achieved using a cavity length scan driven by a PZT glued to the end mirror.

The mirrors used had a specified reflectivity of $R = 0.995 \pm 0.003$ giving an expected $\mathcal{F} \approx 626.7$. We measured a Finesse of 611.2 ± 7.3 ; for comparison, the Advanced LIGO army cavities have a measured Finesse of ~ 450 [87, 122, 123]. The Finesse was measured by carrying out a length scan and fitting an Airy distribution of the form [134]

$$A_t = y_0 + \frac{a}{1 + 4\mathcal{F}^2 \left(\frac{\nu - \nu_0}{\nu_{FSR}} \right)^2} \quad (3.4)$$

to a resonance peak — Fig. 3.5.

Locking was achieved with a Pound–Drever–Hall lock facilitated by a Moku:Lab [95], using the Laser Lock Box instrument [135].




Figure 3.5: The 1.5 m Fabry-Pérot cavity resonance used to measure the Finesse.

This setup gave us experience setting up, characterising and locking a linear cavity. It is the most ‘basic’ optical layout of this kind but many valuable and transferable skills were gained for all the subsequent work.

The 2 m Michelson interferometer

On June 10th 2021 the in air Michelson interferometer was maintained at a DC-offset. This was achieved using a simple feedback loop with a user defined set-point.

Within the bandwidth of the loop, the output (the dark fringe offset) followed the input (the set-point; see Eq. 2.87). This was done using the Moku’s Digital Filter Box instrument [136].

The Michelson interferometer had 2 m arms with 6 m concave end test masses (ETMs). There was no mode-matching required for this, but alignment was tricky. The X-arm (that transmitting through the beam splitter) ETM was placed and aligned first, using irises. The beam splitter was then added, which introduced a small shift in the beam path as expected. The beam splitter was aligned at 45° to ensure the reflection beam (defining the X-arm) followed the table’s screw holes at 90° from the X-arm. Once that was done, the X-arm end mirror was then slightly shifted to put the mirror centre at the beam spot, the mirror alignment was recovered by making sure the reflection from it, and through the beam splitter followed the injection path irises. The X-arm was then blocked and the X-arm end mirror added and aligned using the same injection irises.

Fine alignment was then carried out using the anti-symmetric port. An IR viewer card was placed far from the beam splitter. With the X-arm blocked, the X-arm generated spot on the IR card was recorded as the *target*. The X-arm was then blocked and the X-arm aligned to hit the target. This established a more fine tuning but the final tuning was done by observing the fringe patterns when a DARM length scan was carried out.

An indefinitely stable DC-offset was maintained. A 5-minute data set is shown in Fig. 3.6. The trend is the output with the loop closed. The vertical axis limits are the maximum and minimum of a DARM length ramp measured immediately prior to engaging the fringe control, to show the bright and dark fringe range (~ 40.4 mW). A DC-offset of (11.0 ± 0.4) mW was maintained, meaning a standard deviation $\sim 1\%$ of the fringe range.




Figure 3.6: 5-minutes of data for the prototype in-air Michelson interferometer. The vertical axis limits represent the power of the dark and bright fringes, giving a fringe range of ~ 40.4 mW.

This gives us an idea of the contrast defect. As described in Sec. 2.8, the contrast

defect is the light which is imperfectly destructively interfered at the beam splitter and so leaks out of the anti-symmetric port while at the dark fringe. Unfortunately, the input power was not logged during this measurement so an estimation of losses cannot be made. A well aligned and constructed Michelson interferometer should have a bright fringe very near the input power.

The fringes are measured using a photodetector which outputs a voltage. This voltage is converted to incident watts when measuring powers, but it can also be used to calibrate a sensitivity curve. The voltage of each fringe extreme is recorded and the difference is assumed to represent a differential-mode length change of $\lambda/2 = 532$ nm. While the fringe control is engaged and a DC-offset is maintained, the mean DC value gives the displacement from the dark fringe. A spectrum of this measurement provides the displacement sensitivity as shown in Fig. 3.7.




Figure 3.7: ‘Sensitivity curve’ for the 2m arm length in-air Michelson interferometer. No gravitational waves were detected...

3.4 QUEST’s design

The key aspects to the design of QUEST will be described here keeping in mind that commissioning efforts are still ongoing. Fig 3.8 shows the optical layout of the key components of one of the Michelson interferometers. The Output Mode Cleaner is not described in detail in this chapter, it is the subject of Chapter 4.




Figure 3.8: QUEST's experimental layout — for one of the two identical apparatus. The power-recycling mirror (PRM) forms a cavity with the Michelson interferometer arms, generating power build up when operated near the anti-symmetric port dark fringe. A squeezer is injected into the output Faraday isolator in order to reduce vacuum fluctuation noise. The Output Mode Cleaner (OMC) rejects higher order mode light exiting the instrument before the signal carrying light is split between two photodetectors, for noise cancelling and redundancy purposes. The distance separating the flat PRM to the 50/50 beam splitter is 5 cm. From the beam splitter (BS) to each 6 m radius of curvature (ROC) ETM it is 1.8 m. The power-recycling cavity waist, located at the PRM, is $\sim 968 \mu\text{m}$.

3.4.1 AMY and BOB

QUEST's two instruments are called AMY and BOB. They have been built and used in tandem for different scientific and technical benchmarks. Both vacuum systems are complete and in place, but the optical configurations are different.

AMY is the 'lead' instrument, it is configured as a power-recycled Michelson interferometer and it is the one used to design and commission the control schemes. The work discussed in this chapter was primarily carried out on AMY.

BOB is currently a Fabry-Pérot cavity comprised of (what will be) the power-recycling mirror (PRM) and the X-arm ETM. It is used for testing the high power amplifier and soon, the squeezer. There is also a planned auxiliary experiment regarding testing the coating tolerances on the mirrors. The circulating power will be driven higher and higher until the Finesse of BOB is seen to deteriorate. This will indicate a loss of reflectivity and potentially give an assessment of the coating

tolerance. This thesis does not detail the work ongoing on BOB.

Once these various sub-systems have been understood, and control loops and such are sufficiently robust, BOB will be transformed into a power-recycled Michelson interferometer and AMY will have the separate high power amplifier and squeezer installed. Both will then be identical but independent high power, quantum enhanced power-recycled Michelson interferometers.

For the sake of simplicity, from here on it will be assumed that both AMY and BOB are identical in their configuration.

3.4.2 Co-located

AMY and BOB are identical power-recycled Michelson interferometers which are co-located as closely as is practicable in order to maximise their shared volume of spacetime. The reason for the dual and overlapping design is described in Sec 1.4.5. Each Michelson interferometer defines a causal diamond which restricts the entropy content and thus entangles the quantum spacetime fluctuations expected within. The two Michelson interferometers which share a region of spacetime define a common causal diamond (Fig. 1.8); now, otherwise independent entangled fluctuations are present and correlated to both instruments due to this spacial overlap. The time integrated cross-spectrum of their separate data will establish a suppression of random and uncorrelated noise, while strengthening the common and correlated fluctuation signals.

3.4.3 The power-recycled Michelson interferometers

The following will refer to a single power-recycled Michelson interferometer, but it applies equally to both and it also describes the current configuration of QUEST. Unless otherwise stated, each component mentioned is separate to its respective power-recycled Michelson interferometer and has independent electrical power and controls from the other.

The Michelson interferometer has an arm length of 1.8 m separated by 90° using a 50/50 beam splitter, delimited by ETMs with a 6 m concave radius of curvature. The inclusion of the flat PRM at 5 cm from the beam splitter establishes the power-recycling configuration. The input power of 0.5 W provided by a Coherent Mephisto [96] laser at wavelength $\lambda = 1064$ nm will be amplified to 10 W via the neoVAN Amplifier from neoLASE [137] and will be further amplified via the power-recycling to an expected 10 kW on the beam splitter.

The common-mode power-recycling cavity, defined by the common-mode arm length, CARM (Eq. 1.32) is mode-matched to using a two lens mode-matching telescope and aligned to with two input steering mirrors. The cavity waist, located at the

PRM, is $\sim 968 \mu\text{m}^3$. The laser frequency is locked to CARM using the reflected field and a Pound–Drever–Hall lock acting on the laser crystal’s PZT and temperature (see Sec. 2.9 and [102]).

The differential-mode arm length, DARM is controlled using the ETM longitudinal actuators by maintaining a DC-offset at the anti-symmetric port.

The power-recycled Michelson interferometer optics are housed within vacuum chambers, while the injection and detection optics are not in vacuum.

All of the in vacuum, injection and detection optics including the squeezer (see Sec. 3.4.13) and the laser are housed on the 13-legged, seismic isolated bench — the layout is depicted in Fig. 3.9. Since the frequency of interest is 1 – 250 MHz, seismic noise is not of great concern for QUEST in terms of affecting the overall SNR— this means there is no need to suspend any of the optics or components beyond that of the isolation offered by the bench; and the feedback loops will maintain resonance otherwise.

3.4.4 Bench layout

The compact, table-top design enables straightforward re-configurations of the geometry and optical layout. The first version of QUEST are 90° power-recycled Michelson interferometers, while future configurations include $\neq 90^\circ$ power-recycled Michelson interferometers, power-recycled Michelson interferometers with one or both arms containing an additional bend, and eventually 3-D layouts [48].

³This is calculated using the steps in Sec. 2.4.1.

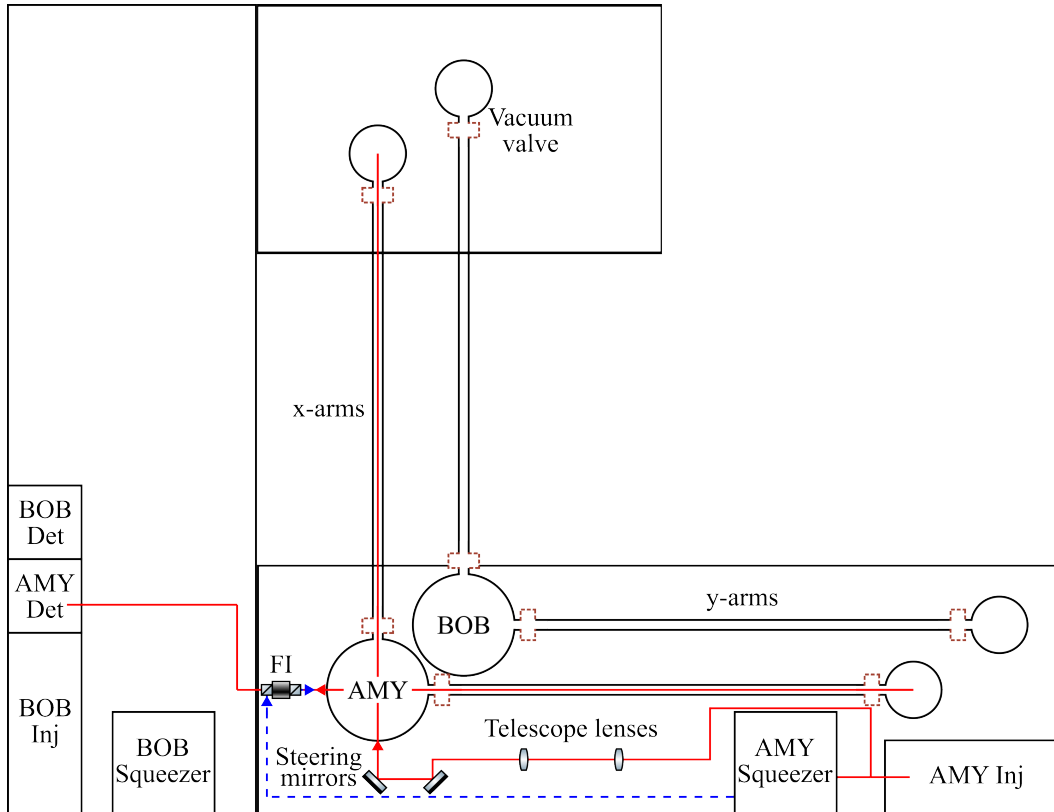


Figure 3.9: Layout of QUEST as on the optical bench. There are three separate benches which have been bolted together to form one rigid bench. The injection breadboard is detailed in Fig. 3.10 and the detection breadboard in Fig. 4.15. The only optics which could not be contained to these two breadboards are the mode-matching telescope and steering mirrors for the power-recycled Michelson interferometers; and aside from them, additional mirrors for directing the beam (including some periscopes because the arms are the same height and the beam must be navigated around them where necessary).

The mode-matching telescope for the power-recycled Michelson interferometer was designed using software called ‘Just another mode matching tool’ (JamMT) [138], and confirmed using the method showcased in Sec. 2.4.1.

3.4.5 The injection breadboards

The injection breadboards are configured as in Fig. 3.10. They are designed to accommodate all of the necessary optics and components to prepare the light for injection into their respective power-recycled Michelson interferometer, and also provide a pick-off to the squeezer. The Pre-Mode Cleaner is shown here, but it is not currently in use. It is a small triangular cavity used to filter the beam of higher order modes prior to use in the power-recycled Michelson interferometers. It is functionally identical and similar in design to the Output Mode Cleaner, Chapter 4, but it is physically larger, has a higher Finesse and a lower bandwidth. It has been constructed but not yet used. Space is reserved for it and the mode-matching telescope optics are in position such that if it is deemed necessary, adding it will not

cause large disruptions downstream.

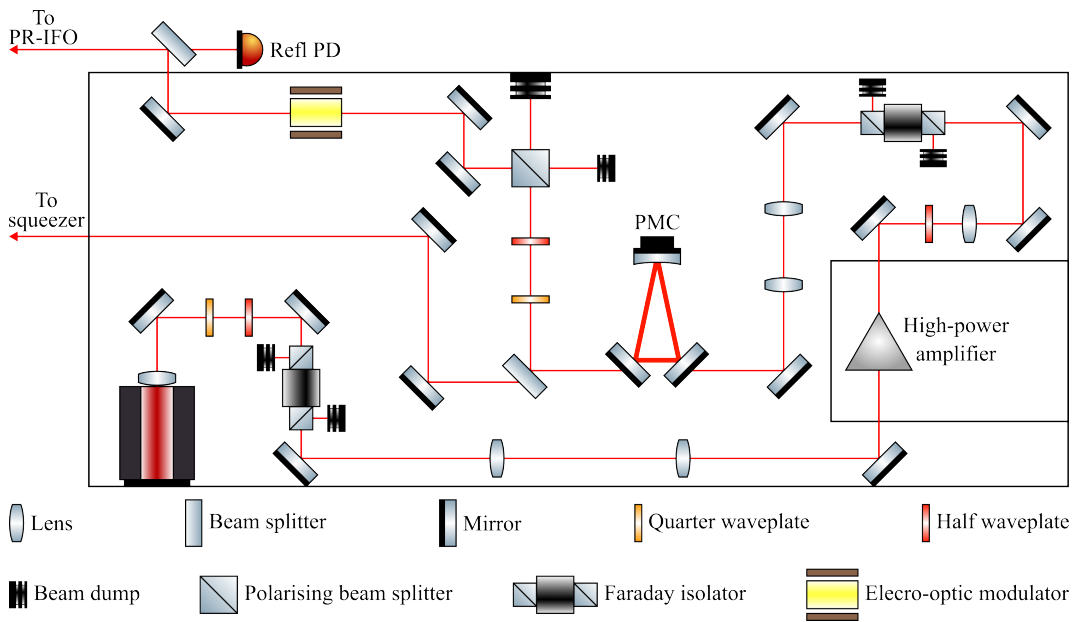


Figure 3.10: The injection breadboard. The black outer rectangle is the perimeter of the breadboard itself, some nearby components are shown for completeness. The Pre-Mode Cleaner (PMC) is shown but it is not currently installed. The Refl PD is the power-recycled Michelson interferometer (PR-IFO) reflection photodetector; it is used for the Pound-Drever-Hall lock.

Each of the mode-matching telescopes were designed with a mixture of JamMT [138] and the Matlab solver of Sec. 2.4.1. With the various lenses included in the layout, modelling the beam is a vital aspect of the design. Profiling the beam output by the laser itself was the first step, since from that, everything else follows. The beam profile method is described in Sec. 3.3.1 and the results are in Sec. 3.4.7.

Fig. 3.10 also shows the electro-optic modulator and reflection photodetector. These are used for power-recycled Michelson interferometer control. The electro-optic modulator produces phase modulated sidebands onto the carrier and the reflection photodetector detects the field reflected from the power-recycled Michelson interferometer — this is described in more detail in Sec. 3.5.

3.4.6 Vacuum chambers

The vacuum chambers shown in Fig. 3.11 were designed with future customisability in mind.




Figure 3.11: The vacuum chambers of QUEST. The smaller one is the end chamber which will house the respective ETM, the larger is the central chamber which will house the PRM and the beam splitter. The many limbs of the chamber are to facilitate future configurations which establish a $\neq 90^\circ$ separation of the arms. The end chamber also contains multiple possible configurations because future configurations include bent arms with various bend angles.

The vacuum system has three distinct pumps which are used in sequence.

1. Primary pumps — one per Michelson interferometer.
2. Turbo pumps — one per chamber.
3. Ionic pumps — one per chamber.

The primary pump establishes a vacuum level of 10^{-3} mbar after which the turbo pumps take over and bring it down to $\sim 10^{-6}$ mbar. Although the turbo pumps are not noisy, they do produce vibrations sufficient to prevent locking the laser to the power-recycling cavity — they are attached to the chamber and so any vibrations are immediately coupled to the optic mounts. Once the pressure reaches 10^{-6} mbar, the ionic pumps will be activated and the turbo and primary pumps switched off. The ionic pumps will bring the vacuum to the final desired pressure of $\sim 10^{-8}$ mbar. Each end of the arms has a vacuum seal valve highlighted by red ellipses in Fig. 3.12. This allows individual chambers to be brought to room pressure as required without losing vacuum everywhere else.

The ionic pumps have been delivered but we do not yet have their controllers so they have not been commissioned. However, the primary and turbo pumps were tested. The goal vacuum was $\sim 10^{-8}$ mbar which was well surpassed. The corner chamber reached 2.88×10^{-8} mbar while 2.73×10^{-8} mbar and 1.82×10^{-8} mbar was reached for x- and y-end chambers, respectively.




Figure 3.12: The (unfinished) vacuum chambers connected by the vacuum pipe. The red ellipses are the seal valves which enable chamber specific vacuum control.

3.4.7 The lasers

The beam profile of the lasers used at QUEST were characterised using a Thorlabs beam profiler [133] — see Appendix 3.3.1 for a full description of beam profile characterisation. The data are in Table 3.1. Values for Laser 1's waist position are not of the same precision as the others because the measurement was made at a different time. Initial measurements showed too poor a roundness and so it was sent back; these values are for the lasers currently in the lab.

During each characterisation, the lasers were outputting approximately 0.5 W since that is the optimum operating power.

Laser 0 is our 'spare', used for auxiliary work. Lasers 1 and 2 are the injection lasers for BOB and AMY, respectively.

		Waist size [mm]	Beam roundness	Waist position [mm]
Data sheet [96]		0.16	< 1.1	-105
Laser 0	x	0.193 ± 0.002	1.119 ± 0.024	-167.212 ± 10.663
	y	0.216 ± 0.004		-178.345 ± 23.042
Laser 1	x	0.133 ± 0.009	1.150 ± 0.090	-138.6 ± 72.7
	y	0.153 ± 0.006		-170.6 ± 40.7
Laser 2	x	0.145 ± 0.013	1.076 ± 0.115	-227.078 ± 118.584
	y	0.156 ± 0.009		-203.571 ± 74.502

Table 3.1: Laser beam profiles.

During operation they are always set to output 0.5 W and left for a few minutes to stabilise prior to use. However, the first optics downstream of the laser output are a quarter-waveplate, half-waveplate and polarising beam splitter sequence (see Fig.

3.10) which is used to control the power for the remainder of the layout. This gives far more power control and prevents instabilities in the laser's output that can occur when not operating at the nominal output power.

3.4.8 The end test mass actuators

The mounts for the in-vacuum optics which form the power-recycled Michelson interferometer have inbuilt pico-motors, see Fig 3.13a — driving the pitch and yaw⁴ adjustments with a step size of $0.7 \mu\text{rad}$ and a range $\pm 61 \text{ rad}$ [139]. Coarse alignment is carried out using these motors — as well as the ordinary manual screw type adjusters. The pico-motors provide a good enough range of adjustment when the vacuum chambers are closed.

The ETM optics are also attached to three separate PZTs arranged as an equilateral triangle, see Fig 3.13b. These PI S-316 PZTs [140] have a $12 \pm 2.4 \mu\text{m}$ displacement range each, providing $12 \pm 2.4 \mu\text{m}$ longitudinal range and $1200 \pm 240 \mu\text{rad}$ angular tilt range in both pitch and yaw. These PZTs have finer displacement than the pico-motors with a step size of around 0.2 nm longitudinal and $0.05 \mu\text{rad}$ angular.

Fig. 3.14 shows the position of the PZTs for each ETM mount. The diagrams are as if looking at the high-reflectivity surface of the mirrors.

⁴Of an aircraft, pitch would be nose up/down while yaw is nose left/right. For a mirror, the reflective surface normal is analogous to the aircraft nose.




Figure 3.13: The ETM mounts and actuators. (a): Optical mount with pico-motors providing ± 61 rad in pitch and yaw. These pico-motors are accessed via Ethernet switches and can be used while the chamber is under vacuum. (b): The PZTs at the front face of the ETMy mount, before the mirror was added. The ETMx mount PZT layout is mirrored, as shown Fig. 3.14. The PZTs are connected to and controlled by the control and digital system (Appendix A.1). They provide $\pm 1200 \mu\text{rad}$ in pitch and yaw and $\pm 12 \mu\text{m}$ longitudinal range.




Figure 3.14: End test mass PZT locations when looking at the high-reflectivity (HR) surface.

3.4.9 Interferometer mirrors

The $R = 99.5\%$ PRM currently in use does not have the final design reflectivity (of $R = 99.9\%$). The control schemes are still being commissioned and a lower Finesse power-recycling cavity is useful for that because it is easier to lock given the wider bandwidth (see Sec. 2.3.3). However, the ETMs in place are the design-reflectivity ones. Table 3.2 lists the mirror specifications, the resulting power-recycling cavity Finesse and other important properties. Values in bold are estimated using the measured mirror reflectivities. Measurements were carried out as depicted in Fig.

3.15.

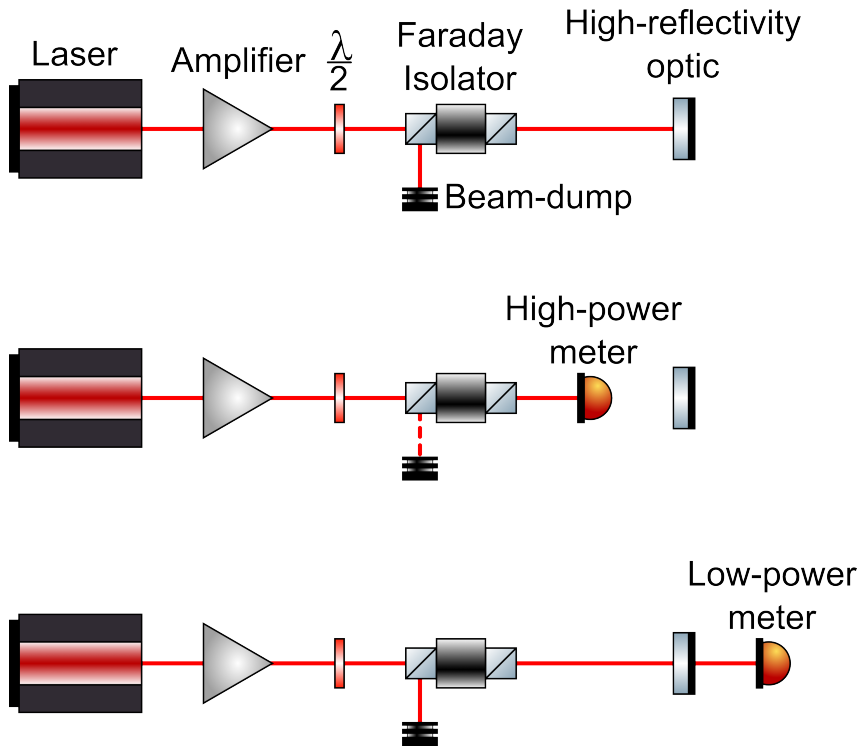


Figure 3.15: Characterisation of high reflectivity optics. The laser amplifier diode current was set to the nominal 10 A. The beam was passed through a half waveplate and a Faraday isolator. The mirror was positioned in the transmission path of a Faraday isolator. A normal angle of incidence was achieved by ensuring the reflection was directed back into the Faraday isolator, identified where minimal beam deformation is visible and power is maximised at the Faraday isolator rejection port. A high-power meter [141] was used to measure the incident power (which was limited to ≈ 9.5 W using the half waveplate), and a low-power meter [142] for measuring power transmitted through the mirror(s).

The power meters each display the incident power so for this measurement there was no requirement to convert from a voltage (as provided by a photodetector for example) to power. Unfortunately, these measurements were made only once so an uncertainty was no recorded.

	Trial spec	Measured	Final spec
PRM (R)	99.5 %	99.53 %	99.9 %
ETMs (R)	99.995 %	99.9987 %	99.995 %
Finesse	~ 1250	1330	~ 6000
Bandwidth	65 kHz	61 kHz	13.5 kHz
P_{BS}	~ 7.8 kW	~ 8.4 kW	~ 36.3 kW

Table 3.2: Table of QUEST optical specifications. P_{BS} values assume 10 W input power.

FSR is based on power-recycling cavity length L_{PRC} , which is the Michelson interferometer common arm length \bar{L} as defined in Eq. 1.32, plus the distance between the

beam splitter and the PRM (5 cm in this case) — then $\text{FSR} = c/2L_{PRC} = 81 \text{ MHz}$ as in Eq. 2.48.

3.4.10 Mode matching telescopes and lenses

There are a few mode-matching telescopes for each power-recycled Michelson interferometer which can be separated into three groups: injection, power-recycling cavity and detection.

The injection breadboards, Fig. 3.10, are designed to accommodate a Pre-Mode Cleaner if it is deemed required. There are lenses already in place to mode-match to that cavity.

Following the Pre-Mode Cleaner, the high-power amplifiers also need mode-matching and therefore a dedicated telescope of their own.

The final mode-matching telescope is for the power-recycled Michelson interferometer itself. This is not currently on the injection breadboard, it is instead on the optical table given space restrictions.

An additional mode-matching telescope is required for the squeezer as well.

Aside from the mode-matching telescopes, there are some individual lenses used for focusing the beam in order to avoid clipping in various strategic places. Key components include the Faraday isolators and the electro-optic modulator used for power-recycled Michelson interferometer control (see Sec. 3.5.1).

The layout was designed using a combination of the Matlab script described in Sec. 2.4 and JamMT [138].

3.4.11 Contrast defect at QUEST

Contrast defect was introduced in Sec. 2.8; it is the light which couples to the anti-symmetric port due to imperfect interference of the fundamental mode and higher order modes generated via optical imperfections and thermal lensing.

At QUEST the design goal for the contrast defect (defined in Eq. 2.109 as $\text{CD} = \frac{P_{AS,DF}}{P_{BS,DF}}$) to be $< 10^{-6}$ for fundamental mode light contribution and $< 10^{-5}$ for the higher order mode contribution. The higher mode content is more lenient because of the Output Mode Cleaner, the subject of Chapter 4.

3.4.12 Readout

Phenomena which generate oscillating DARM changes produce the signals we are attempting to observe⁵. These differential oscillations will generate phase modulated

⁵Gravitational waves and dark matter specifically generate differential-mode length changes. Quantum gravity fluctuations are not directly differential effects, but they will cause random length

sidebands on the carrier and be output at the anti-symmetric port which must be somehow identified as separate to the noise of the contrast defect. The three main methods of reading these signals: heterodyne, homodyne and DC are described in Sec. 2.7. QUEST is using the DC-readout method for reasons as listed in [99] — key motivations include:

- With heterodyne readout there are harmonic sidebands at twice the heterodyne frequency. These contribute only to noise.
- Calibration of the data is simpler because the signals of interest are only modulating one carrier, not two sidebands.
- Spatial overlap of the carrier/local oscillator and the signal field is guaranteed.
- Injecting squeezed light is simpler for a DC-readout than for heterodyne.

When a DC-readout scheme is implemented, a DARM offset is introduced and maintained in the Michelson interferometer. This ensures a small amount of fundamental mode light is coupled to the anti-symmetric port to be used as a local oscillator for signal extraction. This light is called the dark fringe offset or DC-offset. The chosen magnitude of the dark fringe offset is dependent on the contrast defect because the dark fringe offset must be the dominant light source in order to resolve faint signals.

In principle, the dark fringe offset could always be made large enough to dominate the contrast defect, but increasing it introduces new noise sources — in particular, the coupling of common-mode noises to the anti-symmetric port, i.e. laser power noise [99]. It also reduces the power on the beam splitter since it is, in effect, introducing loss in the power-recycling cavity, which lowers the signal-to-shot-noise-ratio (Eq. 1.36).

Another issue with arbitrarily increasing the dark fringe offset (or the circulating power in general) is the saturation limits of the photodetectors. Low noise, high bandwidth photodetectors have a very limited power threshold [143]. If the combined power of the contrast defect and the dark fringe offset were to approach this saturation limit, any further increased in the circulating power would damage the photodetectors. A balance must be maintained, but ideally the dark fringe offset is kept as small as possible while ensuring it is larger than the the contrast defect.

3.4.13 Squeezing

Squeezing⁶ is a method of manipulating the Heisenberg uncertainty relationship between the number of and phase of photons — see Sec. 1.4.2. One of the uncertainties is reduced at the expense of the other which ultimately for QUEST’s use, results in

changes which are uncorrelated in the differential arms, leading to an output.

⁶The use of entangled squeezing [144–146] is also a consideration.

a reduction in the shot-noise⁷.

A Faraday isolator will be positioned between the anti-symmetric output and the detection breadboard. This serves to provide an input point for squeezed vacuum states of light [73, 115, 147, 148].

The first squeezer was delivered in October 2022, it has been installed and is currently being commissioned. The plan is to have the first observation of squeezed light at QUEST in mid 2023.

3.4.14 Scientific data acquisition

The cross-correlated data is the principal scientific data output of QUEST. The acquisition of scientific data was itself a PhD project⁸ given the complexity of handling such huge amounts of data, and having to perform real-time cross-correlation on them. The flow of scientific data acquisition is depicted in Fig. 3.16.

There are two identical 500 MB/s, 16-bit, 4-channel digitisers (NI PXIe-5763), one for each Michelson interferometer. The digitisers each perform FFT using their Xilinx Kintex UltraScale KU035 FPGAs. This FFT is the noise power spectral density (PSD) of each Michelson interferometer, which is expected to equal the shot-noise PSD for ~ 10 kW of circulating power [48].

Real-time cross-correlation and averaging is then performed using an FPGA co-processor module which contains a Xilinx Kintex UltraScale KU060 FPGA (NI PXIe-7915). The cross-spectrum PSD is the geometric mean of the individual noise PSD, which decreases with root of the number of spectra measured.

These digitisers and the co-processor are housed in a PXI chassis (NI-PXIe-1092). The PXI chassis is connected to a DELL PowerEdge R7525 server computer which receives the processed data and controls the digitisers and co-processor. A PXI remote module (NI PXIe-8398) enables communication between the PXI chassis and the server computer.

⁷Squeezing can be performed in order to reduce radiation pressure, but this is not required for QUEST since that noise is below the bandwidth of interest.

⁸S. M. Vermeulen, whose thesis is in progress — expected submission is alongside this one.

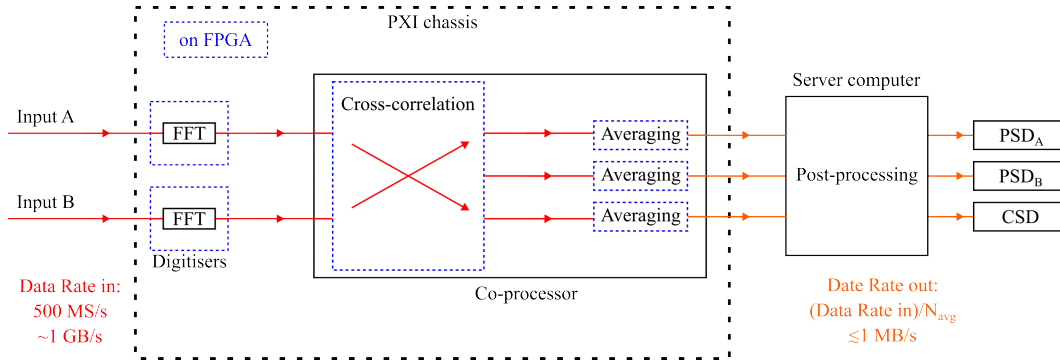


Figure 3.16: Data acquisition signal flow. Each Michelson interferometer inputs data into a standalone digitiser where FFT is carried out producing the individual noise power spectral density (PSD_{A,B}). The FFT data is sent to the co-processor where cross-correlation and averaging takes place. This cross-correlated, averaged data is the cross-spectral density (CSD).

3.5 QUEST’s lock scheme

Control for QUEST is performed using our Control and Data System (CDS) — see Appendix A — and Moku:Labs [95]. CDS has a bandwidth of 32 kHz which is generally fine for most control data needs, but for schemes which require higher bandwidth — e.g. Pound–Drever–Hall locks, the Moku:Lab is used. Moku:Lab is a very versatile tool which will be made reference to at multiple points throughout this thesis. They have an input bandwidth of DC to 200 MHz, an output bandwidth of DC to 300 MHz and an input referer noise better than $30 \text{ nV}/\sqrt{\text{Hz}}$ above 100 kHz [95]. There are several of them at QUEST, exploited in different applications — to drive the electro-optic modulators for all our Pound–Drever–Hall locks, as spectrum analysers, signal generators and oscilloscopes. They will be referred to as “Moku” from here on. For Pound–Drever–Hall lock, their Laser Lock Box instrument [135] is used.

There are two controlled longitudinal degrees of freedom for QUEST — CARM and DARM⁹ — and control of them is acquired sequentially. As is discussed in Sec. 3.4.8, each ETM has three separate PZT actuators arranged in an equilateral triangle at the back of the mirror which are responsible for both longitudinal and angular actuation. There are also the angular degrees of freedom: pitch and yaw, each with both common and differential actuation. There is no alignment feedback so far, but it is likely to be a requirement in the short-term. Two models designed using a Simulink interface (see Appendix A) were built, one for longitudinal/length sensing and control (LSC) and another for alignment sensing and control (ASC).

The signals for the PZTs are passed through a matrix which applies appropriate coefficients in order to get the desired adjustment. Each of the models has a matrix

⁹There is also the power-recycled cavity length, but this is equal to the distance between the beam splitter and the PRM (which is fixed) plus CARM — there is no separate actuator to control this degree of freedom specifically.

for the respective degree of freedom. The matrices are shown in Table 3.3, the values therein are a matter of geometry and determined by the PZT’s position on the mirror (shown previously in Fig. 3.14). For example, if manually driving DARM, the user inputs some value which is cloned into six total signals, three to each ETM. Through the LSC matrix, the PZTs for ETM_x gets the raw signals, those for ETM_y get the negated signals establishing the differential movement. All signals to the PZTs are passed through these matrices whether they be manual signals or those output by control loops.

	LSC Model		ASC Model			
	DARM	CARM	D-PIT	C-PIT	D-YAW	C-YAW
X1	1	1	0	0	1	1
X2	1	1	1	1	-0.5	-0.5
X3	1	1	-1	-1	-0.5	-0.5
Y1	-1	1	0	0	-1	1
Y2	-1	1	-1	1	0.5	-0.5
Y3	-1	1	1	-1	0.5	-0.5

Table 3.3: End test mass PZT signal distribution matrices. PIT and YAW refer to pitch and yaw, D- and C- to differential and common.

3.5.1 Lock acquisition

The layout for QUEST’s lock scheme is in Fig. 3.17.

Since Moku has the bandwidth required to facilitate a Pound–Drever–Hall lock, it is used for CARM control. It drives an electro-optic modulator on the injection breadboard (see Fig 3.10) at 8.1 MHz which generates phase modulation sidebands at a frequency $\nu_c \pm 8.1$ MHz, where ν_c is the carrier frequency of the laser. The photodetector located at the rejection port of the injection Faraday isolator is then used to detect the reflection from CARM, and is demodulated within the Moku to create the Pound–Drever–Hall error signal. This photodetector signal is also sent to the CDS for slow controls and data acquisition.

The CARM error signal is generated and observed but the Pound–Drever–Hall loop is not closed at this point. Next, a manual DARM signal is sent to the ETMs while observing the photodetector at the anti-symmetric port to ensure the dark fringe is approximately reached. This has the effect of increasing the Finesse of the power-recycled Michelson interferometer or CARM cavity by ensuring a maximally reflective ‘end mirror’ of that cavity — see Sec. 2.5. Higher Finesse in the CARM cavity means its bandwidth is reduced; this reflects the sidebands more strongly, enhancing the Pound–Drever–Hall error signal.

The temperature of the laser is then tuned to bring the CARM error signal closer to

a 0 V offset. The lasers have two remote frequency adjustment inputs [96], the PZT which acts directly on the laser crystal (fast control), and the temperature of the crystal (slow control). The PZT has a high bandwidth (100 kHz) but limited range (± 65 MHz) [96] — less than the cavity FSR, see Sec. 3.4.9, while the temperature control has a low bandwidth ($\cong 1$ Hz) but a very high range (30 GHz) [96].

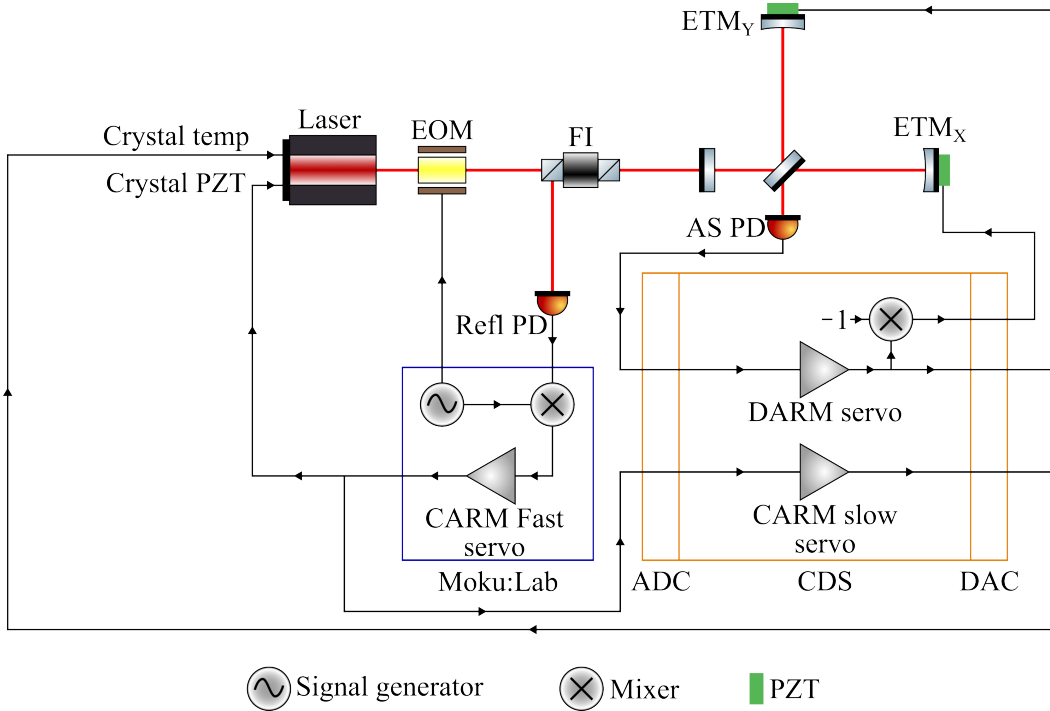


Figure 3.17: The layout for the QUEST lock scheme. The blue box is the Moku:Lab which is used as the fast controller for the CARM Pound-Drever-Hall lock — it contains a signal generator which drives the electro-optic modulator (EOM) which produces phase modulated RF sidebands on the carrier. The reflection photodetector (Refl PD) is positioned at the ejection port of a Faraday Isolator (FI), which is then input to the Moku:Lab and mixed with the signal generator signal for demodulation. The CARM fast servo provides the CARM control signal. This control signal is output to the laser crystal PZT for fast feedback; it is also sent to CDS. The CARM slow servo within CDS filters the CARM control signal with a pole at 0.1 Hz. This provides feedback for the laser crystal temperature controller. The anti-symmetric photodetector (AS PD) is used for the DARM servo. DARM controls are output to the ETM PZTs and the AS PD signal is minimised for the dark fringe. ADC and DAC are the CDS analogue to digital and digital to analogue converters, respectively.

We take the control signal from the Moku and split it; one goes directly to the laser PZT since the bandwidth of the Moku is higher than the CDS — so it is best for higher frequency control — while the other control signal is sent to CDS. With this input, we not only have access to analyse the CARM control signal, but we can also use it to feed to the laser temperature after applying suitable filtering. The way the Moku generates the error signal is to output a triangular ramp to the laser PZT and plot the demodulated reflection photodetector signal (the error signal) as a function of volts sent to the PZT — see Fig. 3.18. Where a resonance is swept through, the signature Pound-Drever-Hall error signal is visible, and its location on the x-axis is

what DC-offset the PZT will require in order for the laser to reach that frequency. The user then selects the desired error signal and the Moku locks to it. The problem with this method is if the chosen resonance is far from 0 V, then the DC-offset output to the PZT can limit its range when the loop is engaged since the output limit of the Moku is only 2 Vpp. The chosen method to reduce this offset is to manually drive the laser temperature with CDS to bring an error signal to the 0 V point on the error display. This ensures the laser's 'resting' frequency is approximately correct and the PZT will begin at its null position, leaving the full range for the loop. With the error signal over the 0 V point, the error signal is selected and the loop is closed, locking the laser to CARM.




Figure 3.18: (a): error signal requiring some offset to the PZT. (b): temperature tuned to center the error signal, the laser is approximately frequency matched to the cavity and the PZT can begin holding resonance with the full range.

The Moku's output — the CARM control signal — is then used in CDS for the LSC closed loop to drive laser temperature. The Moku is responsible for this loop at frequencies above 1 Hz and the CDS for frequencies below it — the pole low-pass filter for the temperature is at 0.1 Hz. The range of the temperature is so large

that it will override the PZT and so it must be attenuated to only drive very low frequencies.

With CARM now locked, the DARM loop can be engaged. A DARM dither signal at ~ 1.5 kHz is sent to the ETM PZTs and the anti-symmetric port photodetector is used for demodulation. The minimum output is reached and maintained using a dither lock — Sec. 2.10.

3.5.2 DC lock

The final stage of QUEST’s lock procedure is to switch to a DC-offset. This locking method ensures a sufficient amount of the fundamental mode is output at the anti-symmetric port so as to be dominant over the noisy light and can therefore be used as the local oscillator for signal extraction.

The DC-offset is achieved as a final lock stage by disabling the DARM dither loop and switching to a simple feedback loop with a user defined set-point. The power at the anti-symmetric port will shift from the previous minimum and once the desired set-point is reached, the loop now operates to maintain it. This DC-offset is carried out in tandem with the CARM Pound–Drever–Hall loop.

Currently the DC lock (and DARM control in general) is achieved using a 90/10 beam splitter pick-off from the beam en-route to the detection bread board as shown in Fig. 3.19. Since this pick-off introduces losses it is not desirable, particularly when squeezing is included in the overall scheme. The prevailing thinking for the final locking scheme is to use the Output Mode Cleaner reflection and transmission photodetectors in combination. The sum of the three rather than just the two transmission photodetectors should establish a decoupling from the Output Mode Cleaner lock status, with only a small increase in overall loss when locked — due to the higher number of round-trips and thus interactions with slightly transmissive mirrors — but this will only be $\sim 1\%$ of the available power (see Sec 4.7.3). A calibration would be necessary since in transmission are the high-bandwidth Newport 1811-FS photodetectors [143] and in reflection is a Thorlabs PDA20CS2 photodetector, but provided the voltage is accurately converted to power, this calibration should be quite straightforward and yield an equivalent assessment of the power at the anti-symmetric port as a single photodetector directly in the anti-symmetric port path.

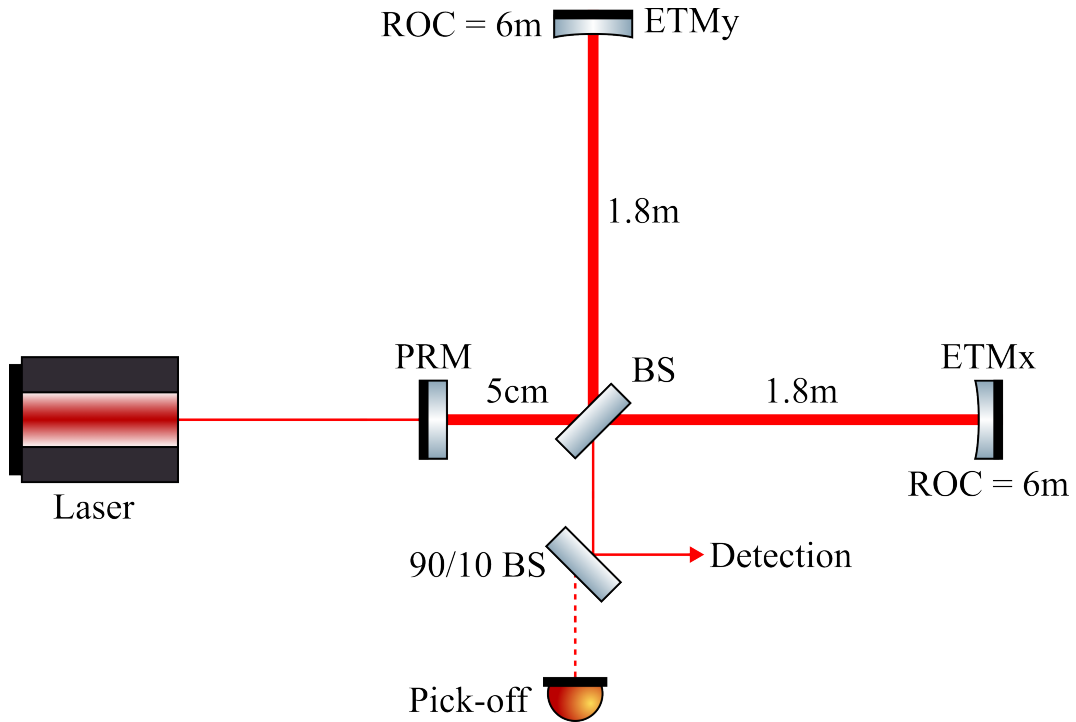


Figure 3.19: Anti-symmetric port pick-off for dark fringe and DC-offset control. The pick-off is achieved using a 90/10 beam splitter, with 90% in reflection (towards the detection breadboard) and 10% in transmission (towards the pick-off photodetector).

3.5.3 Locking interface

Aside from the Moku's direct connection to the laser PZT, all the controls are driven by CDS. Fig. 3.20 shows the overview screen for the LSC loop which is designed to give a graphical representation of the components involved and make it easier to act on the controls.

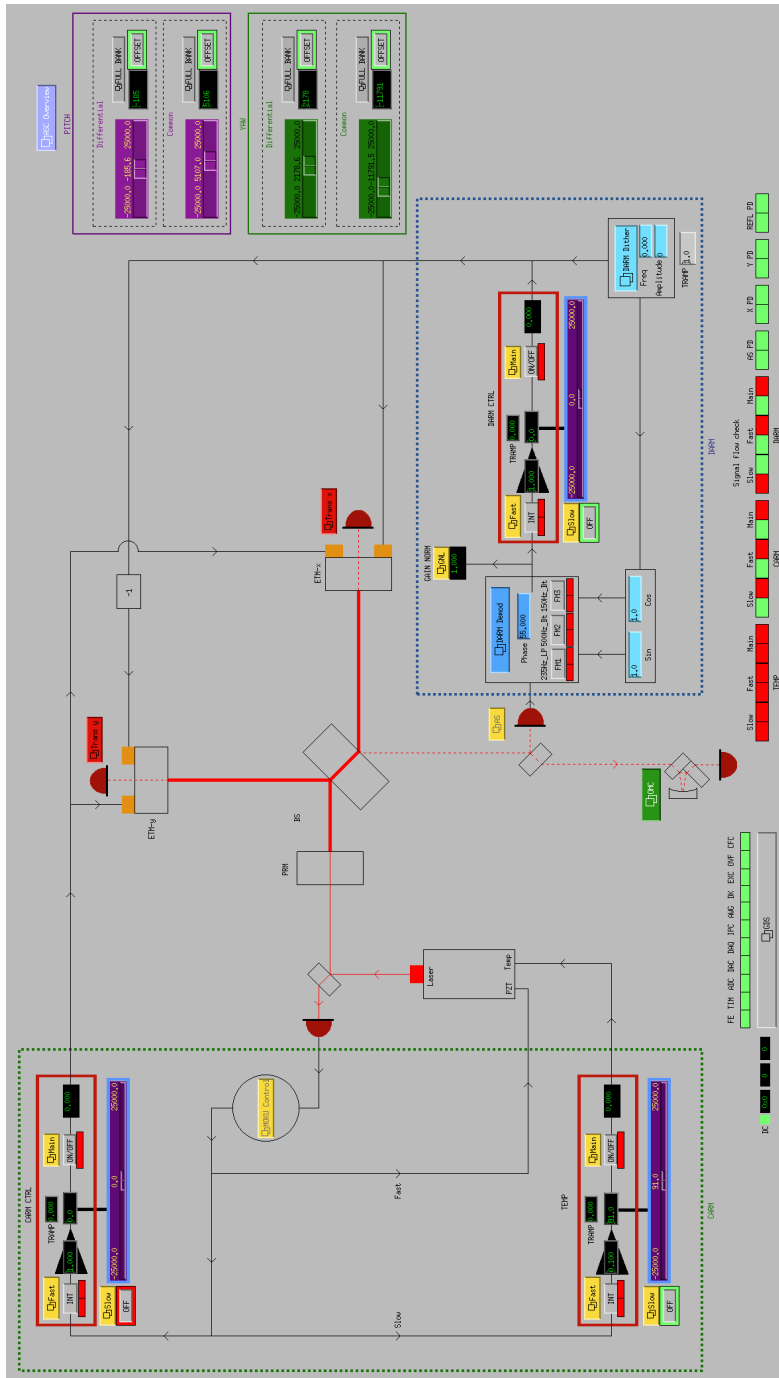


Figure 3.20: LSC overview screen. The central image is the power-recycled Michelson interferometer, with the injection laser and the various important surrounding photodetectors. Top right has quick links to the alignment controls since these are quite often needed when locking and attempting to minimise the dark fringe. Left side in a green dashed box is the CARM section. It includes the laser temperature and CARM signals to the ETM PZTs (though the ETM PZTs are not currently involved in the CARM loop); the Moku control signal is the input to both of these. The lower right blue dashed box is the DARM section. The input is the anti-symmetric photodetector and output are the ETM PZTs. The frequency and amplitude of the dither signal can be interacted with as required; it is always output to the ETM PZTs and to the demodulation part of the loop. The overall model signal flow check bar at the bottom is a fast way of showing the user where signals are being blocked. Each should be green for the signals to be flowing. The green bar to the left, above the ‘GDS’ button is a health status check, it shows various parts of CDS and can alert the user if something has gone wrong within the system.

For each of the separate loops — DARM, CARM (PZTs) and CARM (temperature) — there are currently three filter bank stages (see Appendix A).

1. The ‘Fast’ filter bank contains the loop filters, this is the controller.
2. The ‘Slow’ filter bank is present to facilitate adding an offset without that offset being passed through the controller filters and gain.
3. And the ‘Main’ filter bank is there to give one final off switch to the output.

The signal flow is as in Fig. 3.21.




Figure 3.21: Signal flow for the three filter bank stages in LSC. Fast contains the loop filters and gain, so adding an offset here is not always the best place since it would be multiplied by the gain. Slow provides a location to add an offset which will not be multiplied. The output of Fast and Slow are then summed and passed through the Main bank which provides a point to observe the combined signal and block both at once if necessary.

3.5.4 Angular to longitudinal control coupling

Table 3.3 shows the *basic* matrices which ensure the mirrors receive the correct actuation. This does however assume each PZT is identical in operation and that the equilateral geometry/orientation is perfect. Given there are three PZTs, for purely longitudinal drive all three need, in principle, to be displaced by an equal amount, while in reality, they behave slightly differently to the same voltage. When driving DARM or CARM, there is some coupling to the angular degree of freedom as shown in Fig. 3.22. To investigate this, a quadrant photodetector¹⁰ [149] was positioned at the anti-symmetric port of AMY. With the Y-arm closed and the PRM removed, an 11 Hz longitudinal excitation was sent to drive ETMx. The spectrum of the quadrant photodetector pitch and yaw channels very clearly shows the coupling.

¹⁰A quadrant photodetector is a photodetector whose sensor is split into four quadrants. This provides a translation sensing capability. If the beam moves, some quadrants will receive more power while other will lose power i.e. the spot position can be monitored.




Figure 3.22: An 11 Hz drive was acting on ETMx longitudinal control while a quadrant photodetector at the anti-symmetric port monitored for angular coupling. The coupling is clear in both the pitch and yaw degrees of freedom. This shows when driving DARM, the mirrors are also changing alignment.

This means that when driving DARM to control the anti-symmetric output, this angular coupling is constantly causing imperfect destructive interference at the beam splitter resulting in a prematurely high contrast defect type power. The angular coupling also means the alignment to the power-recycling cavity is not stable because the cavity axis itself is moving. As the alignment accuracy fluctuates, this results in a constantly changing amplitude of the fundamental mode at the anti-symmetric port and higher order mode power being reflected.

A second matrix is included after the ‘basic’ matrix in the model(s) in order to minimise this coupling. Excitation is injected to the longitudinal controls while the anti-symmetric quadrant photodetector spectrum is observed. The second matrix values are then adjusted until no coupling remains. It is not yet fully understood whether minimising the coupling holds for all alignments or just whatever alignment they had during the analysis — commissioning is ongoing.

3.6 Loop characterisation

With the DC lock engaged, an open loop transfer function was taken of the DARM loop — shown in Fig. 3.23. A swept-sine was injected at the ‘DARM Fast’ filter bank excitation point, then the ‘IN1’ and ‘IN2’ test points (Fig. A.7) were used to measure the open loop transfer function — see Sec. 2.6.1 for a more thorough description of measuring a transfer function. The loop DC gain was 100 and a unity gain frequency of ~ 150 Hz with around 80° phase margin was measured.




Figure 3.23: Open loop transfer function of QUEST’s DARM loop during DC lock.

3.7 The lock

The power-recycled Michelson interferometer DC lock has seen some success. The first 2-minute long lock is shown in Fig. 3.24. During lock, the power transmitted through the X-arm ETM was measured to be $P_X \approx 260 \mu\text{W}$. This optic was characterised with a measured transmission of $T_X \approx 27$ ppm, so P_{BS} was estimated to be $P_{BS} \approx 2P_X/T_X \sim 19 \text{ W}$ — the factor two accounts for there being two arms.

The set-point of this lock was 30 mW. The first minute is reasonably stable at $\approx 57 \text{ mW}$ after which a downwards drift occurs until lock was interrupted.

The final goal for QUEST is for the contrast defect and higher order mode content of the beam to be 10^{-6} and 10^{-5} compared to the fundamental mode, respectively [48]. There has been no assessment of the higher order mode content so far because the beam is not stable enough to align to the Output Mode Cleaner yet. This noisy DC lock is likely due at least in significant part to the angular to the longitudinal coupling of the mirrors.




Figure 3.24: Anti-symmetric port power (P_{AS}) during 2-minutes of DC lock. The mean of this data is $P_{AS} = (52.64 \pm 9.8)$ mW, the error being the standard deviation. Using the transmitted power of the X-arm ETM, $P_{BS} \approx 19$ W.

With loop improvements/commissioning there has since been an over-the-weekend lock achieved, with a 10-hour period of reasonably stable DC lock — shown in Fig. 3.25. Though it proves long periods of stability, with the longest of around 10 hours, the high peaks also show that the magnitude of δ_{DFO} was far too large during the ‘stable’ parts. The data is from the X-arm transmission, which has a linear relationship to the circulating power¹¹, and so power spikes in transmission up to ~ 4 times that during lock show P_{BS} could have been at least a factor 8 higher (because there are two ETMs) with a smaller δ_{DFO} .

¹¹See Eqs. 2.42 and 2.43, their ratio is T_{ETM} but since there are two ETMs, it is actually halved.




Figure 3.25: Weekend DC lock data from the X-arm transmission photodetector. Stable regions are DC lock with the longest being around 10 hours. Since this is the X-arm transmission, which is a function of the circulating power, peaks show the DC-offset was too large. Circulating power (and therefore ETM transmitted power) spikes when the δ_{DFO} is minimised. The ~ 4 times increase in power from lock to spikes show that during lock, far more power than is desirable was output.

Improvements to this lock are an ongoing area of the overall commissioning effort. However, the angular to longitudinal coupling as discussed in Sec. 3.5.4 has prevented any further investigations into the lock specifically.

3.8 QUEST Sensitivity (model)

During a stable DC lock, a high-bandwidth photodetector was used to trial assessing the sensitivity to displacement. Details of the powers at the time of measurement are summarised in Table 3.4.

Input power	(357 ± 0.1) mW
Trans ETMy	(923 ± 10) μ W
P_{BS}	(142 ± 2) W (inferred)
DC-offset set-point	30 mW
P_{AS}	avg = (30 ± 15) mW, rms = 33 mW

Table 3.4: Sensitivity measurement details.

A noise budget was produced prior to the measurement, shown in Fig. 3.26.




Figure 3.26: Noise budget for QUEST sensitivity measurement. The dark noise is the measured photodetector voltage with the sensor covered. Moku input is the terminated input noise of the Moku. DC-offset is the measurement of the anti-symmetric port during a DC lock.

The result is the blue trend in Figs. 3.26 and 3.27. The peaks we believe to be the Pound–Drever–Hall sidebands ($\nu = 8.1$ MHz) and their higher order harmonics. These can be minimised by lowering the sideband modulation depth which will reduce the power coupled to the sidebands. The downwards trend from the left is not understood. The output spectrum is expected to be flat, see Fig 3.1 — further investigations and noise hunting are ongoing.

The measurement was calibrated to the expected shot-noise-limited ASD for $P_{BS} = 142$ W using Eq. 1.36. In this frequency range, we assume this is a good approximation given the noise budget. The dark noise becomes suspect around 30-40 MHz after which is comparable to the measurement.




Figure 3.27: Measured displacement sensitivity curve of QUEST’s power-recycled Michelson interferometer named AMY.

3.9 Summary and discussion

QUEST is a pair of co-located power-recycled Michelson interferometer which are mounted on an optical bench that is seismically isolated via air suspension. The current configuration is that of an overcoupled cavity with a Finesse of approximately 1250. This yields a power recycling gain of around 780 (not yet measured). CARM is locked using a Moku:Lab and Pound–Drever–Hall reflection locking where the laser is the actuator, while DARM is locked by maintaining a DC-offset at the anti-symmetric port where the ETM PZTs are the actuator.

There is still a lot of commissioning to be done. The goals for QUEST include 10 W input power, an indefinitely and fully automated stable DC-lock capable of maintaining the desired offset, ~ 10 kW of circulating power and the injection of squeezed states of light. The angular to longitudinal coupling of the ETM PZTs is believed to a major hurdle which will prevent any other progress being made, with regards to the power-recycled Michelson interferometer lock. BOB is currently being used to commission the high power amplifiers which seems to be progressing straightforwardly. But without the PZT coupling problem solved, the lock will not be stable and high circulating power cannot be achieved.

The highest circulating power calculated to date has been ~ 142 W with a measured input power of ~ 357 mW — showing an achieved power-recycling gain of around 400. With the optics in place, a gain of ~ 840 is expected, meaning we should measure a circulating power of ~ 300 W with a stable, low-loss lock. The calculated power is under half of that expected, highlighting how lossy the current instrument

is. The design gain is > 1000 , but this does include a slightly more reflective power-recycling mirror than that currently in use.

Chapter 4

The Output Mode Cleaner

This chapter will motivate and describe the Output Mode Cleaner (OMC) of QUEST, a compact triangular cavity at the anti-symmetric port. The OMC was a largely individual project and one of the primary components required for QUEST to achieve its projected sensitivity.

The OMC is required due to the high circulating power of QUEST. This high power is expected to cause a thermal lensing effect on the interferometer optics and additional higher order modes will be generated within the interferometer as a result [48]. Higher order mode light at the anti-symmetric port will quite readily surpass the power saturation limits and could damage the detection photodetectors, but the OMC will act as a filter. It will suppress noisy light in the form of higher order modes by \sim two orders of magnitude, thereby removing that power from that which falls on the photodetector sensors. Without the OMC, the circulating high power could not be exploited.

Those with a background in the design of gravitational wave detectors and similar experiments will likely be quite familiar with OMCs. They are somewhat ubiquitous and perhaps even expected as a part of the layout. The major unique properties of QUEST's OMCs are the comparatively huge bandwidth, compactness and lack of suspended optics. Some core properties are given in Table 4.1 to give a brief introduction of the OMC and how it compares to the major gravitational wave detectors.

	QUEST	LIGO (A+) [89]	Virgo [150]	GEO600 [151]	KAGRA [152]
Bandwidth	(218.5 ± 8.4) MHz	678.97 kHz	75 MHz	2.93 MHz	256.23 kHz
FSR	(8 ± 0.4) GHz	264.83 MHz	4.16 GHz	454.23 MHz	199.86 MHz
round-trip	(37.5 ± 2) mm	1132 mm	72 mm	660 mm	1500 mm
Finesse	36.4 ± 2.4	390	50	155	780

Table 4.1: OMC properties for QUEST and gravitational wave detectors.

Each of QUEST's power-recycled Michelson interferometers will have a dedicated OMC. This chapter refers to a single OMC since only one has been constructed and characterised at the time of writing. Mirrors for 10 complete OMCs were purchased. When more are put together, they will be identical (in design at least) and will be characterised and controlled using the same methods as will be described.

Photos of the OMC at the time of completing this PhD are in Fig. 4.1 and some higher modes in transmission of the OMC are in Fig. 4.2




Figure 4.1: Photos of the OMC. (a) A British pound coin is included for scale.




Figure 4.2: Higher order modes in transmission of the OMC. These are screenshots of a video recording of a laptop screen — hence the poor quality. (a) The fundamental TEM_{00} mode. (b) The TEM_{10} mode. (c) The TEM_{02} mode. (d) The TEM_{22} mode.

4.1 Motivation for the OMC

The requirement for an Output Mode Cleaner is introduced in Sec. 2.8 but a more thorough argument is given here.

As described in Sec. 1.4, differential changes in Michelson interferometer arm lengths are measured by power fluctuations at the anti-symmetric port. The lower limit for the sensitivity to those length changes an experiment can achieve depends on any irreducible errors in the stability of the power measured where no length changes occur. Part of the power measured is noise and a suppression in that noise improves the experiment’s SNR. Photon shot-noise (Sec. 1.4.2) is the high frequency sensitivity floor for Michelson interferometer based experiments. Given QUEST’s bandwidth

of 1 – 250 MHz [48], it is the only quantum noise we expect to be limited by. Though shot-noise can be somewhat mitigated by for example higher power (Sec. 1.4.2) and squeezing (Sec. 3.4.13), it remains a fundamental sensitivity limiting noise source.

A common method of improving the sensitivity is to include the power-recycling mirror, creating the power-recycled Michelson interferometer configuration described in Sec. 2.5; this improves the SNR by increasing the circulating power per Eq. 1.36. This power enhancement is also true of any higher order modes present in the input beam, and higher order modes created via imperfections in the Michelson interferometer optics. The power-recycled Michelson interferometer configuration also excites additional higher order modes by virtue of the power-recycling cavity; the generation mechanisms are described in Sec. 2.3.7. As a brief recap, poor alignment and mode-matching leads to power in the fundamental mode coupling to non-fundamental modes. The power coupled to them can be reduced with careful construction and control schemes, but a certain amount is unavoidable. These kinds of higher order modes are expected to be mostly reflected by the power-recycling cavity though. The higher power also leads to a thermal lensing effect [48]. The optics of the power-recycled Michelson interferometer are heated and their properties are slightly changed; the beam splitter in particular, since the full circulating power passes through it. But the end test masses will also be effected in their radius of curvature, for example. This will not only cause fluctuations in the mode-matching requirement (which is a common-mode effect and so will be largely reflected from the power-recycling cavity), but also generate light in the form of differential-mode higher order modes, which will couple to the anti-symmetric port. The combined power of these higher order modes serves only to reduce the SNR.

Sec. 2.8 discussed that for a successful detection of faint signals the power in the local oscillator must exceed that of the contrast defect. The contrast defect is the light which leaves the anti-symmetric port due to imperfections in the Michelson interferometer optics. It is a combination of fundamental mode which did not destructively interfere at the beam splitter and of higher order modes.

The design goal for QUEST based on the desired SNR, was given in Sec. 3.4.11; $< 10^{-6}$ for fundamental mode and $< 10^{-5}$ for higher order modes [48]. The higher order mode requirement is less stringent than the fundamental mode one because of the ability to reject this light in the detection path — the fundamental mode contribution can only be reduced with high quality, well cleaned optics. Once the instrument is constructed it is not something that can be easily improved, particularly if using vacuum chambers where opening those chambers is reserved for matters of absolute necessity. Without addressing the higher order mode content though, that would lead to a premature limit to the circulating power for the sake of the photodetectors. With higher circulating power, the DC-offset and contrast defect

(including the higher order modes) powers all increase as well. The combination at the anti-symmetric output will quickly saturate the high-bandwidth low-noise photodetectors. This was one of the limiting factors for the Fermilab Holometer [62] specifically because they did not have any way of suppressing the higher order mode power and so they needed to use a large DC-offset to ensure signal carrying light would be dominant. They were at the photodetector saturation limit as a result of this which restricted their ability to increase the circulating power to improve their SNR.

An OMC is an optical cavity positioned between the anti-symmetric port and the photodetectors, with the direct purpose of filtering the beam to suppress light that is not fundamental mode — see Fig 3.8. As detailed in Sec. 2.3.5, when an impedance matched cavity, in this case the OMC, is made resonant with and thus transmits the fundamental mode, it will act as a reflector for higher order mode fields preventing them from reaching the photodetectors. The fundamental mode contrast defect will unfortunately pass through the OMC unhindered, since it is the same frequency. However with the design goals of QUEST, this power will be very small [48] and should not present limiting technical issues.

The inclusion of an OMC facilitates two features of QUEST, both of benefit to the overall sensitivity:

1. The circulating power can be increased (with higher input power and a higher Finesse recycling cavity) without risking prematurely saturating the photodetectors, since higher order mode power will be mostly rejected.
2. Since the intention is that the fundamental mode contrast defect power is much lower than the higher order mode power, the DC-offset can be smaller because in transmission of the OMC, the DC-offset only now needs to dominate the fundamental mode contrast defect, not the higher order modes as well. This is depicted in Fig. 2.25b.

In principle, whatever higher order mode suppression the OMC establishes implies we can tolerate an equivalent reduction in the amplitude of the DC-offset — which will reduce loss in the power-recycled Michelson interferometer, maintaining higher circulating power.

The OMC is one of the major upgrades to the most comparable experiment to QUEST, Fermilab’s Holometer [62]. QUEST’s expected sensitivity relies on its inclusion and as such, it is a principal component to the design. This chapter is dedicated to detailing the design and implementation of the OMC from the fundamental requirements to the final characterisation. The lock process and how it behaves as a stand alone device will also be thoroughly described. Secs. 3.5.2 and 6.1.2 show how it will be integrated into QUEST and how it is expected to behave

as part of the whole apparatus.

4.2 Design considerations

QUEST's targeted sensitivity to differential-mode arm length fluctuations extends to 250 MHz [48], meaning that requires all components involved to have a sufficient bandwidth. The OMC is no exception to this and the original goal was to achieve a bandwidth of ≥ 200 MHz. This serves as an adequately high bandwidth, with any attenuation at the higher end being tolerable and recovered through simple filters. The other principal requirement is to suppress higher order spacial modes by at least an order of magnitude with respect to the fundamental mode.

The design choices made to achieve these two fundamental requirements will be discussed in detail in the following sections.

4.2.1 Geometry

A linear (Fabry-Pérot) cavity can be used for a mode cleaner but it is undesirable since the reflection is directed back towards the source. A non-normal angle of incidence on the input mirror of the mode cleaner is preferable since it directs the reflections away from the source. Two common geometries with this feature are the triangular and bow-tie [89] cavities.

The triangular cavity generally has flat input and output mirrors, with a single curved mirror to complete the triangle and focus the beam back to the waist¹. A bow-tie has four mirrors, either comprising two flat and two curved mirrors, or three flat and a single curved mirror. At least one curved mirror is required to focus the beam back to the cavity waist and create a stable eignemode.

The decision on the geometry was made for the sake of the physical size of the cavity. The bow-tie provides approximately double the round-trip-length for a similar bench footprint as the triangular. As will be described shortly, the round-trip-length is on the limit of what three separate optics can reasonably facilitate, and so the bow-tie was not considered.

4.2.2 Finesse

Finesse and its relationship with higher order mode suppression was introduced in Secs. 2.3.3 and 2.3.5, respectively. But that description was in terms of a linear cavity specifically. For a triangular cavity with one curved mirror, further components must be included.

¹It could have two or three curved mirrors, but the configuration of flat input and output mirrors with a single curved mirror makes linearising the cavity much simpler. Via symmetry, it can be equated to a linear cavity of half the round-trip-length.

Firstly, there is a single curved mirror at a non-normal angle of incidence which generates an astigmatism in the cavity eigenmode. With respect to the curved mirror, for S-polarised light as will be used here, the plane of incidence of the beam is called the sagittal plane (y, z) and the plane perpendicular to it is the tangential plane (x, z). For the tangential plane the radius of curvature of the mirror is

$$RoC_t = RoC_C \cos(\beta_C) \quad (4.1)$$

and for the sagittal plane it is

$$RoC_s = \frac{RoC_C}{\cos(\beta_C)} \quad (4.2)$$

where RoC_C is the radius of curvature of the curved mirror and β_C is the angle of incidence of the beam on that mirror [67, 78]. From Eq. 2.73, where the cavity waist is derived from the cavity geometry, it is possible to see that since each plane has a distinct radius of curvature, the waist requirement is different for the planes. This astigmatism results in each plane having separate eigenmodes and thus different mode-matching requirements.

Secondly, there are an odd number of mirrors meaning an odd number of reflections per round-trip in the triangular cavity. The distributions of Fig. 2.3 shows the mode field amplitude is flipped from lobe to lobe. Since reflections cause flips about the sagittal plane only [90], then modes with amplitude distributions asymmetric to the sagittal axis are effected and acquire an additional $\pi/2$ phase shift in their round-trip [90, 153]. If $p = n + m^2$, then modes of the same value p are only degenerate or co-resonant (excluding the aforementioned astigmatism) where n is even because in that case the field amplitude pattern is symmetric with reflection.

The number of nodes splitting the mode in the x, z (tangential) and y, z (sagittal) planes are given by the subscripts n and m , respectively with $n, m \in \mathbb{Z}^+$.

With both of these factors, Eq. 2.52 from Sec. 2.3.5 (spacial mode filtering of a linear optical cavity):

$$\frac{T_{nm}}{T_{00}} = \frac{1}{\left[1 + \left(\frac{2}{\pi} \mathcal{F} \sin\left(\frac{2\pi}{L} \Delta\nu_{nm}\right)\right)^2\right]^{1/2}} \quad (2.52)$$

is modified for the triangular cavity. Now, the higher order mode transmission T_{nm} relative to the fundamental mode transmission T_{00} is

$$\frac{T_{nm}}{T_{00}} = \frac{1}{1 + \left(\frac{2\mathcal{F}}{\pi} \sin\left[\left\{n \arccos(\sqrt{g_t}) + m \arccos(\sqrt{g_s})\right\} + \frac{\pi}{2} \frac{(1-(-1)^n)}{2}\right]\right)^2}. \quad (4.3)$$

²As introduced in Sec. 2.2, n and m are the number of mode nodes in the tangential and sagittal planes, respectively.

Eq. 4.3 is a modified version of equation 14 in [90] which shows the field amplitude while Eq. 4.3 is power, and the cavity g-factors corresponding to their relevant plane g_t and g_s (defined below) have been separated. This separation accounts for the astigmatism introduced by the non-normal angle of incidence on the curved mirror. \mathcal{F} is cavity Finesse.

g_t and g_s are the tangential and sagittal cavity g-factors, respectively, modifications of the g-factor defined in Eq. 2.44.

$$g_t = 1 - \frac{L}{RoC_C \cos(\beta_C)} \quad ; \quad g_s = 1 - \frac{L}{RoC_C / \cos(\beta_C)} \quad (4.4)$$

A simplified version of Eq. 4.3 could reduce the term in curly brackets to: $\{(n + m) \arccos(\sqrt{g_C})\}$, where g_C is the average cavity g-factor which does not account for the astigmatism of the planes. However, the astigmatism has the effect of separating modes of the same order — neglecting this could result in greater higher order mode transmission than expected.

The last term in Eq. 4.3's denominator involving the exponent n provides the $\pi/2$ phase shift for the odd numbered n^{th} modes.

A comparison of the simplicity of Eq. 4.3 is shown below in Fig. 4.3. Fig. 4.3a is Eq. 4.3 with the simplified g-factor and does not account for the asymmetric reflections. Fig. 4.3b includes the astigmatism, separating g_t and g_s , but does not account for the asymmetric reflections. Fig. 4.3c is the complete representation of Eq. 4.3 — the importance of including astigmatism and asymmetric reflections is clear, the shift of the odd n^{th} order higher order modes, in particular.

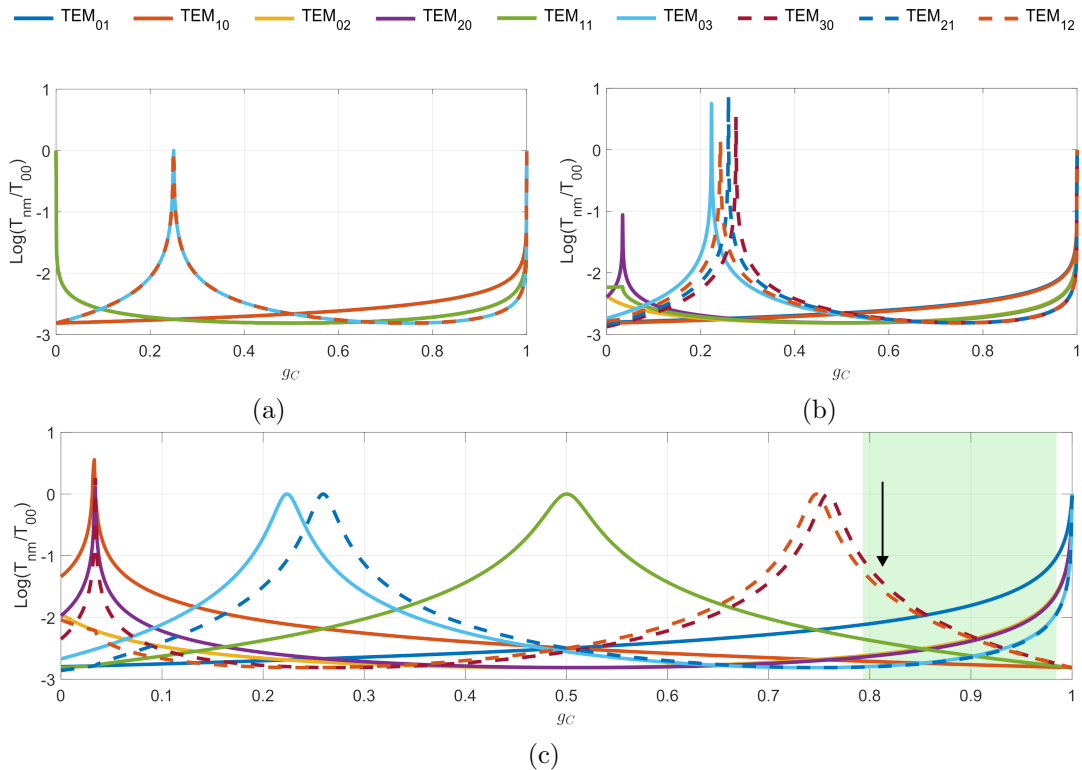


Figure 4.3: Higher order mode transmission comparison. In (a) the average cavity g -factor is used and parity is not accounted for. The 1st, 2nd and 3rd order modes all form a single respective line and could lead one to believe the transmission will be far less than is the case. (b) and (c) show higher order mode transmission with the angle of incidence on the curved mirror = 15° . (b) only accounts for astigmatism. (c) accounts for astigmatism and asymmetric reflection, this is the more accurate representation of higher order mode transmission. The green region in (c) highlights the right most g -factor range (~ 0.793 to 0.985) which keeps transmission below an order of magnitude and the arrow is the g -factor that was chosen. The decision was based on manufacturing tolerances and costs of the mirror curvature.

These plots identify useful g -factor choices — peaks are where the fundamental mode co-resonates with the specific higher order mode(s), and low points are minimal higher order mode co-resonance. $\mathcal{F} \sim 11$ is a minimum in order that low points provide an order of magnitude suppression but $\mathcal{F} = 40$ was decided as a good benchmark since this yields \sim two orders of magnitude higher order mode suppression with an appropriate g -factor. As described above, the Finesse and pre-determined ≥ 200 MHz bandwidth give the required cavity length (Eqs. 2.47 and 2.48) while the chosen g -factor gives the radius of curvature of the curved mirror.

\mathcal{F} is physically determined by the mirror reflectivities per Eq. 2.50. For the three mirror cavity, this is specifically [78]

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2 r_3}}{1 - r_1 r_2 r_3}, \quad (4.5)$$

where r_1 , r_2 and r_3 are the field amplitude reflectivities of the three cavity mirrors.

Round-trip-length and radius of curvature

In a linear cavity, the round-trip-length L_{rt} is two times the cavity length, i.e. $L_{rt} = 2L$. For the triangular cavity the round-trip defined by the triangular circuit, and the *length* of the triangular cavity is half of that (so $L_{rt} = 2L$ still holds), measured from midway between the flat input/output mirrors — where the cavity waist will be — along the beam path to the concave mirror; see Fig. 4.4.




Figure 4.4: The lengths of the linear and triangular cavities.

Since ν_{cav} is a prerequisite determined by QUEST’s design sensitivity, the choice of $\mathcal{F} = 40$ establishes the required $L_{rt} = 37.5\text{mm}$ by Eqs. 2.47 and 2.48. With this, appropriate transmission low points and corresponding g-factors provide choices for radius of curvature by Eq. 2.44.

The right-most low point of Fig. 4.3c (g-factor = 0.856, within the green region of the figure) was chosen as the goal, giving us a suitable g-factor in the range of 0.793 — 0.985 for one order of magnitude suppression. These values correspond to a concave mirror radius of curvature between 0.091m and 1.249m with an ideal radius of curvature of 0.130m. $RoC = 0.1\text{m}$ was chosen for the sake of manufacturing costs and tolerances (indicated by the arrow of Fig. 4.3c).

Mirror coatings

With the Finesse selected, Eq. 4.5 establishes mirror reflectivity. The concave mirror would ideally be 100% reflective but there are manufacturing and budget constraints — so $r_3^2 = 0.9999$ was chosen.

The critical coupling configuration³ is the most suitable for the mode cleaner application since it provides maximum transmission — see Fig. 2.8 blue curve. This means that the input and output mirror reflectivities ought to be equal, resulting in $r_1^2 = r_2^2 = 0.925$, for $\mathcal{F} = 40$.

Manufacturers can only guarantee a certain portion of the mirror’s surface. There are

³The coupling configurations introduced in Sec. 2.3.3 were in terms of a linear cavity. For this example of a triangular cavity, where the concave mirror is made as reflective as possible, it is only the input and output mirror which are used to determine the cavity coupling.

two factors, the surface quality and the coating. Surface quality refers to how rough the surface is — how many ‘scratches’ and ‘digs’ there are [154]. More perfect surface quality reduces scattering. Coating is then the reflectivity, which can be specified for power reflectivity of a particular wavelength, angle of incidence and polarisation. Laseroptik [155] were used for the mirrors of the OMC and they guarantee 80 % of the mirror’s surface will be to specification.

4.2.3 Angle of incidence

The most important angle of incidence (AOI) is that of the curved mirror⁴ and there are two competing factors which must be considered, astigmatism and the counter rotating field. Increasing the angle of incidence leads to greater astigmatism and higher order mode transmission since any degenerate higher order modes now split and co-resonate with the fundamental mode at slightly different cavity g-factors. The BRDF determines the amount of scattered light in a given direction, it can be used to estimate back-scatter [156, 157] which is particularly important. It generates a counter-rotating field within the cavity. Back-scatter decreases with increasing angle of incidence — as shown in Fig. 4.5 in which a clearly defined desirable lower bound of $\sim 5^\circ$ is visible.




Figure 4.5: BRDF vs. angle for a typical Newport SuperMirrorTM[158, 159]. The original data could not be found, this plot is data extracted using an online algorithm (WebPlotDigitizer) [160] from a picture in Newport’s brochure [158], which is also figure 3 in [159].

The power in this counter rotating field is proportional to the BRDF, as [159]

$$\frac{P_{cr}}{P_{inc}} = \frac{\text{BRDF}(\beta_C)}{(1 - r_{1,2}^2)^2} \left(\frac{\lambda}{L}\right) \sqrt{\frac{1 - g_C}{g_C}}. \quad (4.6)$$

⁴This is because flat mirrors do not alter the beam other than its propagation direction. However, a curved mirror also changes the beam’s wavefront radius of curvature and divergence/convergence.

Where β_C is again the angle of incidence on the curved mirror. An angle of incidence just over 5° would have been ideal but $\beta_C = 15^\circ$ was chosen in order to prevent the beam impinging close to an internal mirror edge, as described in the following section.

4.2.4 Cavity Layout

To prevent the reflection beam from impinging near the corner or un-coated edge (introduced in Sec. 4.2.2) of the input mirror, the input and output mirrors were positioned as in Fig. 4.6. Given the size of the cavity, the proximity of the mirrors lead to the wedge design of the output mirror. An appropriate wedge was cut from one edge in order to hold the mirrors together — this has the added benefit of facilitating their gluing together if necessary.

In the process of design, these relationships for lengths and angles as a function of the angle of incidence on the curved mirror (β_C) were derived — they are identified in Fig. 4.6.

$$x = \frac{L_{rt} \sin(\beta_C)}{1 + \sin(\beta_C)} \quad ; \quad y = \frac{L_{rt}}{2} \left(1 - \frac{\sin(\beta_C)}{1 + \sin(\beta_C)} \right) \quad ; \quad k = \frac{x}{2 \cos\left(\frac{90+\beta_C}{2}\right)} \quad (4.7)$$

$$a = 90 - \beta_C \quad ; \quad b = \frac{90 + \beta_C}{2} \quad ; \quad c = \frac{90 - \beta_C}{2} \quad (4.8)$$






Figure 4.6: OMC angles and lengths. L is the length of the rectangular mirrors, 25.4mm. w is the thickness of the rectangular and concave mirror, 6mm. RoC is the radius of curvature of the concave mirror. x and y are the parameters of the triangular path the beam takes ($L_{rt} = x + 2y$). Green shows the guaranteed surface area of the mirror, red is the safety region to avoid. $L - w \cdot \tan(\beta_C)$ is the full length of the inside edge of the output mirror, not just the green section.

The green in Fig. 4.6 displays the *clear aperture*, it is the region of the optic guaranteed by the manufacturer to be at the specified values for reflectivity and surface quality. For these optics it was quoted as being 80% of the surface. The red regions of Fig. 4.6 (and Fig. 4.7) are outside of the clear aperture, this represents the surface of the mirror that is not guaranteed to be at the specified reflectivity and surface quality — the beam should be prevented from impinging in these regions. A perimeter of 10% the surface length was assumed to be *out of bounds* — the size of the clear aperture is represented here as $0.8[L - w \cdot \tan(\beta_C)]$.


In deciding the angle of incidence on the curved mirror, the distance of the beam to the red region shown in Fig. 4.6 was a deciding factor — with $\beta_C < 5.6^\circ$ the beam impinges within this region. The distance inside the clear aperture is shown in Fig. 4.7a.




(a) The distance the beam will impinge from the clear aperture perimeter as a function of β_C . This is $ap = k - 0.1[L - w \cdot \tan(\beta_C)]$ where k is the length defined in Fig. 4.6.



(b) Beam size when propagated $x/2$ from the waist position. Blue is the projection of the spot radius on the mirror surface,
 $\omega_p = \omega_{x/2} / \cos([90 - \beta_C]/2)$



(c) Blue is a simplified spot size, assuming a radius equal to ω_p . The black ring is an aperture with radius equal to ap .



(d) Clipping as a function of β_C ,
 $\exp(-2ap^2/\omega_p^2)$

Figure 4.7: Angle of incidence on the concave mirror, and how it influences spot position within the OMC. Green and red represent the regions of manufacturer guaranteed surface finish and coating specification, and regions which are not guaranteed and should be avoided, respectively.

For each β_C , the spot radius $\omega_{x/2}$ at the mirror was calculated. The projection of this spot size on the mirror surface (blue in Fig. 4.7b) is $\omega_p = \omega_{x/2} / \cos([90 - \beta_C]/2)$. It was assumed that the distance from the beam to the edge of the clear aperture could be approximated as the radius of an aperture through which the beam propagates, having radius $ap = k - 0.1[L - w \cdot \tan(\beta_C)]$ where k is the length defined in Fig. 4.6. With this, the amount of power lost to clipping due to proximity to the edge can be approximated as $\text{clip loss} \sim \exp(-2ap^2/\omega_p^2)$.

By $\beta_C \sim 6^\circ$, clipping falls to 0% (see Fig. 4.7d) — for an extreme safety margin⁵,

⁵The safety margin being so large was motivated by the fact mode-matching (and alignment) are tricky, particularly for such a small cavity waist. This analysis was carried out assuming idealised

$\beta_C = 15^\circ$ was chosen. By Fig. 4.6, this means there must be a 15° wedge cut from the output mirror in order that they form the correct inter-mirror angle.

4.2.5 Vacuum requirement analysis

This analysis aims to assess whether the OMC should be housed in vacuum. When not in vacuum, changes to the refractive index of the air filling the OMC could cause a lock loss if the round-trip-length is sufficiently altered. If so, there is a case to be made that the OMC warrants its own vacuum chamber. The following does not account for any feedback or actuation, it merely serves to model the required refractive index change to drive a sufficient round-trip-length perturbation for a lock loss.

As a simplified model, assume the cavity will lose lock with a round-trip change equivalent to half a bandwidth, the maximum tolerable δL_{rt} is computed as $1/(2\mathcal{F}) \cdot \lambda/2 = 0.0125 \cdot 532 \text{ nm} = 6.65 \text{ nm}$. The refractive index change required to result in δL_{rt} is: $\delta n = (L_{rt} + \delta L_{rt})/L_{rt} = 1.0000002$ — i.e. if the refractive index changes by $2 \cdot 10^{-7}$ it is likely that the OMC will lose lock.

The Edlén equation

The Edlén equation [161] describes the refractive index of air as a function of temperature, pressure, humidity, CO₂ concentration and radiation wavelength. A Matlab script available online [162] models this equation, with inputs given in units of °C, kPa, percentage, ppm and nm, respectively.

Change in pressure

Leaving all other parameters fixed, a plot with changing pressure shows that $\delta n = 2 \cdot 10^{-7}$ results from a pressure variation of 7.5 Pa (see Fig. 4.8a) which is roughly equivalent to an altitude change of 6 m.

An engineering calculator [163] gives 131.5 dB as a sound required to generate such an air pressure fluctuation. A jet engine taking off or an air raid siren are ~ 130 dB — so, other than perhaps the fire alarm, it is highly unlikely to have a sound causing sufficient refractive index fluctuation in the QUEST lab.

Change in temperature

The same analysis was carried out for the temperature, showing that $\delta n = 2 \cdot 10^{-7}$ results from a 0.109 °C change in temperature (see Fig. 4.8b). This appears quite concerning because though the temperature of the lab is controlled it seems a small drift is sufficient to cause lock loss. It is however, not something which fluctuates

conditions, in reality, the waist size and position will not be perfect and so some forgiveness in these kinds of requirements is beneficial where it is not too detrimental otherwise.

quickly and constantly. There may be slow drifts throughout the day, but this is nothing a good control scheme can not respond to efficiently.




Figure 4.8: Index of refraction of air changing with (a): pressure, (b): temperature.

Fabry–Pérot test case

One of the prototypes to QUEST was an in-air Fabry–Pérot cavity (described in Sec. 3.3.2) with which an *indefinitely* stable lock was achieved. Via a length scan, the cavity was measured to have $\mathcal{F} = 611.2 \pm 7.3$ with $L_{rt} = 3$ m, which results in $\delta L_{rt} = 4.4 \cdot 10^{-10}$ m and $\delta n = 1.5 \cdot 10^{-10}$.

It seems reasonable to expect the OMC lock to be capable of withstanding any refractive index and pressure-induced vibration fluctuations of the air, particularly since its Finesse is over an order of magnitude lower than this Fabry–Pérot cavity.

4.2.6 Housing

The round-trip-length of the OMC is small enough to require the input and output mirrors be in contact. The mirrors forming a symmetric shape risks the reflection beam exiting through the corner/edge of the input mirror. To prevent this, the slightly offset design as in Fig. 4.6 was decided.

The wedged mirror was intended to facilitate gluing them together forming a rigid mirror pair. However, the housing designed to fix them in place as the glue cured (Fig. 4.10) was realised to be an ideal permanent housing, removing the gluing requirement. With this housing, the concave mirror is on a separate mount glued to the actuator; more information about the actuator will be given in Sec. 4.2.7.

During commissioning, the housing design was *upgraded* to form more of a monolithic structure — the principal motivation for this was for the sake of ensuring producing multiple OMCs with very similar round trip lengths. With the two Michelson interferometer configuration(s) of QUEST, it is of course preferable to ensure spec-

ifications are as equal as possible for all components. The design of two separate platforms, positioned by eye, lends itself to unequal round trip length between two OMCs; which has a large bearing on their bandwidths. If Finesse is unchanged, then per Eqs. 2.47 and 2.48,

$$\delta\nu_{cav} = \frac{\delta\nu_{FSR}}{\mathcal{F}} = \frac{c}{\delta L_{rt}\mathcal{F}}. \quad (4.9)$$

This relationship is plotted in Fig. 4.9.




Figure 4.9: Fluctuation in OMC bandwidth as a function of % round-trip-length.

A monolithic cavity in which uncertainty comes from manufacturing tolerances of the components and the thickness of glue layers, is a far more robust method of installing uniform OMCs. This was realised with the updated housing design shown in Fig. 4.11; dimensions are given in Appendix C.3.




Figure 4.10: OMC original housing. This was designed initially as a method of holding the flat mirrors together while glue cured. They sit in shallow grooves with the rectangular input mirror held in place against the L shaped post, and the wedge shaped output mirror held against the input one, by nylon tipped screws through the rectangular posts. The lines etched into the base were markers for the input/output beams, and to aid positioning the concave mirror. The groove cut out of the rear post is for unobstructed propagation of the reflected beam — removing a large source of scattering.




Figure 4.11: OMC updated housing. This was designed in order to establish a reliably reproducible round-trip-length. The flat mirrors are held in place by nylon tipped screws. The concave mirror is glued to the PZT and spacer, the spacer is then held in place by a nylon tipped screw. Dimensions are in Appendix C.3.

4.2.7 The OMC actuator

The actuator for the OMC is a PA44M3KW ring PZT from Thorlabs [164], shown in Fig. 4.12, it was chosen for its displacement range and the wide central aperture. Dimensions for the gluing jig are in Appendix C.2.




Figure 4.12: The actuator for the OMC, the PA44M3KW ring PZT from Thorlabs [164]. 15 mm outer diameter with 9 mm central aperture. $3.9 \mu\text{m}$ ($\sim 7 \nu_{FSR}$) displacement range. (a) is from the Thorlabs website. (b) and (c) are photos during the gluing process. The metal ‘jig’ (dimensions are in Appendix C.2) was designed to keep the mirror, PZT and brass spacer commonly centralised while the glue cured. Thin ‘arms’ of the jig are to prevent it from touching the glue, to ensure it can be removed once the glue is cured. The pink surface of the mirror is ‘First Contact’ [165], an adhesive fluid from used for cleaning and protecting the mirrors.

The PZT has a range of $3.9 \mu\text{m}$, which covers approximately $7 \nu_{FSR}$ and a 9 mm diameter central aperture — a wide aperture is required if light transmitting through the curved mirror is to be analysed. The PZT has a natural resonance frequency of 260 kHz with no load [164]. The data for load resonance [164] is in Fig. 4.13. The OMC’s concave mirror will be glued to the PZT. It has a measured mass of $\sim 1.9 \text{ g}$ suggesting the actuator’s resonant frequency under this load is $\gtrsim 160 \text{ kHz}$. The CDS has a bandwidth of 16 kHz so there is no concern about reaching the actuator’s load resonance to first order.




Figure 4.13: The PZT resonance as a function of load [164]. (a) is full data, (b) is the region of interest. There are only two data points available for this load range [164], however the value for a 1.9 g load lies between them.

The PZT is a high voltage actuator, requiring > 25 V to cover $3\nu_{FSR}$ ($\sim 1.6 \mu\text{m}$) of the OMC. This means it must pass the CDS output through a high voltage amplifier (described in Appendix A.1). The high voltage amplifier output has the added benefit of there being a 50 V DC output, meaning the PZT can be both expanded and compounded from this ‘starting’ position. Displacement vs volts data from Thorlabs [164] is shown in Fig. 4.14, it shows a clear hysteresis. This means that for any characterisation length scans, converting from volts to meters may not be straightforward. The workaround of this hysteresis is described in Sec. 4.7.2.




Figure 4.14: PZT hysteresis curve [164].

4.2.8 Design specifications summary

Sec. 4.2 has detailed the predominant design considerations and the choices made to facilitate the requirements of the OMC. Photos of the final design are in Fig. 4.1 and a summary of the major design specifications are in Table 4.2. Transverse mode spacing is the frequency between modes described in Eq. 4.22.

Bandwidth		200 MHz	AOI	Input/output	37.5°
Finesse		40		Concave	15°
FSR		8.0 GHz	g-factor	Mean	0.8126
round-trip-length		37.5mm		Tangential	0.8060
Reflectivity (R)	Input/output	0.925		Sagittal	0.8190
	Concave	0.9999	Waist		115 μ m
Concave radius of curvature		0.1m	Transverse mode spacing		1.1 GHz

Table 4.2: OMC design specifications.

4.3 Detection breadboard

The intention with the detection breadboard layout is for it to house all components relating to detection (other than the mirrors directing light from the anti-symmetric port to the board). This is motivated by the fact two detection paths will be present and future configurations may require them being moved on the bench. Having all components together should result in a strongly reduced effort for future updates. The only physical changes should be the mode-matching telescope lenses, and then alignment can be re-achieved with the dedicated steering mirrors which will not have moved.




Figure 4.15: Detection breadboard. The green area shows the permitted regions for the mode-matching telescope lenses. The mirrors leading to the OMC outside of the green area are the steering mirror. Mode-matching lenses are not permitted downstream of the first steering mirrors to prevent deformations as a result of the beam being moved through a lens. There is a reflection photodetector positioned close to the OMC. The transmission is split 50/50 at the beam splitter between two detection photodetectors. Each photodetector has its own steering mirror and lens to focus to the sensor. The breadboard itself is 45x45cm with a 2.5cm separated grid of M6 tapped holes.

Mode-matching telescope solutions are generated as described in Sec. 2.4, where, for an initial attempt, the green regions of Fig. 4.15 are the only spaces permitted for the lenses to be placed. This is simply to attempt having them on the breadboard. If no telescope solutions are possible under this condition, then the space between the anti-symmetric port and the board can be included. Originally, the full propagation distance, including after the steering mirrors was allowed, but lenses downstream of the steering mirrors means the beam will move around the lens, which is not ideal and can lead to deformations/astigmatism.

One additional requirement is for the steering mirrors to be as separate as possible. The more separation between them the better. By simple trigonometry their angular actuation increases linearly with the distance they are from the OMC, the more separation, the greater alignment control can be achieved.

Each detection breadboard will have two detection photodetectors for the sake of power saturation, noise reduction and redundancy. With two, the transmitted power of the OMC is split between them, allowing double the input power (which is the anti-symmetric port output power). It also provides the opportunity to characterise correlated and uncorrelated noise which will facilitate further noise reduction with time integration [89]. It finally also provides a redundancy. If one fails the other can at least serve to hold resonance even if data acquisition is temporarily paused.

4.4 OMC Control

The OMC is controlled using a dither lock. The process of the dither locking scheme is detailed in Sec. 2.10 but as a brief reminder; an excitation is injected to the OMC actuator to drive a small length fluctuation, the transmission of the OMC is then demodulated at the excitation frequency. The error signal is generated via this demodulation and the OMC length can be held at the peak of resonance. The CDS (Appendix A.1) is responsible for this control scheme. CDS is entirely digital; filters, signal generators, demodulation, analysis etc. is all done within the software.

In this section, the practical implementation of this dither scheme will be described.

There are two phases of the lock scheme.

1. Lock acquisition — including initial lock and automatic re-lock following lock loss. The feedback loop is *open* in this phase.
2. Feedback control — to maintain lock. The feedback loop is *closed* in this phase.

The screen shown in Fig. 4.16 was designed in order to depict the loop. This screen is an interactive tool which allows the user to update certain values, turn filters on/off and block/engage signal flow.

To the right of the left most ‘Sum Display’ box, there is a graphic of a transmission peak – the top \sim third is green and the lower part is red — it represents the split between the two phases. Above the green level, the OMC is close to resonance and feedback can be enabled (phase 1 of the above list). Below the green and the OMC is not on resonance, so a slow length scan is initiated to push the concave mirror close to a resonance (phase 2).

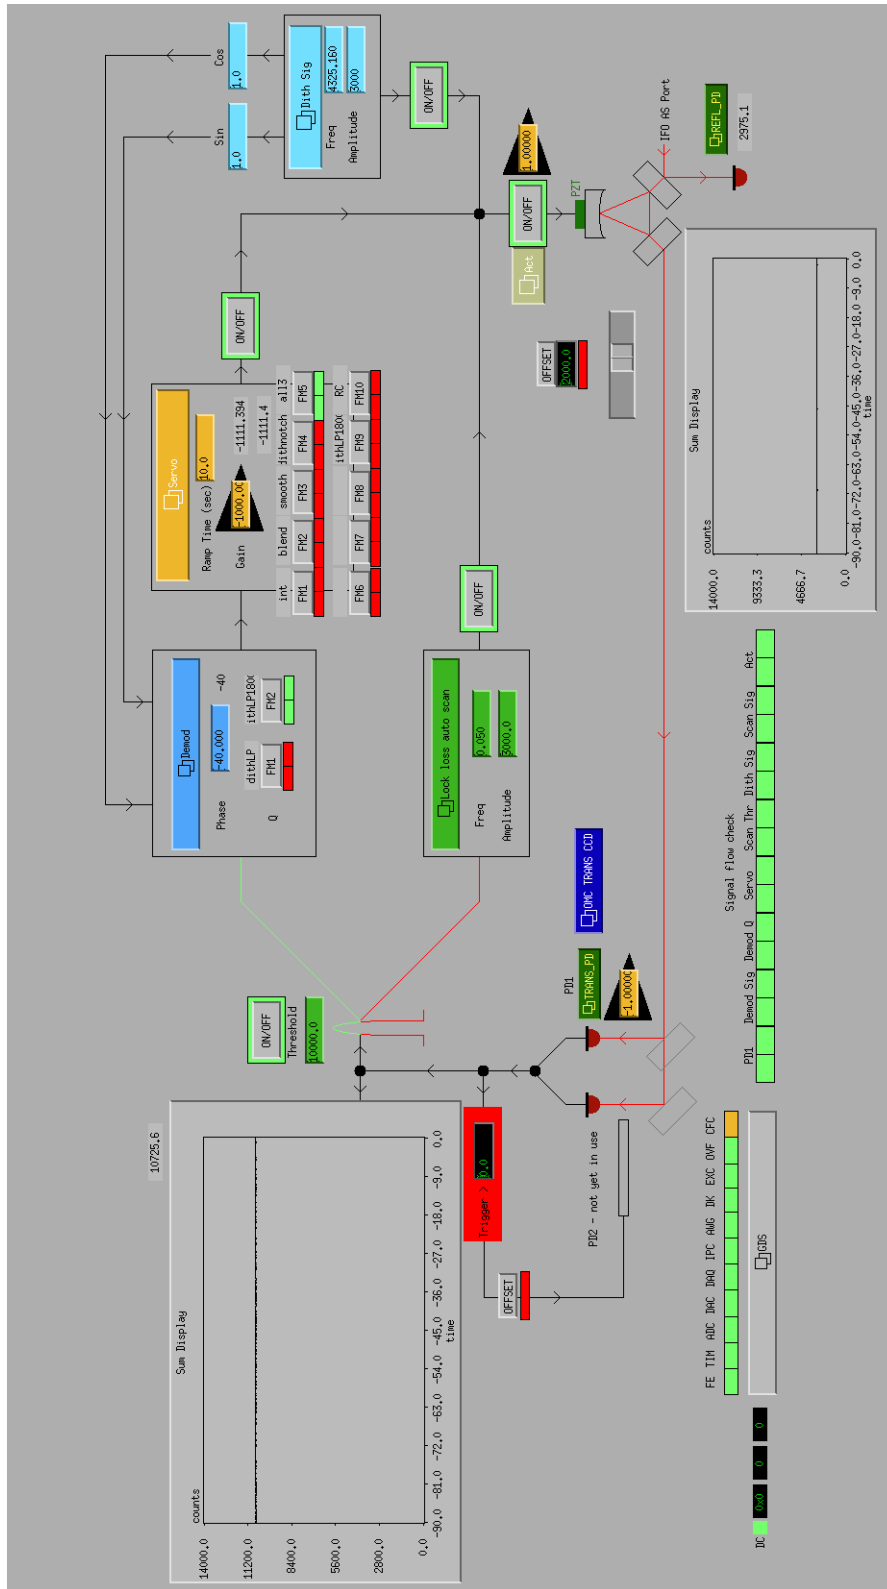


Figure 4.16: OMC Overview CDS screen.

4.4.1 Lock acquisition and auto-scan

Lock acquisition needs to be a somewhat automated process. If a user needs to be present to lock a device every time, experimental work would be extremely laborious and it would be slow progress indeed. The possible lock acquisition methods for a loop with a single actuator are limited to just one — in this case, a length scan; which is used to find the cavity resonance.

The red and green transmission peak graphic in Fig. 4.16 displays the condition for this length scan to kick in (a flow chart of this process is also in Fig. 4.17). If the transmitted power falls below some threshold (details in the following paragraph), the cavity is assumed to be unlocked. At this point the control signal is blocked from reaching the actuator (the loop is opened) and it is replaced by a low frequency, high amplitude sinusoid. The properties of this sinusoid depend on the loop’s ability to ‘catch’ a lock. This requires some element of trial and error.

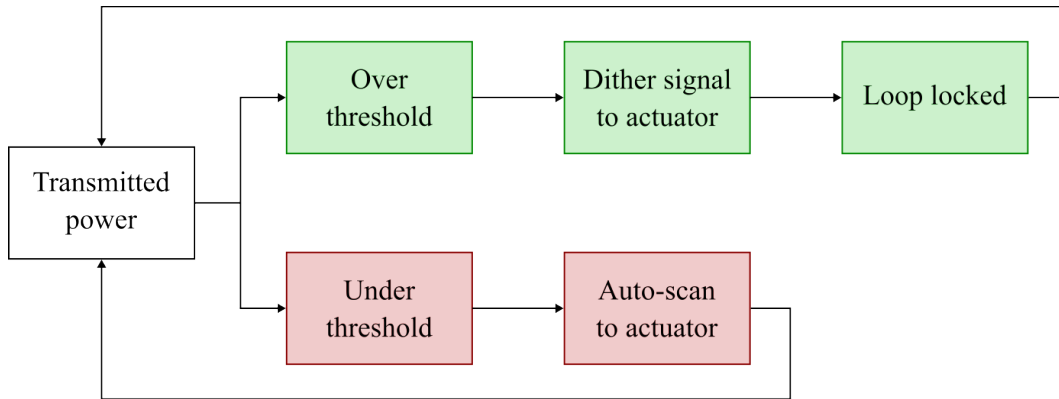


Figure 4.17: The OMC lock sequence as a flow chart. It is a simple logic switch. When the transmitted power is above the threshold, the dither signal is sent to the actuator and the lock loop is closed. When the transmitted power is below the threshold, the dither signal is blocked and the automatic length scan is sent to the actuator. The length is pushed until a resonance is reached and the transmitted power increases to the threshold. After which point the scan is blocked and the dither is re-engaged closing the lock loop.

The auto-scan frequency for this loop is 50 mHz and it covers approximately $1.5 \nu_{FSR}$. It needs to cover more than $1 \nu_{FSR}$ to ensure at least one transmission peak will be fully scanned. Catching lock is very sensitive to the threshold power level. If the threshold is too low then the dither will not cover the peak of the resonance, which is required for it to function (see Sec. 2.10). If it is too high then the loop may not be fast enough to catch it before it falls off the resonance peak. From trial and error the threshold generally works best at $\approx 90 - 95\%$ peak transmission. At the point of reaching the threshold, the loop switches back to the control signal and disengages the scan. The reason the loop can struggle to catch lock is because the scan is effectively finding a new set-point — the position of the mirror for the OMC round-trip to be correct for the ‘current’ frequency of the field. Once in the correct

position, the PZT needs to be held there and use that position as a set-point for the loop. Unfortunately, pausing the scan to hold the PZT's position, and have the control signal operate from that position is not a straightforward option and was not attempted. Instead, the control signal needs to be strong enough to catch that set-point when the scan is disabled.

For example, if during the auto-scan a PZT voltage of 25 V is required, once the transmission threshold is reached the scan is disabled and that 25 V would be removed from the PZT, pulling it away from the correct round-trip-length position. The control signal needs to be ready to apply that 25 V and maintain resonance from there. In a sense, it is like having a loop which needs to find its own set-point each time it is engaged. If the set-point is well defined and will not change, the error signal is driven to 0; but in this case, the error is driven to equal the set-point and fluctuations about that set-point are what is driven to 0.

4.5 Loop characterisation

This section describes some important stages of designing and then characterising the loop.

4.5.1 Loop electronics

Electronic circuits generate all sorts of poles and zeros (introduced in Sec. 2.6.1) owing to all the resistors, capacitors and other components present. A starting point for understanding how to design an appropriate servo for the desired open loop transfer function (which contains the electronic components) is to first characterise the electronics that will be involved in the loop. This will highlight any poles and zeros present due to the electronics that cannot be avoided, but might be unwanted in the open loop transfer function.

The electronic circuitry for the OMC lock loop's actuator consists of the CDS output, a high voltage amplifier and the OMC's PZT. The high voltage amplifier used to drive the PZT has a pole at 7 Hz and a zero at 70 Hz (see Fig. 4.18). Additionally, that high voltage amplifier contains a $R = 10 \text{ k}\Omega$ resistor which together with the PZT's capacitance of $C = 2.2 \mu\text{F} \pm 15\%$ [164] forms a RC circuit. The low-pass cut-off frequency for this RC circuit is $\nu_{\text{cutoff}} = 7.23 \text{ Hz}$ according to [166]

$$\nu_{\text{cutoff}} = \frac{1}{2\pi RC}. \quad (4.10)$$

The transfer function of this circuit was measured by sending a swept sine excitation to the PZT⁶; the general layout of filters when measuring this transfer function is

⁶This did not involve the OMC as a cavity, it was only the electronic circuit that was measured. The effect the PZT had on the mirror was not a factor.

shown in Fig. 4.18.




Figure 4.18: The OMC electronic circuit as organised when measuring the electronics transfer function. The high-voltage amplifier generates a pole and a zero. The OMC actuator (the PZT) forms a RC circuit with the high-voltage amplifier’s resistor adding a further pole. The filter gain shapes are shown at the relevant part of the circuit. The circuit cascades the filters with the resulting filter shape shown in the right most box. DAC and ADC are the digital to analogue and analogue to digital converters, respectively.

The result is plotted in blue in Fig. 4.19. Using trial and error, the RC’s cut-off frequency was adjusted until the model matched — a good fit for gain and phase, which requires $\nu_{\text{cutoff}} = 12.3 \text{ Hz}$, is in Figs. 4.19b and 4.19d, respectively.




Figure 4.19: PZT and CDS high voltage amplifier RC circuit transfer functions (blue) and fits (orange). The fit of (a) and (c) are the model Bode plot for the manufacturer quoted electronic component properties. (b) and (d) is a better fit, found by tuning the RC circuit cut-off frequency. The required PZT capacitance to achieve that cut-off is listed.

4.5.2 First servo design

The electronic circuit open loop transfer function of Fig. 4.19 shows the features which cannot be removed because they are inherent in the circuitry — there is a pole at 7 and 12.3 Hz and a zero at 70 Hz. In order that these are not present in the overall OMC lock open loop transfer function, a filter with zero's at 7 and 12.3 Hz and a pole at 70 Hz should be included in order to counteract the electronics.

For a loop primarily designed to suppress noise, an open loop transfer function of a single pole is usually a good foundation. A 1 Hz pole was added to create a single pole low-pass open loop. Appropriate gain in the servo can result in a unity gain frequency of 10 Hz, which for a dither lock of a low Finesse cavity, is sufficient.

The dither scheme requires a low-pass for demodulation, in order to isolate the DC part of the signal which is linear to the noise (see Sec. 2.10). Since poles also introduce a phase delay, it was deemed prudent to make sure the pole frequency was sufficiently far from the unity gain frequency to not eat into the phase margin; 1.8 kHz was chosen.

A dither frequency of 4325.16 Hz was chosen. There is also no special reason for this choice, it relies mostly on being aware of frequencies which it should not be. It cannot be below the unity gain frequency. It should be a value that far exceeds the loop unity gain frequency such that it does not cause instabilities by peaking above a gain of 1. It is best to avoid integer multiples of the mains electric AC frequency of 50 Hz for the UK because unwanted resonances could occur. It should not be close to, or an integer multiple (or division) of the actuators resonance frequency because that could cause undue instability via mechanical resonances. The resonance of this actuator with the load of the concave mirror is $\gtrsim 160$ kHz (see Sec. 4.2.7) which is far outside the bandwidth of CDS. The frequency cannot lie outside the sampling rate of CDS (16 kHz), otherwise the drive signal would not reach the PZT. Likewise, it cannot be above the actuator's bandwidth, though this information could not be found; you would know if you reached it because nothing would happen when trying to drive it at that frequency. The frequency should also not be similar or a multiple of any of the other control signals used in the rest of the experiment. These listed points are the general restrictions for the choice, apart from these the ultimate chosen frequency is somewhat arbitrary.

With these initial loop features the OMC was locked and stable enough for an open loop transfer function of the loop to be measured. This facilitates loop analysis, in order to begin the process of commissioning.

4.5.3 Final loop transfer functions

The open loop transfer function was taken when the OMC was locked. In terms of a loop diagram, the OMC loop looks like Fig. 4.20. Here, the controller $H(\omega)$ contains the loop servo and dither lock features. The measurement is slightly different to that described in Sec. 2.6.1 because the loop is closed and there is no set point.

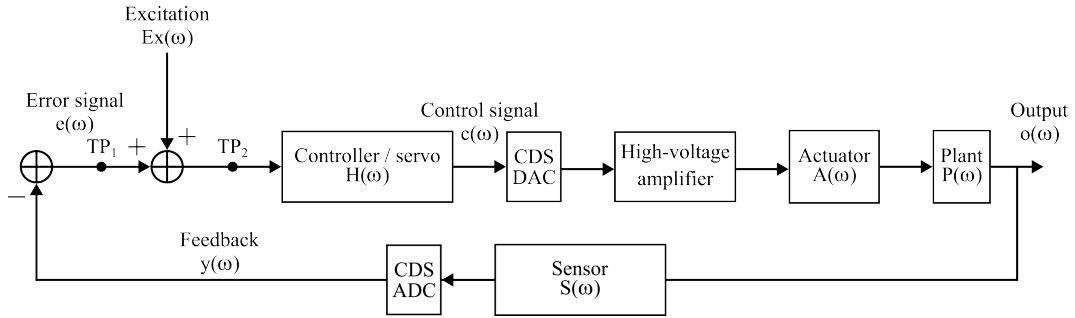


Figure 4.20: OMC dither lock loop transfer function measurement. The test points TP_1 and TP_2 are built into the CDS models — see Fig. A.7. A swept sine excitation $Ex(\omega)$ is injected between the test points. The controller $H(\omega)$ contains the loop servo and dither lock features, the actuator $A(\omega)$ is the PZT and the plant $P(\omega)$ is the OMC itself. The sensor $S(\omega)$ is the transmission photodetector. The CDS analogue to digital and digital to analogue converters (ADC and DAC, respectively) are shown for completeness.

Using the same approach and ‘loop maths’ described in Sec. 2.6.1, the test points measure

$$TP_1 = \frac{-Ex(\omega)G(\omega)}{1 + G(\omega)} \quad (4.11)$$

$$TP_2 = \frac{Ex(\omega)}{1 + G(\omega)} \quad (4.12)$$

which provides the measurements of the open loop transfer function:

$$\frac{TP_1}{TP_2} = -G(\omega) \quad (4.13)$$

and

$$\frac{TP_2}{Ex(\omega)} = \frac{1}{1 + G(\omega)}. \quad (4.14)$$

The result of Eq. 4.14 is valuable because it provides a measurement (not just a model which could be calculated from Eq. 4.13) of what the loop does to the noise of the loop per Eq. 2.89 (repeated below):

$$y(\omega) = z(\omega) + \frac{N(\omega)}{1 + G(\omega)} \quad (2.89)$$

Fig. 4.21 shows the results of the latest measurement.




Figure 4.21: Measured transfer functions of the OMC dither loop. Higher frequencies could not be measured because instability became too great and the loop lost lock. There is a function to set each test frequency's amplitude specifically, it was attempted but a suitable value was not found. The coherence of the open loop shows reliability up to the last two or three points but the trend is clear.

Fig. 4.21 shows the measured unity gain frequency of ~ 10 Hz with a phase margin of $\sim 50^\circ$.

The measured open loop transfer function is compared to the model designed in CDS in Fig. 4.22. It is designed as a low-pass filter with an additional strong cut-off for the dither demodulation. The demodulation pole, or low-pass is provided by an elliptical filter because of the associated notches. The second notch is at the dither frequency and another specific notch was added on top, to enhance the suppression of the dither in the loop.




Figure 4.22: Open loop transfer function data (blue) with the CDS model (orange). The model includes all of the relevant filters in the loop, including the RC circuit of Fig. 4.19. The gain features above 1 kHz are for demodulation purposes. An elliptical filter was used for the demodulation low-pass since it is a hard cut-off and has ‘harmonic’ notches — band-stop type features. The filter was designed such that the second harmonic matched the dither frequency. An additional notch filter was added at the dither frequency as well, to further suppress it, this is visible in the small ‘wiggle’ of the phase. Phase is greatly affected by the elliptical filter, but it far exceeds the unity gain frequency and so is not a problem for the loop.

4.5.4 Maximising the error signal

This section describes the process used to tune the demodulation phase. It is placed here because it was carried out after the lock was quite stable. Prior to this, the demodulation phase was manipulated only to the point of achieving a successful lock.

The output of the demodulation process is two potential error signals which are 90° out of phase — see Sec. 2.10.1. One of them should be maximised by tuning the demodulation phase.

In order to maximise the chosen error signal, a small (with respect to the dither signal) excitation at 124.5 Hz was injected. Like the choice of the dither frequency, the choice of this excitation was also quite arbitrary; keeping it below the dither frequency, and it not being a multiple of 50 Hz were pretty much the only criteria. With the OMC locked, each of the error signal quadratures (I and Q) were monitored and a power spectrum recorded for several different demodulation phases. Fig. 4.23 shows the power spectra of the error points after demodulation with phases -60° , -50° , -45° , -40° , -35° , -30° and -20° . The Q quadrature was chosen to be the maximised one so the I quadrature was minimised; demodulation phase $= -40^\circ$ is the best choice — it appears to be a local minimum, with the I component increasing

quite symmetrically either side.

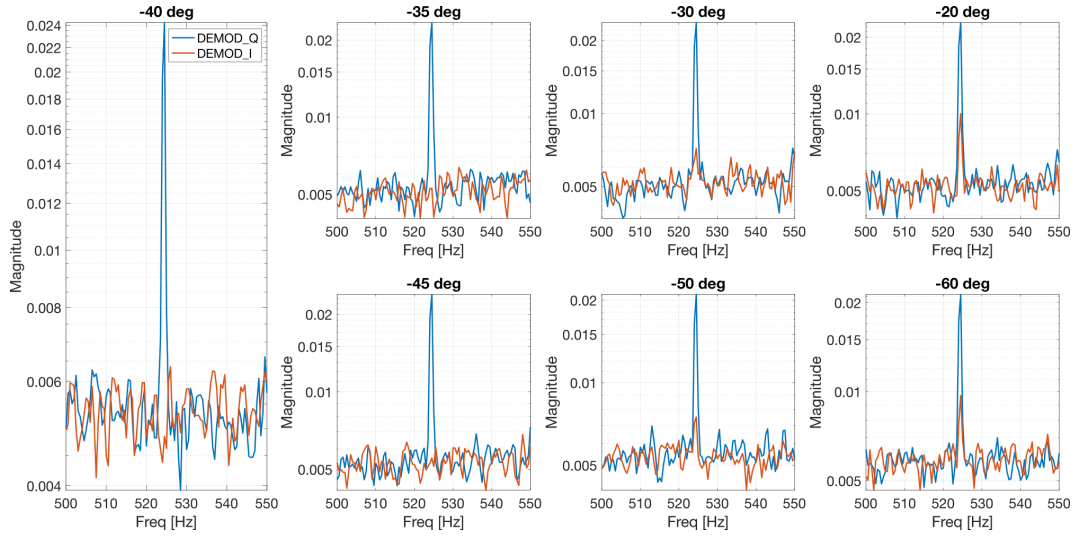


Figure 4.23: Spectra of the two demodulated outputs with different demodulation phases. An excitation was injected at the actuator. Blue is the Q path, orange the I path. Phase angles either side of the chosen one, -40° , are shown for reference. Deviation from -40° shows decreasing Q and increasing I, which is expected when the demodulation phase is not the best choice. By $+140^\circ$ the amplitude of each component would be the reverse of what is seen here.

4.6 The lock

Putting these practices into place lead to a successful *indefinitely stable* lock with an automatic re-lock capability. As a recap:

- A 4325.16 Hz dither signal is injected into the actuator.
- The transmission of the OMC is demodulated at that dither frequency, generating two components which are linear to the error.
- The demodulated signal is then calibrated such that one of the components linear to the error is isolated and maximised.
- The open loop gain is an approximately low-pass filter shape with a unity gain frequency of 10 Hz and phase margin of $\sim 50^\circ$ meaning it is unconditionally stable to noise up to 10 Hz.
- There is an automatic lock acquisition mechanism which engages a slow length scan covering $1.5 \nu_{FSR}$ when the transmission power falls below some threshold: $\sim 90\%$ peak transmission.
- Once the threshold is reached due to the slow scan climbing back up a resonance peak, the scan is disabled and the control signal is reengaged to hold resonance.

- The position of the scan when resonance was reached becomes the loop's new set-point, about which it suppresses noise.
- The continued dither and demodulation maintains transmission at the peak; meaning if the laser frequency (and thus the set-point) slowly drifts, the loop can follow it by continuously modifying the set-point — keeping the OMC locked.

With these practices in place, the OMC lock is complete. Fig. 4.24 shows 4 days of data.

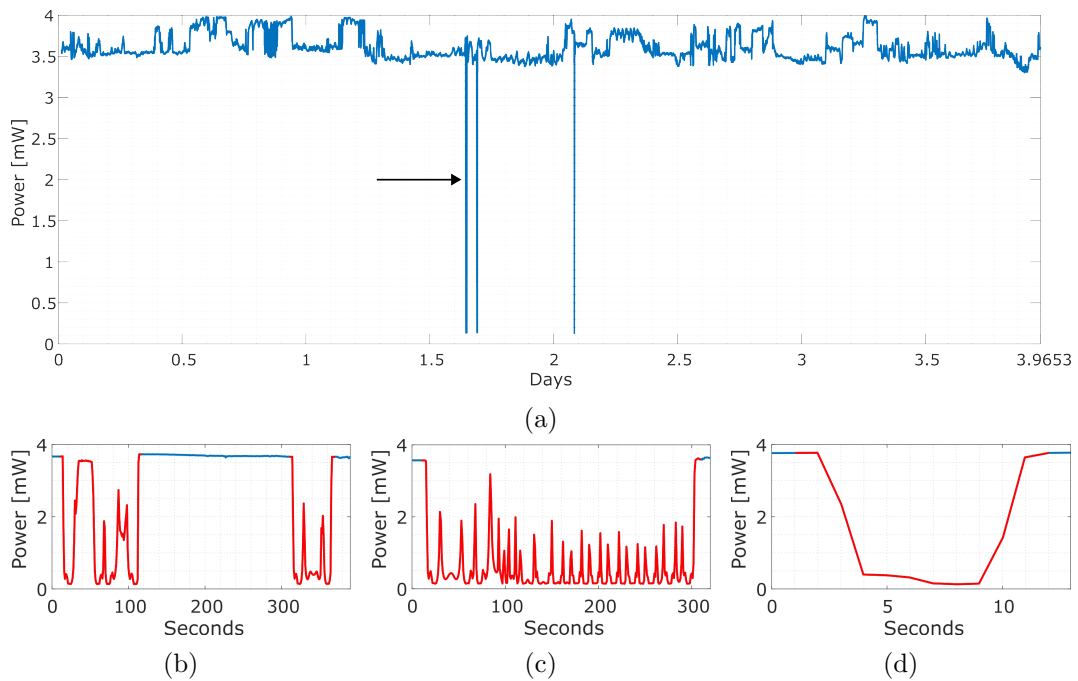


Figure 4.24: Just shy of four days of OMC lock, 342600s, totally automated. These plots show the transmitted power; (a) is the full time series, the three vertical lines where power drops to near 0 W are lock losses. The first lock loss (indicated by the arrow in (a)) actually contains two separate lock losses which occurred ~ 200 s apart, this is shown in (b). (c) and (d) are the third and fourth lock loss, respectively.

Fig. 4.24 is data with a 1 Hz acquisition rate. The reason the third lock loss in particular (Fig. 4.24c) took so long to regain lock is not fully understood. It is expected to be a laser power or polarisation drift. This would prevent the power threshold which triggers the loop to switch between the slow scan and active feedback from being reached. Until it is reached, the control signal is not engaged and the scan continues indefinitely. The four day lock shown in Fig. 4.24 had the setup shown in Fig. 4.25. It used a test laser and the polarisation control for power management was on the detection breadboard. 4 mW was chosen since it is quite a low power and it is good to test that the OMC can remain locked to low power.

Fig. 4.24 shows four lock losses totalling ~ 462.2 s. From the 342600 s of data this

equates to a duty cycle of around 99.87%.




Figure 4.25: Test setup for the four day lock of Fig. 4.24. Since the laser used was a test laser, there was no polarisation control and so that was done on the breadboard itself. A polarising beam splitter was placed following the half wave plate and the power in reflection was measured at 4 mW, then the polarising beam splitter was removed. This is because a polarising beam splitter reflects S-polarised light and transmits P-polarised. So setting the reflection ensures all the light following the half wave plate, and which reaches the OMC when the polarising beam splitter is removed, is S-polarised.

Given the polarisation control is so close to the OMC, the power instability during lock visible in Fig. 4.24 suggests either laser power or polarisation drift⁷. An analysis investigating this is shown in Sec. 4.6.1.

4.6.1 OMC lock test with a polarising beam splitter

The instability of the transmitted power of Fig. 4.24 lead to an investigation into why it was so unstable. Previous records of the long term output power stability of the lasers in use at QUEST (results not shown in this thesis) do not point to output power drift being the culprit. Some kind of polarisation instability seemed to be the other possibility. To test this, a polarising beam splitter was used.

The polarising beam splitter was placed in front of the half wave plate as shown in Fig. 4.26. The polarising beam splitter transmits P-polarised light which is not what the OMC is designed for; it can still be locked, it just has a lower Finesse (see Sec. 4.8 for more details) which was an acceptable temporary feature for this analysis. Using the reflected S-polarised light would mean replacing a steering mirror with the

⁷If the polarisers were further from the OMC, polarisation drifts could be due to some other factor, but the polarisation being set mere cm from the OMC suggests it is likely the laser itself at fault.

polarising beam splitter, or introducing several more mirrors to route that reflected light to the OMC.




Figure 4.26: Polarisation stability test setup.

The OMC was left locked overnight and data for 13.9 hours from the OMC transmission and polarising beam splitter (PBS) reflection photodetectors recovered, Figs. 4.27a and 4.27b, respectively. The correlation is clear, they are near mirror images of one another. Fig. 4.27c is the difference between the normalised power of each photodetector, $\text{norm}(\text{PBS}) - \text{norm}(\text{Trans})$, which aims to highlight how correlated they are.

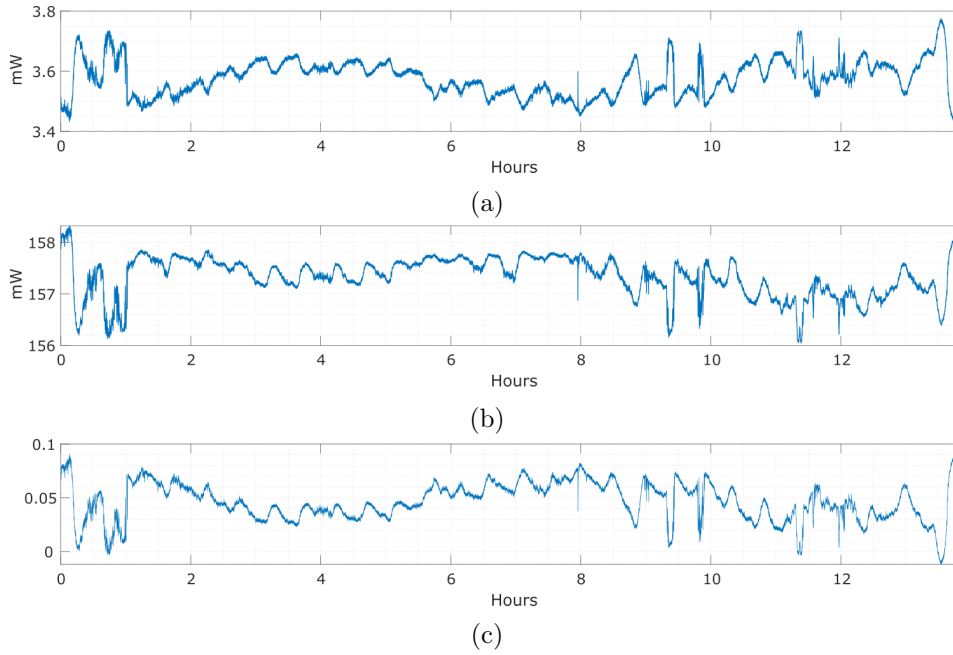


Figure 4.27: Polarisation stability test results. (a) Shows the OMC transmitted power (Trans) and (b) is the polarising beam splitter reflected power (PBS). (c) Shows the difference between the normalised powers in (a) and (b), i.e. $\text{norm}(\text{PBS}) - \text{norm}(\text{Trans})$

The FFT of both the transmission and polarising beam splitter reflection photodetectors are shown in Fig. 4.28. Unfortunately, there is no clearly dominant frequency in the spectra though one was expected — in Fig. 4.28 there does seem to be an oscillation occurring once every half an hour or so.




Figure 4.28: Polarisation stability test spectra. This is the power spectral density of the OMC transmission photodetector and the polarising beam splitter photodetector data shown Fig. 4.28. Features were expected at around $\sim 5.56 \times 10^{-4}$ Hz (half and hour), but unfortunately this appears quite featureless.

4.6.2 Further ambient condition test

During another weekend long lock some interesting features were noticed in the control signal of the loop. The full time series data for the control and actuator signals⁸ are shown in Figs. 4.29a and 4.30a, which highlights the features. Figs. 4.31a and 4.32a are the time series data for the transmission photodetector and error point signal, respectively. The FFT of these recorded time series are in Figs. 4.29b, 4.30b, 4.31b and 4.32b. The FFT data is generated because some oscillatory signals were expected and having a better understanding of their frequency could aid in noise hunting.

The data includes a few lock losses and the auto re-lock slow scan at 50 mHz is clear in the spectra. The scan's higher order harmonics are also visible in the case of the control and trans photodetector spectra.




Figure 4.29: Weekend ambient condition test - Control signal. (a) is the time series data, (b) is the FFT of (a). Small features are visible at ~ 0.49 mHz and ~ 0.66 mHz. The auto re-lock slow scan at 50 mHz and its higher order harmonics are very clear.

⁸Reminder: the difference between these two signals is lock status dependent. When locked the control signal is added to the dither signal which together become the actuator signal. When unlocked the actuator signal is only the slow length scan.




Figure 4.30: Weekend ambient condition test - Actuator signal. (a) is the time series data, (b) is the FFT of (a). Small features are visible at ~ 0.49 mHz and ~ 0.66 mHz. The auto re-lock slow scan at 50 mHz is very prominent. The peak at 40 mHz is not understood.




Figure 4.31: Weekend ambient condition test - Transmission signal. (a) is the time series data, (b) is the FFT of (a). The features at ~ 0.49 mHz and ~ 0.66 mHz are not visible as they were in Figs. 4.29b and 4.30b. The auto re-lock slow scan at 50 mHz and its higher order harmonics are very clear.




Figure 4.32: Weekend ambient condition test - Error signal. (a) is the time series data, (b) is the FFT of (a). The features at ~ 0.49 mHz and ~ 0.66 mHz are visible as in Figs. 4.29b and 4.30b. The auto re-lock slow scan at 50 mHz and its higher order harmonics are not clear here.

For the control, actuator and error spectrum data, there are some peaks at around 0.49 mHz and 0.66 mHz which are oscillations of approximately 34 and 25 minutes, respectively. To attempt finding the source of the signals, the actuator time series data was divided up into parts; using trial and error until one or both of the features vanished from the spectrum. The features were found to be present in the yellow and orange coloured time series sections in Fig. 4.33a. These sections are re-plotted in Figs. 4.33b and 4.33c, respectively. The spectrum of each of the time series data are in 4.33d, 4.33e and 4.33f; the red and blue vertical lines of these spectrum plots highlight the 0.49 MHz and 0.66 MHz signals, respectively, which are oscillations of approximately 34 and 25 minutes, respectively.

These \sim half an hour oscillations are what was expected in the polarisation test of the previous section however only the photodetector signal data were saved for that test and further analysis could not be done. Since these two tests have very similar features in the transmission photodetector, the oscillations are expected to be of the same origin.

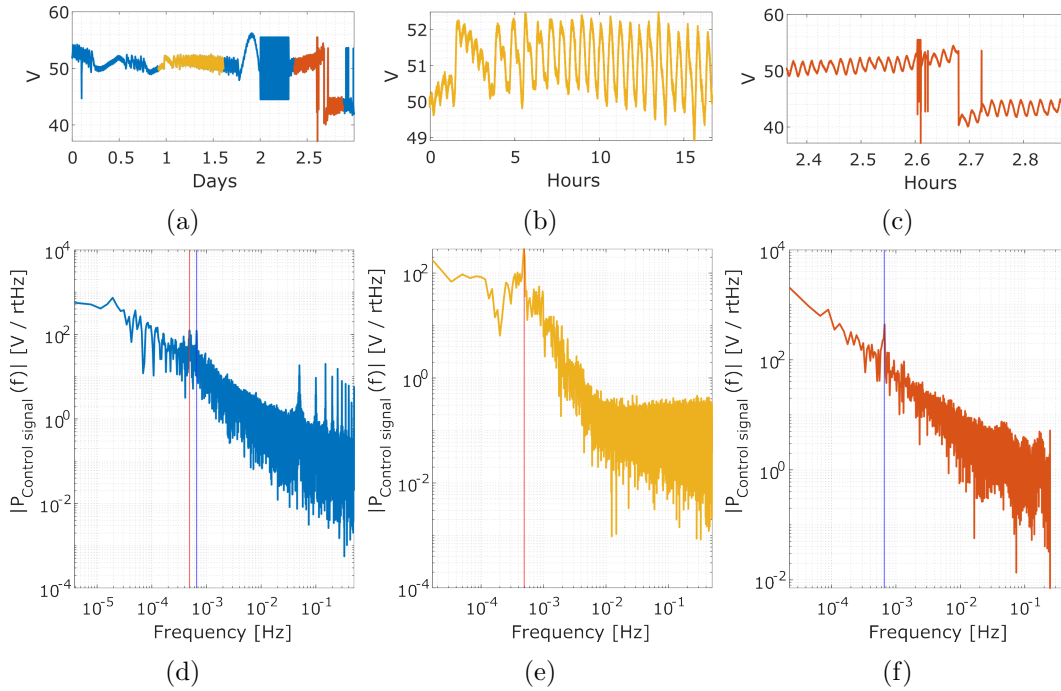


Figure 4.33: Isolating frequencies in the spectra. (a) Shows the full time series data. (b) and (c) are the coloured segments of (a). (d), (e) and (f) are the FFT of (a), (b) and (c), respectively. The red and blue lines of (d) highlight features at $\sim 0.49 \text{ mHz}$ and $\sim 0.66 \text{ mHz}$, respectively. The segments in (b) and (c) show the region of the full data where those two different spectrum features are present; highlighted by the red line of (d) repeated in (e) and blue line in (f).

Whatever caused the slightly different oscillations appears to be a day or so apart. It could be correlated to the lab’s air conditioning and/or temperature systems. With a change of temperature from day to day this could explain the slightly different frequency the lab needs to regulate heating, if that is what it is. This will be investigated along with the polarisation drift discussed in Sec. 4.6.1, it is very likely to be the same cause.

There is also a feature in the actuator spectrum at almost exactly 40 mHz , a period of 25 s . It is not obvious what this is. It is present in all three segments of Fig. 4.33 even where no auto-scan is occurring. It should not be linked, but it is oddly close to the scan frequency. This feature was discovered during the writing of this thesis and no commissioning or noise hunting was attempted due to a lack of time.

4.7 OMC characterisation

4.7.1 Mirror reflectivity

The reflectivities of the OMC mirrors were measured using a test Coherent Mephisto laser [96]. The mirror was positioned approximately at normal incidence and the beam was confirmed to be central. An iris was closed tightly before the mirror and an IR camera was used to observe when the reflection was directed back through

the it. This was considered to be the normal incidence reference angle. Afterwards, the mirror was rotated to the required angle of incidence — 37.5° for the input and output mirrors and 15° for the concave mirror. The power in reflection and input were recorded at 3.33 Hz for 30 s, giving 100 data points each. The reflection was recorded first since the sensor being at normal incidence to the beam is important for operation. After the reflection measurement the input power measurement was done; this was much more straight forward to ensure normal incidence and could be done quickly in order to minimise effects due to laser power drift. Results are in Table 4.3.

	Measured	Spec
Input	0.91786 ± 0.00034	0.925
Output	0.91544 ± 0.00023	
Concave	0.99557 ± 0.00014	0.9999

Table 4.3: Table of OMC mirror power reflectivities.

The estimated Finesse (per Eq. 4.5) from these reflectivities is $\mathcal{F} = 35.187 \pm 1.184$. This value will be compared to the measured Finesse which is in Table 4.4.

4.7.2 Length scan analysis

The Finesse and bandwidth of the cavity were calculated using a triangular signal sent to the actuator to carry out a length sweep, a section of the data is in Fig. 4.34.




Figure 4.34: OMC length scan. The actuator was driven by a 10 mHz, 36.6 Vpp with 57.3 V offset triangular signal which covers $3\nu_{FSR}$. Data acquisition rate was 4 kHz.

Given the hysteresis of the actuator (Fig. 4.14), attempting to convert from volts to meters did not provide a linear x -axis. To calibrate the x -axis more appropriately,

the transmission peaks of the scan were assumed to be $\nu_{FSR} = \lambda/2 = 532 \text{ nm}$ separate. Each peak's voltage was identified (orange dots of Fig. 4.35) and these values plotted against the distance the PZT *should* have been displaced. The orange dashed fit line of Fig. 4.35 was generated by Matlab's curve fitting tool and has the form $m = -2.307 \times 10^{-10} \text{ V}^2 + 7.1697 \times 10^{-8} \text{ V} + -2.6617 \times 10^{-6}$. This enabled an appropriate calibration of the x-axis for the length scans, blue in Fig. 4.35. This calibration was carried out for each half cycle of the length scan. The result of the calibration for one scan is in Fig. 4.36a.

Each transmission peak was isolated to measure the bandwidth, assuming a Lorentzian fit [167] of the form

$$\frac{A}{\pi} \frac{\nu_{cav}/2}{(x - x_0)^2 + (\nu_{cav}/2)^2} \quad (4.15)$$

where A is the amplitude scaling, ν_{cav} the bandwidth and x_0 the peak position. This is shown for one transmission peak in Fig. 4.36b.




Figure 4.35: $3\nu_{FSR}$ section of the length scan. Blue: Transmitted power. Orange: Dots show the peak voltages and the dashed fit assumes these peaks are $1\nu_{FSR}$ separate.




Figure 4.36: (a): $3\nu_{FSR}$ section of the length scan with x-axis calibration. (b): Blue is one isolated transmission peak, orange dashed is a Lorentzian fit.

The 30 total analysed resonance peaks yield an average bandwidth of $\nu_{cav}^- = 0.027\nu_{FSR}$ with a standard deviation of $\sigma = 0.001\nu_{FSR}$.

ν_{FSR} was also estimated using the round-trip-length with an assumed tolerance of 2.0mm, based on manufacturing tolerances and glue layer depth. Cavity Finesse was then calculated, results are in Table 4.4.

	Measured	Specification
Bandwidth	(218.507 ± 8.355) MHz	200 MHz
Free Spectral Range	(8.000 ± 0.427) GHz	8 GHz
Finesse	36.612 ± 2.403	40

Table 4.4: OMC optical properties.

The Finesse calculated here is slightly higher than that estimated from mirror reflectivity measurements (Table 4.3, $\mathcal{F} = 35.187 \pm 1.184$). The discrepancy is likely to be primarily due to the mirror reflectivity measurement. If we assume all the reflectivity loss is in the curved mirror r_3 , comparing the Finesse values can give an estimate of the power drift it would have required to produce the same result.

$$\mathcal{F}_{scan} = \frac{\pi\sqrt{r_1 r_2 r_3}}{1 - r_1 r_2 r_3} = 36.612 \quad ; \quad \mathcal{F}_{coat} = \frac{\pi\sqrt{r_1 r_2 r_3}}{1 - r_1 r_2 r_3} = 35.187 \quad (4.16)$$

$$(r_1 r_2 r_3)_{scan} = 0.9178 \quad ; \quad (r_1 r_2 r_3)_{coat} = 0.9146 \quad (4.17)$$

$$\left(\sqrt{0.925 \cdot 0.925 \cdot R_3}\right)_{scan} = 0.9178 \quad ; \quad \left(\sqrt{0.925 \cdot 0.925 \cdot R_3}\right)_{coat} = 0.9146 \quad (4.18)$$

$$R_{3,scan} = 0.9845 \quad ; \quad R_{3,coat} = 0.9776 \quad (4.19)$$

Assuming both measured a reflected power P_r of 100 mW, the input power P_i can be calculated

$$P_{i,scan} = \frac{P_{r,scan}}{R_{3,scan}} = \frac{100 \text{ mW}}{0.9845} \quad ; \quad P_{i,coat} = \frac{P_{r,coat}}{R_{3,coat}} = \frac{100 \text{ mW}}{0.9776} \quad (4.20)$$

$$P_{i,scan} = 101.57 \text{ mW} \quad ; \quad P_{i,coat} = 102.29 \text{ mW} \quad (4.21)$$

This model shows that to go from the Finesse measured during the coating analysis to that measuring with the scan analysis, only a 0.7% drift in laser power would have had to occur⁹ during the measurements of the coatings. This is effectively certain to have happened without any power stabilisation.

However, the measurement calculated from the scan also relies on an approximation of the PZT displacement. The fit line of Fig. 4.35 is generated from only three possible points. But, given the linearity of the x -axis this calibration provides and the amount of the transmission peaks analysed, this is the chosen Finesse value to use.

As part of this length scan, the position of the higher order modes were identified using the modified version of equation 14 in [90] — as described in Sec. 4.2.2

$$\nu_{nm} = \frac{\nu_{FSR}}{\pi} [n \arccos(\sqrt{g_t}) + m \arccos(\sqrt{g_s})] + \frac{\nu_{FSR}(1 - (-1)^n)}{4} \quad (4.22)$$

which is the frequency difference between the higher order modes of order n , m and the fundamental mode. Each resonance peak has ‘positive’ and ‘negative’ higher order modes associated with it, depending on which side of the peak the scan is. For each of the three resonance peaks of Fig. 4.37, $\pm\nu_{nm}$ for each mode of order $n + m = 1$ up to $n + m = 3$ was identified. The green values of Fig. 4.37 are for the left most resonance at $\nu_{FSR} = 0$, the black are for the central resonance at $\nu_{FSR} = 1$ and the red are for the right most resonance at $\nu_{FSR} = 2$.

⁹This result can be achieved at Eq. 4.19 by $100(R_{3,scan} - R_{3,coat})$.

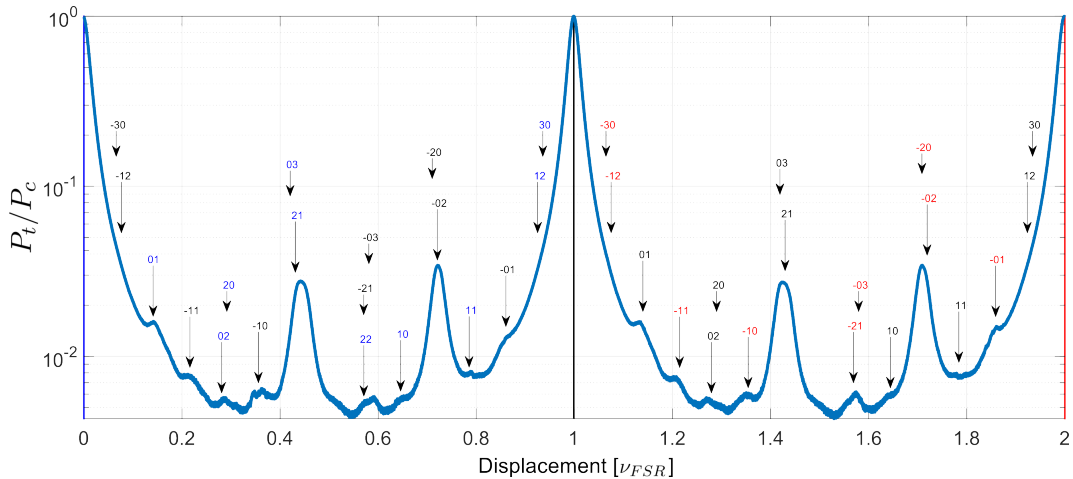


Figure 4.37: OMC scan with higher order modes identified. The numbers represent the order, mn . Black are linked to the central resonance. Green are those of the left most resonance and red the right resonance. The arrows are plotted at the point representing the solution to Eq. 4.22, their not aligning to the peaks exactly is likely due to the x axis linearisation method described above.

This length scan, particularly in the log scale of Fig. 4.37, shows the higher order modes which are excited by the cavity itself. These modes are not necessarily high amplitude in the input beam, but rather due to misalignment and mode mismatch coupling light from the fundamental mode to these other modes, as described in Sec. 2.3.7. These imperfections cause the particular modes to resonate at their specific eigenmode/cavity length. The OMC could in principle be locked to these modes.

This length scan showcases how the OMC will reject modes. If the same order modes were generated within the power-recycled Michelson interferometer cavity¹⁰, they'd be present in the anti-symmetric port beam and could have similar relative amplitudes as in Fig. 4.37. When that beam is to pass through the OMC though, those modes fall as far outside the OMC's fundamental mode bandwidth as do the modes of Fig. 4.37, and so will not be transmitted through the OMC. The low points of the scan are below two orders of magnitude, it is reasonable to expect any higher order modes present in the anti-symmetric beam will be rejected by at least amount.

4.7.3 OMC losses

The loss of the OMC can be calculated by assuming one of the mirrors is responsible for all of the cavity loss. For an optic, $R + T + P = 1$ where R and T are the power reflectivity and transmission coefficients already introduced [67][77]. The new factor is the loss, included here as the coefficient P . We can update the original Finesse equation for a three mirror cavity, Eq. 4.5 ($\mathcal{F} = (\pi\sqrt{r_1 r_2 r_3}) / (1 - r_1 r_2 r_3)$), to include

¹⁰There will be differences in their position since Eq. 4.22 is specifically for a triangular cavity with one concave mirror at non-normal incidence. But any differences in mode spacing will actually be beneficial since the same order modes will now be degenerate (see Sec. 4.2.2).

loss for one of the mirrors, i.e. $r_1 = \sqrt{1 - T_1 - P_1}$. However, since this mirror is assumed to be the only one with loss, this P_1 is recast as P_{OMC} since it is now attributed with the overall OMC loss:

$$\mathcal{F} = \frac{\pi \sqrt{\sqrt{1 - T_1 - P_{OMC}} \sqrt{R_2 R_3}}}{1 - \sqrt{1 - T_1 - P_{OMC}} \sqrt{R_2 R_3}}. \quad (4.23)$$

Solving this for P_{OMC} gives

$$P_{OMC} = \frac{1}{2} \left(2 - 2T_1 - \frac{2\mathcal{F}^4 + 4\pi^2 \mathcal{F}^2 + \pi^4}{\mathcal{F}^4 R_2 R_3} + \frac{\pi \sqrt{(2\mathcal{F}^2 + \pi^2)^2 (4\mathcal{F}^2 + \pi^2) R_2^2 R_3^2}}{\mathcal{F}^4 R_2^2 R_3^2} \right). \quad (4.24)$$

The specification values are $R_2 = 0.925$, $R_3 = 0.9999$ and $T_1 = 1 - 0.925$. Measured $\mathcal{F} = 36.612 \pm 2.403$. Plugging these values in gives us an OMC loss in terms of the specification Finesse of $P_{OMC} = 0.01 \pm 0.41$.

The identification of the higher order modes shown in Fig. 4.37 also facilitates estimating a reasonably accurate throughput of the OMC. The power of each higher order mode related to the central peak was summed (totalling ≈ 0.38 mW) and subtracted from the input power (≈ 3.33 mW). This $\approx 3.33 - 0.38 = 2.95$ mW is the power which ought to remain in the fundamental mode. Normalising the transmitted power to 2.95 mW gives the data of Fig. 4.38.




Figure 4.38: Length scan where transmitted power is normalised to (input power - total power in the higher order modes). Peak average is $\approx 98.8\%$.

Of the three peaks in Fig. 4.38, this approach gives an average transmission of the fundamental mode of $\approx 98.8\%$ suggesting an OMC loss of $\approx 1.2\%$, or $P_{OMC} \approx 0.012$; which is consistent with the loss estimated from Eq. 4.24.

4.8 Commissioning lessons

There are two major lessons learned during the commissioning phase. The plots shown here are screenshots from CDS, which originally have a black background, so please excuse the poor quality — the data could not be recovered in order to generate the plots properly.

1. Input polarisation resonance dependency.

The S- and P-polarisations have separate resonances and display different bandwidths. Fig. 4.39 shows the transmission when scanning the OMC length with 45° polarisation injected. The OMC mirrors were designed for S polarisation. The different bandwidth has to be a result of the mirror reflectivities; ν_{FSR} is identical since L_{rt} is unchanging leaving only $r_{1,2,3}$ able to influence \mathcal{F} . The separate resonances is due to unequal effective depths of the coatings for S- and P-polarisation. This results in changes to the reflected field phase for each polarisation and so in different resonances. Interestingly, the higher order modes for the wider resonances (P-polarised) are significantly larger than for the S-polarised resonances. Deeper analysis was not carried out to identify them but they are likely mode-matching related.

Viewed on a CCD the transmissions look identical, and their different bandwidth is not immediately apparent. The lesson is to be careful to correctly polarise the input light.




Figure 4.39: OMC length scan transmission spectrum with 45° linearly polarised input light. Each peak is a fundamental mode resonance, but close examination reveals two distinct bandwidths. These are the resonances of the S- and P-polarisations of the field.

2. Laser crystal temperature mode hops.

At various crystal temperatures, the laser exhibits a mode hop and changes frequency [96]. These mode hops are not discrete as appears in Fig. 4.40. If operating near one of the mode hop temperatures, there can be dual fundamental mode outputs of different frequencies. Fig. 4.41 was recorded when operating at 26°C. The polarisation is still at 45° here, so there are also separate transmissions for each S and P pol. The highest peaks of roughly equal amplitude are one fundamental mode. The lower peaks are of a different fundamental mode. This was conformed using a CCD in transmission. Tuning the temperature to 28°C removed the second fundamental mode, resulting in the transmission of Fig. 4.39.




Figure 4.40: Coherent Mephisto laser frequency hop regions as a function of crystal temperature [96].




Figure 4.41: OMC length scan showing four fundamental transmissions. Higher peaks are one frequency split by their polarisation. Lower peaks are a different frequency (due to the temperature operating at a mode-hop temperature), and there are also two due to the polarisation split.

4.9 Summary

This chapter has been dedicated to the OMCs of QUEST. The motivation for why an OMC is required as part of the strategy to meet QUEST's design sensitivity was described. It is a vital component which facilitates the high circulating power of the power-recycled Michelson interferometers, which improves the shot-noise-limited displacement sensitivity.

The specifications were detailed and each important factor in the design were described, including how they impact the overall properties and why certain values were chosen.

A characterisation of the OMC using a test laser and the digital CDS showcased

- The working dither feedback loop with a unity gain frequency of 10 Hz, providing indefinitely stable locking and automatic re-lock capability, with a duty cycle upwards of 99.8 %.
- A length scan providing analysis of the Finesse and bandwidth.

The OMC was specified to suppress higher order modes by an order of magnitude (requiring a Finesse > 11), and have a bandwidth of ≥ 200 MHz. The OMC built and characterised was measured to have a Finesse of 36.612 ± 2.403 and a bandwidth of (218.507 ± 8.355) MHz. Losses in the OMC have been estimated using a length scan at ≈ 1.2 %.

Chapter 5

Fibre noise cancellation method for constructing an X-arm vs science laser error signal at LIGO

The following is a brief synopsis of the lock scheme to provide some context to the work of this chapter. The goal here is to develop an overarching picture of how Advanced LIGO acquire a comparison between the science laser frequency and the common-mode arm length; the core of this project.

5.1 Advanced LIGO

Advanced LIGO has five longitudinal degrees of freedom:

1. Michelson interferometer length, $\text{MICH} = l_X - l_Y$
2. Power-recycling cavity length, $\text{PRCL} = l_p + (l_X + l_Y)/2$
3. Signal-recycling cavity length, $\text{SRCL} = l_s + (l_X + l_Y)/2$
4. Common-mode arm length, $\text{CARM} = L_X + L_Y$
5. Differential-mode arm length, $\text{DARM} = L_X - L_Y$

As in Fig. 5.1: $l_{X,Y}$ is the distance from the beam splitter to the respective input test mass, l_p is the distance from the power recycling mirror to the beam splitter, l_s is the distance from the signal-recycling mirror to the beam splitter and $L_{X,Y}$ is the length of the respective arm cavity — the distance between the input test mass and the end test mass. The first three listed degrees of freedom, MICH, PRCL and SRCL form a dual-recycled Michelson interferometer and are referred to here as the

corner station degrees of freedom or corner station cavities. The last two, CARM and DARM are related to the arm cavities.

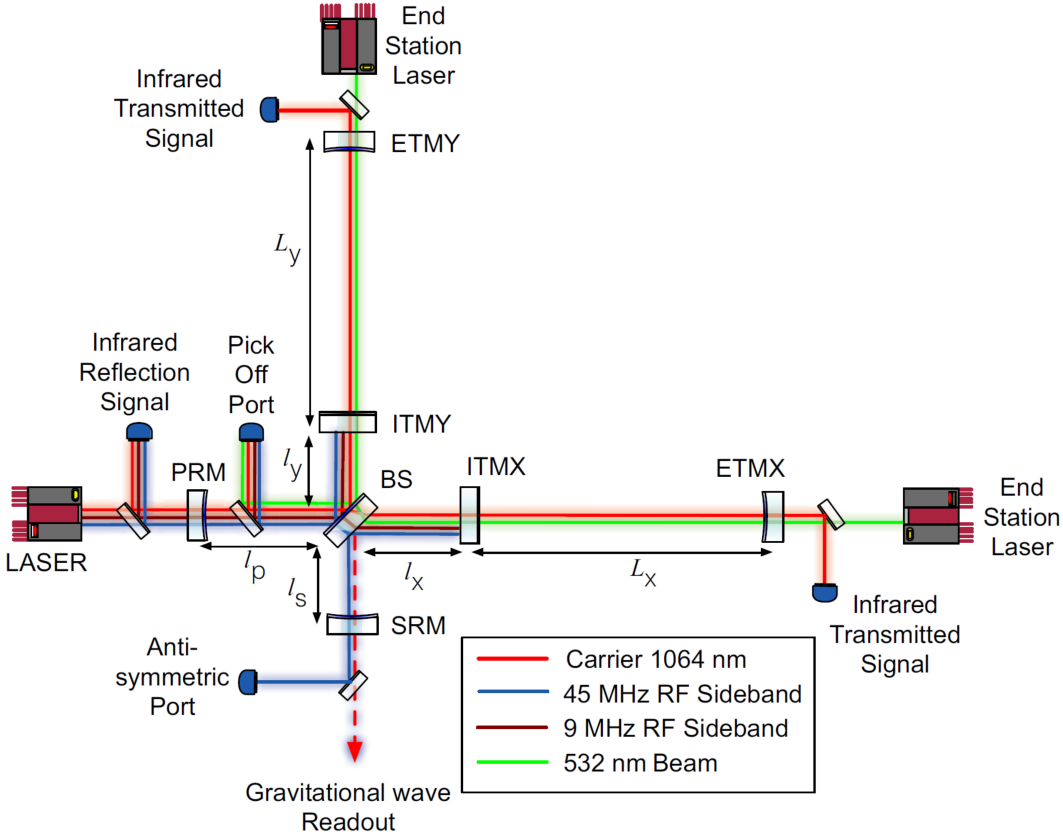


Figure 5.1: The 5 longitudinal degrees of freedom for Advanced LIGO [168]. $l_{x,y}$ is the distance from the beam splitter to the respective input test mass (ITM), l_p is the distance from the power recycling mirror (PRM) to the beam splitter (BS), l_s is the distance from the signal-recycling mirror (SRM) to the BS and $L_{x,y}$ is the length of the respective arm cavity: ITM to end test mass (ETM). Where each laser are resonant is shown here. The 1064 nm carrier is the fundamental mode of the science laser (labelled ‘LASER’). The 45 MHz and 9 MHz sidebands are used to lock the signal- and power-recycling cavities, respectively. The 532 nm green is output from the end station lasers for arm cavity control.

The five degrees of freedom are highly coupled and each can strongly impact the resonance state of the others [168]. Since the signal-recycling mirror was added to the configuration, reliable control of all of Advanced LIGO’s degrees of freedom by a single laser is not feasible, effort must be made to somehow decouple the degrees of freedom. Advanced LIGO uses the science laser and various sidebands to control the corner cavities and, as a way of decoupling the arm cavities from the corner cavities, two new lasers have been added; one at each end station. These end station lasers (shown in Fig. 5.1) are used to lock the arm cavities with a frequency offset to the science laser, preventing the science laser becoming resonant with the arms. This is required because of the phase flip of the reflected/transmitted fields when resonance is achieved (see Fig. 2.26). In this case, that would also flip the signs of the control signals for the corner station cavities causing instabilities in those

degrees of freedom. For this reason the arm cavities must be held off resonance to the science laser field in the initial stages of lock acquisition. At later stages of the lock sequence this frequency offset is steadily removed in order that the science laser does become resonant in the arms.

To achieve this controlled relationship, the frequency of the science laser and the end station lasers must be very well coupled. If they were independent then it would not be possible (using this approach of separate lasers) to reliably hold the arm cavities off resonance to the science laser. The lasers are frequency coupled initially using a fibre optic link and ultimately using heterodyne mixing. For the purposes of this project, the objective is in the construction of a science laser vs common-mode arm length error signal. The fundamental requirement for this is for the eventual science laser and arm cavity/common-mode arm length resonance.

For the below only one end station will be discussed, but the same happens at both.

A 1064 nm pick-off of the main science laser is sent to the end station via the fibre optic link. The output is mixed with the end station laser and a phase-locked-loop [169] maintains a coupling between the science and end station lasers. A low noise voltage controlled oscillator drives an acoustic-optical modulator which introduces a frequency set-point into the to the end station phase-locked-loop — the end station laser's frequency is not driven to equal the science laser's frequency; instead, their beat (difference) frequency is made to follow the acoustic-optical modulator output frequency. The end station laser is now following the science laser with that offset.

The end station laser has two outputs, the 1064 nm infrared output and a frequency doubled clone, the 532 nm green output. The 1064 nm output is the one used in the end station phase-locked-loop but the 532 nm is used for locking the arm cavity. The arm mirrors (the input and end test masses) have dichroic coatings meaning a different reflectivity for each wavelength. This unequal reflectivity results in a different cavity Finesse for the two wavelengths; the 1064 nm Finesse is around 450 [170] while the 532 nm Finesse is more like 13 [168]. The substantially lower Finesse makes the arm cavity far easier to lock with green light, given the wider bandwidth (see Sec. 2.3.3).

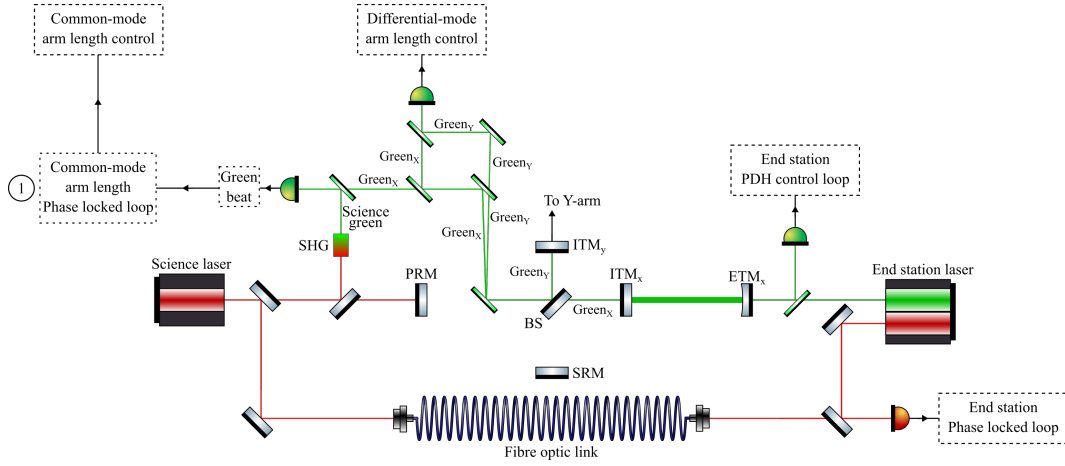


Figure 5.2: Layout of the arm length stabilisation system, inspiration taken from [168]. This is a simple block diagram of the key parts of the scheme. A pick-off of the science laser is directed to the end station via a fibre optic link. At the end station, the end station laser is phase-locked to the science laser. The end station laser outputs green light which is made resonant with the arm cavity using a Pound–Drever–Hall lock. The green light from both arm cavities (Green_X and Green_Y for the X- and Y-arm, respectively) is extracted from the centre cavities and used for both the differential- and common-mode arm length control loops. For differential-mode control, the green light from the X- and Y-arms are superimposed and beat together to give an error between the arms. For the common-mode control, the X-arm is superimposed with a pick-off of the science laser which has been converted to green by a second harmonic generator (SHG). This gives an error between the science laser and the X-arm. That error is fed back to the science laser frequency stabilisation loops. For information about the loop specific actuators see [103, 168, 171], they are beyond the scope of this project and are not shown here. The circled number 1 highlights the signal point for the green line of Fig. 5.13a.

Locking the arm to the end station green laser stabilises the positions of the arm mirrors using green light where the cavity Finesse is lower than at the 1064 nm wavelength of the main laser. This green light lock slows the mirrors down so that it is easier to lock to the science laser at 1064 nm, where the arm cavity has higher Finesse, later in the locking sequence. This locking strategy requires phase locking of the main science laser to the green laser at the end station. This is possible because the end station laser has both 1064 nm and frequency doubled green light outputs, so that the 1064 nm light can be beaten against the main laser light to establish phase synchronisation. In the current Advanced LIGO configuration, the beating is achieved by transmitting science laser light to the end station laser using a fibre optic link. However length fluctuations of the fibre optic link introduces frequency noise $\nu_{\text{Fib}}(\omega)$ onto the science laser [103, 172], which corrupts the end station phase-locked-loop. The frequency noise introduced is a function of the fibre’s length L_{Fib} [103]:

$$\nu_{\text{Fib}}(\omega) = i\omega \frac{2\pi}{\lambda} L_{\text{Fib}}(\omega) \quad (5.1)$$

where ω is the angular frequency and λ is the wavelength of the light.

Given this additional noise in the end station phase-locked-loop, there must be a direct comparison between the science laser and the end station laser in order to

effectively couple and control their frequency offset. The two light sources which are used for this are:

1. The X-end¹ station green light that is transmitting through the X-arm input test mass into the corner station — ‘Green_X’ in Fig. 5.2.
2. A science laser pick-off which is frequency doubled to 532 nm green light via a second harmonic generator — ‘Science green’ in Fig. 5.2.

These two green light sources are superimposed and the resulting beat signal is sent to the common-mode arm length phase-locked-loop. This ‘green beat’ is the ultimate comparison between the lasers, or technically, the science laser and the X-arm length.

From here, control signals to maintain the frequency relationship between the science laser and the common-mode arm length are sent to the science laser’s initial frequency stabilisation control schemes. These schemes involve a temperature controlled fused silica (fixed length) reference cavity [173, 174], and a triangular input mode cleaner [175]. These are outside the scope of this project and will not be described further.

5.1.1 Challenges

A key reason that the Advanced LIGO lock scheme requires the green beat and common-mode arm length phase-locked-loop is because of the frequency noise introduced by the fibre optic link. If the fibre noise was not a factor, the end station laser (and therefore the X-arm/common-mode arm length) could be held at a stable enough offset from the science laser that nothing further would be necessary.

Acquiring the green beat involves a convoluted beam path, which uses many auxiliary optics, introducing multiple unnecessary sources of loss. A diagram showing the green beam path is in [176].

The core of the project described in this chapter is the trialling of a new method of acquiring the error between the science laser and the common-mode arm length, without using the green beat. It is based on the principle of double sampling the noise, providing two signals that can be combined in such a way as to ‘subtract’ the noise. This method is described next.

¹The common-mode arm length control loop uses the X-arm light and superimposes it with the science laser. The green light from both arms are also superimposed in the DARM control loop (see Fig. 5.2). When both loops are active, the X- and Y-arm cavities are coupled and so for common-mode control, sampling only one arm is equivalent to sampling both.

5.2 Fibre noise cancellation method

In order to make the green beat section redundant, the fibre noise must be otherwise suppressed. A solution is to sample the noise twice by sending light in both directions and combining the signals from either end in such a way as to cancel the common noise². A version of this method was proven possible using 4.6 km of test optical fibre [103, 172]. A layout of the test is in Fig. 5.3.

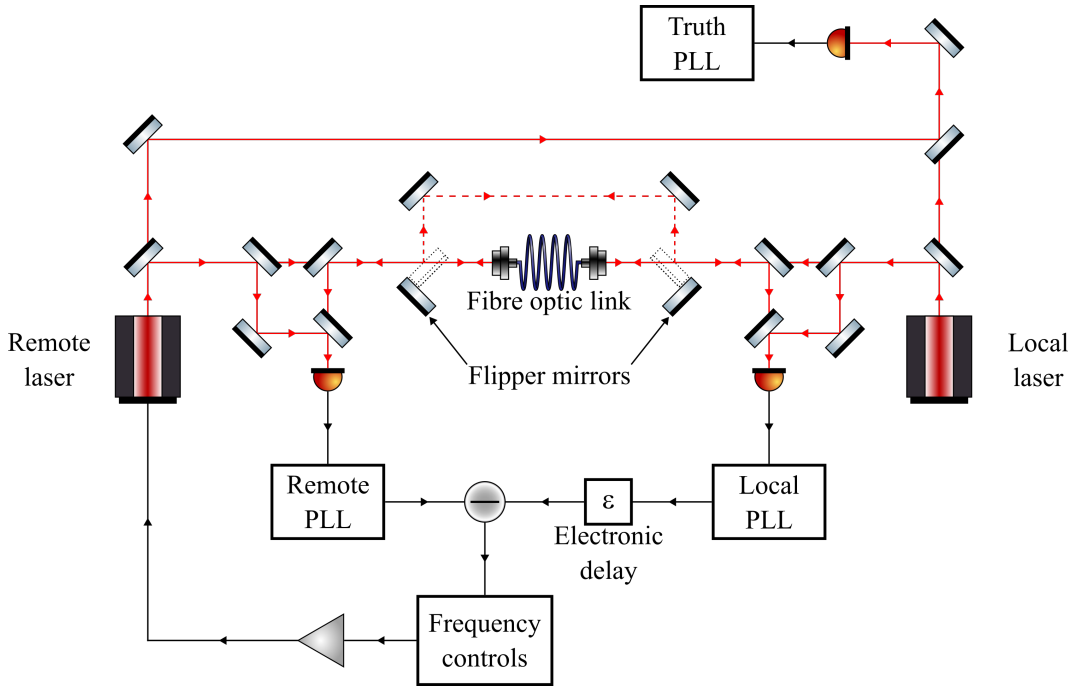


Figure 5.3: Proof of concept experimental layout [103, 172]. The local and remote lasers represent Advanced LIGO’s science and end station lasers, respectively. The fibre optic link was 4.6 km in length but unlike at Advanced LIGO, the couplers either end were not 4 km apart. The truth phase-locked-loop (PLL) was the passive witness path, it provided a frequency comparison between the lasers and was not used for any feedback — this provided the beat signal that was recorded for results. The baseline experiment was to use the flipper mirrors to bypass the fibre link and use the remote laser phase-locked-loop alone to match their frequency. The Advanced LIGO representative experiment moved the flipper mirrors which coupled both lasers to the fibre link. Only the remote phase-locked-loop was used to frequency match the lasers. The trial method used the same configuration as the Advanced LIGO representative experiment but this time the local phase-locked-loop was also included. The electronic delay simulated the delay expected at Advanced LIGO given the laser separation. The signals from the remote and local phase-locked-loops both contain a sample of the fibre noise and in taking their difference, the common fibre noise is attenuated [103]. Results are in Fig. 5.4.

During the experiment in [103, 172] shown in Fig. 5.3, two lasers ‘Science’ and ‘Auxiliary’ were frequency matched using separate phase-locked-loops. Three methods for minimising their frequency difference were compared. The comparison was measured using a ‘Truth’ path with its own phase-locked-loop. This was not used for any feedback, it served as a witness to the true frequency difference between the

²A similar scheme is used at KAGRA [177, 178]

lasers for all three cases. Data from the truth phase-locked-loop was sampled at 100 kHz and an amplitude spectral density of 10s of data is shown in Fig. 5.4.

The first method is the benchmark, it involved a direct optical mixing of their fields; this is the case of Fig. 5.3 where the ‘Flipper mirrors’ are put into the beam path, bypassing the fibre optic link. The auxiliary laser was then phase locked to the science laser using the auxiliary phase-locked-loop. This is trend (c) in Fig. 5.4. The second method removed the Flipper mirrors, coupling both lasers into the 4.6 km fibre optic link. Now, each laser propagates through the fibre before reaching the phase-locked-loop at the other end. The Auxiliary laser was again phase locked to the science laser using (only) the auxiliary phase-locked-loop, but this time the fibre optic frequency noise is present in the loop. This is trend (a) in Fig. 5.4.

The third method is the trial solution of suppressing the fibre frequency noise. This time, the science laser phase-locked-loop was also engaged. The control signals from both phase-locked-loops now contain a sample of the fibre frequency noise. Their difference removes the common fibre noise [103] and is fed back to the auxiliary laser’s frequency control. This is trend (b) in Fig. 5.4.




Figure 5.4: Proof of concept results [103, 172] recorded from the truth phase-locked-loop (see Fig. 5.3). Line (c) is the baseline experiment where the fibre link is bypassed and the remote phase-locked-loop is used for frequency matching the remote laser to the local laser. Line (a) is the Advanced LIGO representative result where the fibre link is included and only the remote phase-locked-loop is used. Line (b) is the fibre noise cancellation method, the fibre link is included and both the remote and local phase-locked-loops are used and subtracted. Line (d) is the Advanced LIGO frequency noise requirement.

With the fibre noise subtraction, their frequency difference noise fell by orders of magnitude [103, 172].

5.2.1 Fibre noise cancellation at LIGO

The work involved in this project has a similar core working principle to the above experiment (and that used at KAGRA [177, 178]) — sampling the fibre noise twice by having light travel down both directions and combining the signals from either end to filter out the common noise. There are two control loops at Advanced LIGO which already contain a sample of the fibre noise: the end station phase-locked-loop and the Pound–Drever–Hall loop used to lock the end station laser’s green output to the X-arm. The Pound–Drever–Hall control loop was used.

To generate a signal at the corner station end of the fibre link, the fibre breadboards at both ends needed modification. In the original Advanced LIGO scheme (described in Sec. 5.1) the science laser was coupled to fibre and sent to the end station while the end station laser was not coupled to the fibre at all — see Fig. 5.2. Also, there was no need for a photodetector on the corner station fibre breadboard since no lasers were being superimposed there. With this trial scheme, the end station laser was coupled to the fibre link and sent to the corner station — see Fig. 5.5. A photodetector was added to the corner station fibre breadboard, because now the 1064 nm lasers are superimposed in the corner station as well as the end station — this is shown in Fig. 5.5 as the ‘IR beat’. This additional corner station photodetector measures the beat between the lasers, corrupted by the fibre noise; providing a new signal with a sample of the fibre noise.

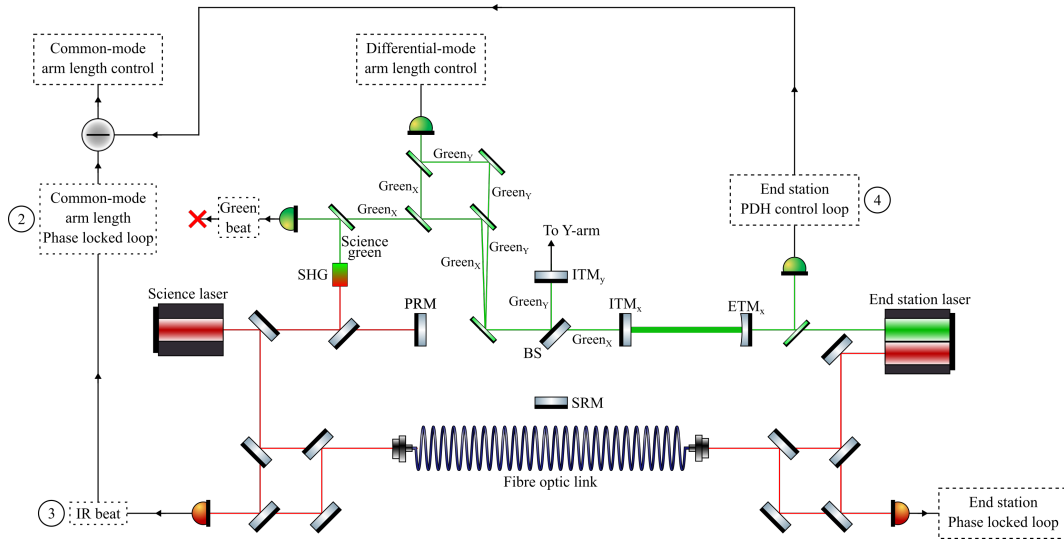


Figure 5.5: Modification of Fig. 5.2 to include the new laser paths facilitated by the updated fibre optic breadboards. 1064 nm light from the end station is now coupled to the fibre and a photodetector is added to the corner station breadboard. The previously used green beat is no longer the input to the common-mode arm length phase-locked-loop. Now, the input is the IR beat between the science laser and the end station laser which is output from the fibre into the corner station. For information about the loop specific actuators see [103, 168, 171], they are beyond the scope of this project and are not shown here. The circled numbers 2, 3 and 4 highlight the signal points for the red line of Fig. 5.13a, and the magenta and blue lines of Fig. 5.13b, respectively.

The final step was to switch from using the green beat to using the IR beat as an input to the common-mode phase-locked-loop, and then sum the result to the end station Pound–Drever–Hall loop control signal — shown in Fig. 5.5 by the red cross cutting connection for the green beat and the new difference junction before the common-mode arm length control box. This provides the common-mode arm length control loop with a ‘clean’ comparison between the science laser and the X-arm without the need for the green beat.

5.2.2 Summary

Fig. 5.6 is a comparison list between the two methods. Fig. 5.6a is the original Advanced LIGO strategy, Fig. 5.6b is the modified strategy which uses fibre noise cancellation method.

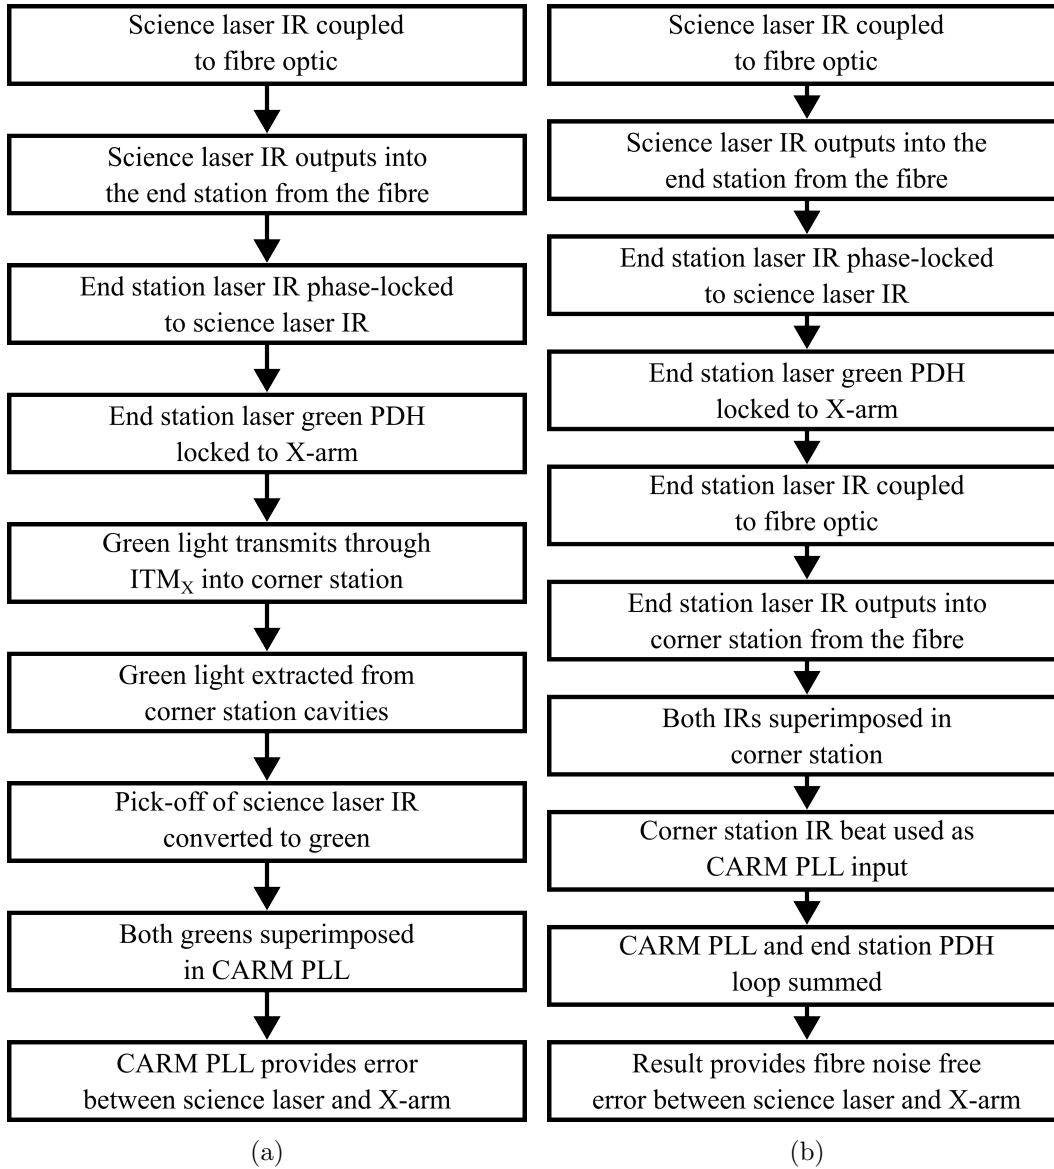


Figure 5.6: (a) A flowchart of the key steps to construct the science laser vs X-arm error. (b) Flowchart of the fibre cancellation method of constructing the science laser vs X-arm error.

5.3 The work undertaken

This project involved designing suitable fibre optic breadboard layouts to replace the existing ones at both the corner station and end stations. The major modification to the corner station was to include a photodetector for beating the output of the fibre which did not exist before. For the end station, the modification involved coupling light to the fibre to send it to the corner station. Both required a polarisation control method for maximising the beat signal.

The function of the breadboards was the same at both ends — to provide a connection to the 4 km fibre link, a point to superimpose the two lasers establishing a beat

signal, polarisation control to align the fields and an RF photodetector to measure the beat signal.

The work also required characterising the photodetector to ensure their electronic noise was sufficiently low, constructing and testing the breadboards in particular the polarisation control, installing the breadboards, and finally testing the new scheme. The new fibre cancellation method was implemented as the LSC lock scheme and the frequency noise between the science laser and the X-arm was compared to that of the previous scheme. Also, the correlation between the two signals containing the fibre noise was assessed in order to observe the frequency region showing most correlation i.e. where the fibre noise is equally measured and can be cancelled.

5.3.1 The breadboard(s) design

Fig. 5.7 shows the final layout of the fibre optic breadboards designed and installed at both the corner and X-end stations of LIGO Livingston. Dashed lines are within fibres, solid red lines are in free space.. Given the symmetry of the scheme, the breadboards at the corner and end stations can be identical.




Figure 5.7: The fibre breadboard layout. The splitters work as in [179] — each path is an input and an output. Each input on one side is split according to the ratio into the two paths at the other side. The remote laser is directed to the photodetector without interruption, a Faraday isolator (FI) prevents it reaching the local laser. The local laser is directed into the fibre optic and towards a fibre polarisation controller (which works via stress-induced birefringence produced by wrapping the fibre a certain number of times around a spool — twice for a quarter waveplate equivalence [180]) which introduces a fibre-quarter wave plate and fibre-retroreflector (F-Refl) in order to align the polarisation of the lasers. The retroreflector directs the local laser back towards the splitter, where it is then mixed with the remote laser and the mixed light is output to the photodetector (PD) via an in air half waveplate (HWP) and polarising beam splitter (PBS).

The idea behind the design is for the remote laser³ to pass through the splitter [179] and arrive at the photodetector. The local laser passes through the splitter and meets a fibre-quarter wave plate and a fibre-retroreflector. The fibre-quarter wave plate is a Thorlabs Polarization Controller [181], where the fibre is looped an appropriate number of times about a paddle, two and four for quarter and half waveplate equivalency, respectively; they work via stress-induced birefringence [180]. These paddles are shown in Fig. 5.12. The paddle can then be tilted to alter the

³This depends where the breadboard is located. For the breadboard installed at the corner station the remote laser is the end station laser. For the one installed at the X-end station, it is the science laser.

polarisation of the field within the fibre. The field then meets the fibre-retroreflector, passes again through the fibre-quarter wave plate after which it is mixed with the remote field within the splitter. The mixed fields then generate the heterodyne beat signal at the in-air photodetector. A half waveplate and a PBS between the fibre output and the in-air photodetector rejects light that is not aligned to the beat.

The fibre-quarter wave plate functionality is required in order to align the polarisation of the local and remote fields — if the fields were orthogonally polarised, there would be no beat signal at all. A maximum amplitude beat signal is achieved where the fields are aligned. Since the remote field is very susceptible to losses during the 4 km transit, the local field was chosen to pass twice through the splitter (since each time, 50% of the field is lost to the other output path).

The polarisation angle of the remote field is also not known given this 4 km transit. However, it is assumed to still be linear. It is therefore simply a matter of ensuring the local field can be aligned to match the remote field. The fibre-quarter wave plate, fibre-retroreflector, fibre-quarter wave plate sequence operates in the same way as a half waveplate (analysis showing this is in Appendix B), it was the only feasible design — seen at the bottom left side of the breadboard in Fig. 5.7. The transmission of the PBS, which is P-polarised light, is used for detection.

5.3.2 The photodetectors

The photodetectors were Newport 1811-FS types [143], which have a DC and an AC output. The transimpedance gain (provided by a resistor) of the DC output is $G_{PD-DC} = 1 \text{ k}\Omega$ and for the AC output it is $G_{PD-AC} = 40 \text{ k}\Omega$. The two noises expected from the photodetectors are:

1. ‘Photocurrent noise’. It is calculated as $\sqrt{2eIG_{PD-DC}}$ where e is the charge of the electron, I is the current at the photodetector, calculated by V/G_{PD-DC} , V is the voltage output by the photodetector and G_{PD-DC} is the transimpedance gain of the photodetector’s DC output.
2. ‘Johnson noise’. This is the thermal noise of the resistor providing the transimpedance gain, calculated as $\sqrt{4k_B T G_{PD-DC}}$ where k_B is Boltzmann’s constant and T is the temperature in Kelvin [68].

To analyse these, the DC output was used to calculate the power incident on the sensor while the AC output was used to measure a noise spectrum. The first step was to convert from DC volts to watts. The data of that test are in Fig. 5.8. A Thorlabs S132C [142] power meter was used for this analysis.

For ambient room temperature of 21°C this gets a Johnson noise of $\sim 4 \text{ nV}/\sqrt{\text{Hz}}$. These two noise sources can be equated to find the estimated voltage where shot-

noise equals Johnson noise

$$V_{DC} = \frac{2k_B T}{e} \quad (5.2)$$

assuming a resistor temperature of approximately 294.15 K (21 °C) we have an expected voltage of ~ 50 mV required to be shot-noise-limited. Going by the fit lines in Fig. 5.8 this equates to 0.5 mW. Above 0.5 mW and the shot-noise should surpass the photodetector Johnson noise meaning it is shot-noise limited.




Figure 5.8: Watts measured at the sensor vs DC volts output by the photodetector.

Each photodetector was characterised to find the shot-noise-limit for the 80 MHz region as shown in Fig. 5.9.




Figure 5.9: Photodetector shot noise limit test layout.

The AC output was mixed with a 80 MHz sine wave provided by a low noise signal generator, and the mixer result passed through an analogue DC — 1.9 MHz low-pass filter (Mini-Circuits[®] [182]). This *demodulation* isolates the 80 MHz region of the photodetector, the part we are most interested in.

A spectrum of 1-801 Hz with 800 bins and 10 averages was analysed for each incident power. The mean was then used as a measure of the noise. The result is shown in Fig. 5.10. It is evident that at low powers, increasing power has no effect on the mean noise at the 80 MHz region of the AC output. Above a certain power, the noise

steadily increases above the dark noise, the cross-over shows the minimum required power.

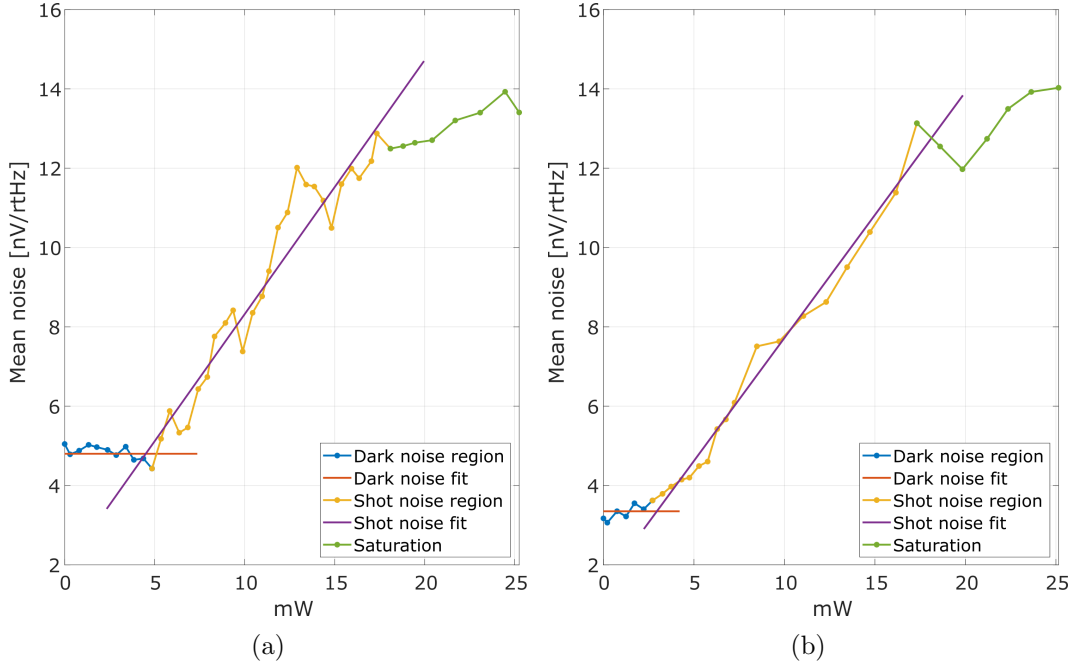


Figure 5.10: Photodiode shot noise limit required power. The blue section displays where the power is not sufficient to dominate the photodetector dark noise. Yellow is the shot-noise-limited region, where the Johnson noise has been surpassed and the shot-noise is the dominant noise. Green is the saturation power of the photodetector. (a) is the breadboard placed in the science laser enclosure, (b) is breadboard placed at the X-end station.

Compared to the science laser enclosure’s breadboard photodetector (Fig. 5.10a), the photodetector for the X-end station breadboard (Fig. 5.10b) is not as obvious at the low end. But the section of linearity is reasonably clear and the deviation at the low end must be due to dark noise.

The result tells us that a power of at least 5 mW is required in the beat signal to be in the shot-noise-limited regime. The power in the beat signal, when simply summing the fields and isolating the combined part, is

$$P_B = 2\sqrt{P_{local}P_{remote}} \quad (5.3)$$

which highlights the lower limit of the powers required at the mixing point.

5.3.3 Installing the breadboards

The breadboard in the science laser enclosure (Fig. 5.11a) was installed on the optical table the science laser itself is mounted on. The breadboard in the X-end station (Fig. 5.11b) was positioned on three *stilts*, given the lack of space on the optical table itself. The field inside a fibre is susceptible to phase and polarisation

fluctuations if the fibre is moved, hence the fibres are taped down quite thoroughly. The fibres used are polarisation maintaining single-mode types.




Figure 5.11: (a): The fibre breadboard installed in the science laser enclosure. (b): The fibre breadboard installed in the end station.

After installation, the beat at each end was maximised using the fibre-quarter wave plate.




Figure 5.12: Fibre-quarter waveplate paddle. The fibre is wrapped around the drum and fastened at either side of the controller to prevent it unravelling. 2 wraps are required for quarter waveplate, 4 for a half waveplate.

5.4 Results for the fibre noise cancellation scheme

The noise in the two schemes is analysed by measuring an amplitude spectral density of the common-mode arm length phase-locked-loop output signal, after the loops are closed. This shows the noise in the coupling between the science laser and the X-arm while the schemes are in operation. The original Advanced LIGO scheme (Sec. 5.1)

data was provided by archived measurements, it is the green trend in Fig. 5.13a. The fibre cancellation trial method data is the red trend of Fig. 5.13a. The trial method's noise is significantly worse than the original showing that the experiment was not successful.

The two signals which carry the fibre noise: the IR beat which was input to the common-mode arm length phase-locked-loop, and the end station Pound–Drever–Hall lock control signal are in Fig. 5.13b, magenta and blue, respectively. These data were recorded without the common-mode arm length loops being closed i.e. the X-end station laser green was locked to the X-arm and the corner station fibre breadboard measured the IR beat, but the feedback of the common-mode arm length control loop was not provided. This is recorded in order to assess any correlation between the signals; common noise features are expected to be attributed to the fibre noise.

The coherence of the two fibre noise carrying signals is shown in Fig. 5.13c while the phase between them is Fig. 5.13d. Where high coherence is accompanied by a 180° phase between the signals, then it can be assumed the common noise will be cancelled.

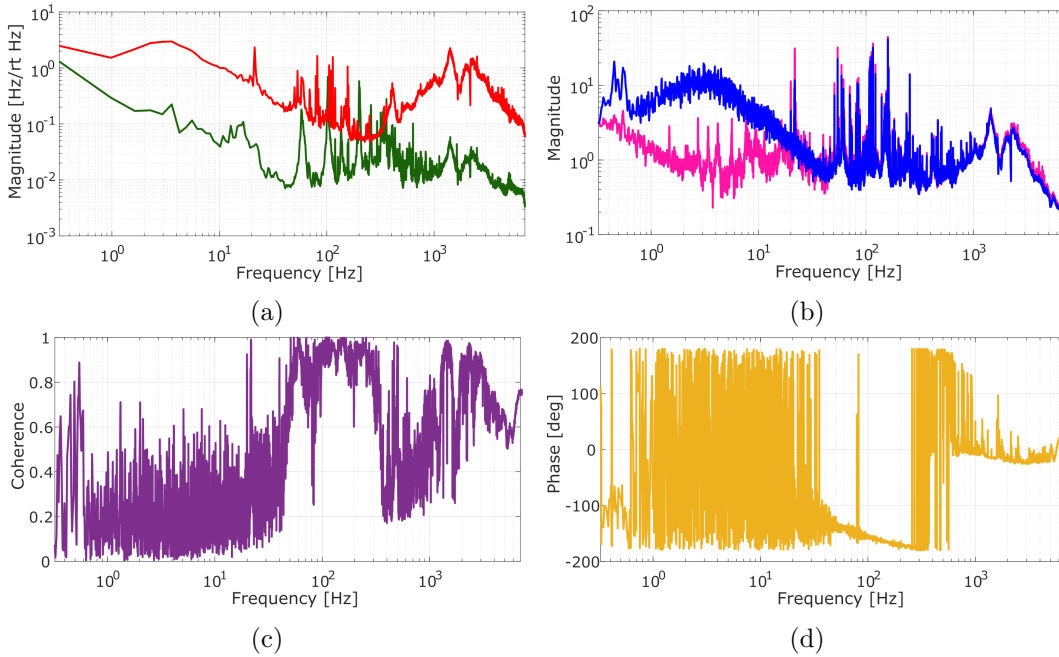


Figure 5.13: Subplot (a) shows the a noise spectra of the common-mode arm length phase-locked-loop output signal — the comparison between the science laser and the X-arm — using the original Advanced LIGO scheme (the science laser) in green (measured at point 1 of Fig. 5.2), and using the fibre cancellation scheme (the IR beat) in red (measured at point 2 of Fig. 5.5). Subplot (b) shows the noise spectra of the two signals used which contain the fibre noise: the corner station IR beat signal in magenta (measured at point 3 of Fig. 5.5) and the end station Pound–Drever–Hall control signal in blue (measured at point 4 of Fig. 5.5). Subplots (c) and (d) are the coherence and phase between the two signals in (b) respectively, where the Pound–Drever–Hall control signal is divided by the IR beat signal. 5.5 and 2 of fig:new drawing for ligo chapter

In plots Figs. 5.13c and 5.13d, the combination of high coherence and $\sim -180^\circ$

phase shows the region of 100 – 250 Hz to be the area of most success. This is partially reflected in the noise spectra in Fig. 5.13a — the trends are maximally comparable around 200 Hz (ignoring the low resolution area of < 1 Hz, though the phase and coherence are not alarmingly bad for ~ 0.5 Hz either). However, the noise in the fibre cancellation scheme (red in Fig. 5.13a) is significantly greater than the original scheme. This trial update has not succeeded.

The somewhat large hump shape which peaks at ~ 3 Hz in the blue trend in Fig. 5.13b is not yet fully understood. It is potentially due to a transient period of frequency instability in the auxiliary laser.

Ultimately, Fig. 5.13 shows that the fibre cancellation method has not been successful. The 100–250 Hz region does prove that there is a coherence between the signals however, and with some simple filter shaping, calibration or noise hunting, a wider bandwidth is possible. With an increased region of coherence between the signals, it should become comparable to the original scheme in terms of noise between the science laser and the X-arm. It is regrettable that 4-months was the length of the fellowship.

5.5 Summary

This chapter has described an investigation into a possible modification to the Advanced LIGO Livingston lock scheme. One of the key points of the scheme is to acquire an error between the science laser and the X-arm length. This is required in order to eventually bring the science laser and common-mode arm length to resonance. They are not in resonance from the start due to the strong coupling between the five longitudinal degrees of freedom (common- and differential-mode arm lengths, Michelson interferometer length, power- and signal-recycling cavity lengths) and the difficulties of using a single laser for them all. Decoupling the CARM and DARM length degrees of freedom is achieved using auxiliary lasers at the end stations. These auxiliary lasers are frequency coupled to the science laser via a fibre link and are subsequently made resonant in the arm cavities. These auxiliary lasers then keep the cavities off resonance to the science laser during lock acquisition. The fibre introduces too much frequency noise for a direct comparison between the science laser and the X-arm at this stage. The light transmitting into the corner station out of the arm cavity needs to be sampled, and is used as the X-arm reference. This light is then superimposed with the science laser creating a difference or beat signal between them, which provides the error between the science laser and X-arm.

The scheme is optically expensive, requiring many mirrors, lenses and dedicated controls to obtain and maintain the overlapped fields. It is required because of the fibre noise.

The modification has the goal of establishing two signals which contain the fibre noise; this creates a noise cancellation opportunity since the noise is sampled twice. It was trialed using only the X-arm. The end station Pound–Drever–Hall loop locking the end station laser to the arm cavity already contains the fibre noise, this remains in place and provides the first of the two signals. The new signal is generated by way of sending the end station laser back down the same fibre link towards the science laser. The science laser can now be directly beat against the end station laser, which is resonant in the X-arm, i.e. this is a science laser vs X-arm error signal which is corrupted by the fibre noise. There are now two signals which contain the fibre noise and they can be combined in order to cancel out the fibre noise. This is the essence of the update — obtaining the science laser vs X-arm beat in a manner that does not require sampling light from within the corner station cavities.

The modification was carried out to limited success. There is some evidence of noise cancellation in a small bandwidth, but it is not sufficient at this stage to be relied on as a new lock scheme for Advanced LIGO. Unfortunately the modified scheme did not succeed in reducing noise in the control loop locking the end station and science lasers, though high coherence between the corner station IR beat signal and the end station Pound–Drever–Hall control signal between 50 Hz and 300 Hz does indicate that the modified scheme could potentially be made to work in that frequency range. Results from KAGRA [177, 178] indicating that schemes of this type can be made to work would encourage further iterations on the modified scheme in future, perhaps focusing on adjustments to the loop gain of the servo in the modified scheme and renewed efforts to perfect the hardware where the fibres are combined, shown in Fig. 5.11.

Chapter 6

Conclusion and outlook

6.1 QUEST and the OMC

QUEST is a pair of co-located table-top power-recycled Michelson interferometers which will employ a DC-offset readout scheme. It aims to be the detector most sensitive to displacements in the 1 – 250 MHz bandwidth. It is also sensitive to entangled spacetime noise given the co-located design. Compared to other similar experiments, the Fermilab Holometer for example, it will have a higher power and therefore lower quantum noise floor. With that higher power comes the requirement to include an Output Mode Cleaner in order to reduce the noisy light, in the form of higher order modes, reaching the photodetectors. Without the OMC, a higher DC-offset would be required to ensure signal carrying light is dominant over the noisy light, and this will surpass the photodetector power threshold.

The work carried out during this PhD was primarily dedicated to QUEST and the OMC. Though QUEST (Chapter 3) is still being commissioned, there is a sensitivity curve result and a proven stable lock scheme with an automatic re-lock capability. With an input power of ≈ 357 mW, a circulating power of ≈ 142 W has been reached without use of the input amplifier. The goal circulating power is of the order of 10 kW. This will be achieved with a higher reflectivity power-recycling mirror and the use of an input power amplifier.

The OMC (Chapter 4) has been commissioned and characterised. It was designed to have a bandwidth of > 200 MHz and Finesse of > 11 . Measurements show a bandwidth of (218.507 ± 8.355) MHz, a Finesse of 36.612 ± 2.403 and a throughput of $\approx 98.8\%$. The dither locking scheme has been measured to have a unity gain frequency of 10 Hz and has shown successful lock of the order of days, with a working automatic re-lock function.

6.1.1 Future work — QUEST

There is still a lot of ongoing commissioning for QUEST. Some jobs are well understood and on the docket, others are potentials and some will undoubtedly spring out of their hidey-holes right at the least opportune moment.

The top priority is the angular to longitudinal drive coupling of the end test masses discussed in Sec. 3.5.4. Without this issue solved, it will not be possible to integrate the OMC in the lock acquisition process. The problem lies in the three PZT actuators used for longitudinal and angular controls, not the pico-motors which are attached to the mirror mount. These are used for coarse control and once alignment is achieved as best can be with them, they are no longer powered and so should not be contributing to these problems.

More comprehensive automation in the lock scheme will be needed sometime soon. There are several stages of lock, and with the OMC it could become quite complex. Many lock loops will need to interact and work in tandem and ideally this will eventually be a wholly automatic procedure. The operator chooses the lock level to reach and the loops do the rest, sequentially working through their internal logic to reach the desired point.

Potentially related to the angular coupling issue is the optical bench resonances. ‘The bench’ is actually three separate benches which are fixed together. But there has not yet been a thorough investigation into whether a detrimental amount of residual twisting/bending about the joins between the benches occurs.

The squeezer commissioning is ongoing, with the target the inject squeezed light by mid 2023.

The power-recycling mirror will need to be replaced with the higher reflectivity one to achieve the final sensitivity. This will require realignment and tuning of the loops/locking scheme.

The ionic vacuum pumps need to be commissioned. When the controllers arrive that will be done promptly. It is likely once they are operational, QUEST will be brought to vacuum and commissioning will continue under vacuum.

There is a possibility of including a Pre-Mode Cleaner on the injection breadboard. The Pre-Mode Cleaner is a triangular cavity used for filtering the laser of its inherent higher order modes prior to it entering the power-recycled Michelson interferometer. The principle is identical that of the OMC in Chapter 4, but the specifications are different. It is a larger cavity to the OMC (24 cm round-trip length vs the OMC’s 3.75 cm), has a higher Finesse (613 vs 36) and a substantially lower bandwidth (2 MHz vs 218 MHz). The Pre-Mode Cleaner has been assembled but it was not installed on the bench or characterised — it was originally assembled as a learning

tool in preparation for the OMC, but the OMC mirrors arrived in good time and so the Pre-Mode Cleaner was not used. Further, the beam shape in injection is so far very good, which suggests there is no real need for the Pre-Mode Cleaner at this stage (though this may change with higher power) and any higher order modes in the injection beam will be filtered by the power-recycled Michelson interferometer pole. However, the mode-matching telescope lenses for the Pre-Mode Cleaner are in place on the injection breadboard in case we do decide it is required.

The high-power amplifiers are still being commissioned. The one for AMY is not yet in position but with the ongoing work on BOB which involves its amplifier, it should be straightforward to install and operate AMY's when the time comes.

6.1.2 Future work — The OMC

The OMC has not yet been integrated into the QUEST experiment due to angular coupling to the longitudinal drive of the Michelson interferometer's end test masses. This is detailed in 3.5.4. The end test masses are driven by three separate PZT actuators who are responsible for length and angular degrees of freedom. Their unequal response to voltages causes the beam positions to jitter or wobble when they recombine at the BS where longitudinal controls are used as in the DC-lock loop. Aside from the impact that has on the DC-lock itself, it also causes the exiting beam position to be sufficiently unstable that aligning to the OMC was not possible.

When that angular coupling is no longer an issue, the light at the anti-symmetric port will be aligned to the OMC and a higher order mode suppression estimate can be carried out. During a DC-lock, the OMC length will be scanned in order to align and mode-match, this will provide data for throughput of the fundamental mode. With that and the power in the reflected light, the higher order mode content of the anti-symmetric port beam can be assessed. There is even potential to construct a second OMC and position it in reflection of the 'main' one, a length scan of that second OMC would split them higher order modes and their relative power could be estimated. This may not be necessary though.

With an OMC locked to the anti-symmetric port light, the merger of the lock schemes can begin. It is expected that the reflected and two transmission PDs will be summed, the sum of all three is roughly equivalent to a single PD directly in the path of the anti-symmetric port. They can then be used as the sensor for the DC lock as introduced in Sec. 3.5.2. This will however mean the OMC lock is dependent on the output behaviour of the power-recycled Michelson interferometer and as such, there will need to be a normalisation of power implemented. The OMC lock will need to determine if the OMC transmitted power has dropped because the OMC has lost resonance or because the DC-offset/anti-symmetric port output has fallen. This should not be too difficult to achieve, the transmission of the end test

masses is monitored by PDs and they are already used to estimate the power on the BS. This information can readily be used for some kind of threshold normalisation to prevent the OMC losing resonance if the DC lock has a hiccup.

The overall lock process also depends on the eventual contrast defect achieved by QUEST. There may be sufficient fundamental light for the OMC to lock at an early stage of the combined scheme. If the OMC can be locked after CARM is locked, which will always use the power-recycled Michelson interferometer reflection light in the Pound–Drever–Hall scheme, then the OMC transmission PDs alone can serve as the DARM sensors; as long as the OMC can remain on resonance as the power-recycled Michelson interferometer is driven to its DC-offset.

There is only one OMC at the moment. At least one more will need to be assembled including gluing the concave mirror to the PZT, soldering the PZT cables to a BNC connector and having a second metal housing unit made. Gluing and soldering will be done in the lab, the housing is fabricated at Cardiff University’s mechanical workshop. They should still have the 3-D drawing and so making a second should not take too long.

As well as the second OMC, the second detection breadboard has yet to be assembled. The component parts were bought and are ready to be positioned. The major effort will be solving the mode-matching telescope but that is quite routine practice once propagation distances are known.

There is potential to upgrade to an adaptive mode-matching telescope whereby either the lens focal lengths or their positions can be driven by a feedback loop. This was in question a year or so ago, but the variable focal length lenses were just far too expensive for whatever benefit they’d bring. Remote translation stages are possible however. This has not been seriously considered yet, in terms of thinking about how to actually implement adaptive mode-matching. But it could involve some kind of dither mechanic where perhaps the reflection PD is used to minimise reflection.

Adaptive mode-matching naturally leads onto the possibility of alignment control as well. This could involve using wavefront sensors as is quite common in gravitational wave detectors. Though both adaptive mode-matching and alignment control are more of a QUEST wide future work potential.

6.2 LIGO

The LIGO project carried out at LIGO Livingston was an investigation into a potential modification to the lock scheme. During lock, the main science laser (PSL) is frequency stabilised using the X-arm cavity while not being the resonant laser in that cavity. An auxiliary laser is used to lock the arm cavity and through resonance

in the arm, auxiliary light is available to mix with the PSL in order to generate an error between the PSL and the X-arm frequencies. This method involves PSL light being sent to the X-end station via a fibre link. The comparison between the PSL and the resonant auxiliary light is required because of the frequency noise of the fibre. A new method involved sending auxiliary light back down the fibre link towards the PSL. This establishes two samples of the fibre noise, and signals from either end of the link are combined to suppress the noise. This gives a comparison between the PSL and X-arm without the need of accessing light from within the corner station cavities. Success in this project would be a frequency match between the PSL and X-arm of comparable stability/noise to the original scheme.

The project was not successful but does show promise. A ~ 150 Hz bandwidth region of the coherence and phase between the two signals containing the fibre noise has the required features — high coherence while being out of phase, thus leading to a cancellation. But this feature was not strong enough, nor was it of sufficient bandwidth to enable the noise cancellation the project sought to establish.

6.2.1 Future work — LIGO

For this method to be successful it must be understood why such low a bandwidth is seen between the two signals which carry the fibre noise. This will require noise hunting and possibly some filter shaping.

A coherence measurement of standalone fibre lasers at either end could show the coherence one can expect from this fibre link.

Appendix

Appendix A

CDS and models in Simulink

A.1 CDS

QUEST will be controlled almost entirely digitally; CDS is responsible for this. CDS is an acronym for the LIGO 'Control and Data Systems'. The hardware, firmware and software is a copy of a portion of that developed for LIGO by Rolf Bork, Alex Ivanov, Jay Heefner and others. In this thesis, CDS is the collective term used when describing any of the computing done for control and data acquisition unless otherwise stated, for example when using the Moku:Lab or an oscilloscope. QUEST's scientific data is acquired with a stand alone separate system and is described in Sec. 3.4.14. Any mention of "data acquisition" will refer to technical data related to control loops and such.

As well as the computers, CDS includes the interface between the analog real world, to the digital computer world, and back out to the analog. External devices such as the photodetectors provide inputs to CDS via the analog to digital converter (ADC) cards. Then the CDS outputs are sent to devices such as the PZT actuators via the digital to analogue (DAC) cards. The ADC has an associated input quantisation noise since computers only have a certain bit rate. The computer cannot store or record continuous data, it must do so in discrete steps, this introduces noise. This quantisation or ADC noise is a white noise, i.e. it is flat and not frequency dependent. If an input's signal power is below the input quantisation noise of the CDS then when being output by the DAC, the signal now includes the ADC noise (see Fig. A.1 top). Whitening filters are essentially high-pass filters to boost frequencies above the input noise. This is done prior to the analogue to digital conversion and so the ADC noise does not pass through the filter, it remains flat. To maintain signal integrity, this of course needs to be undone and so de-whitening is the reverse process (see Fig. A.1 bottom). The two filters cancel out for the data, but the ADC noise only gets passed through the de-whitening stage. This has the effect of lowering the ADC

noise and *hopefully* bringing it below the output signal.

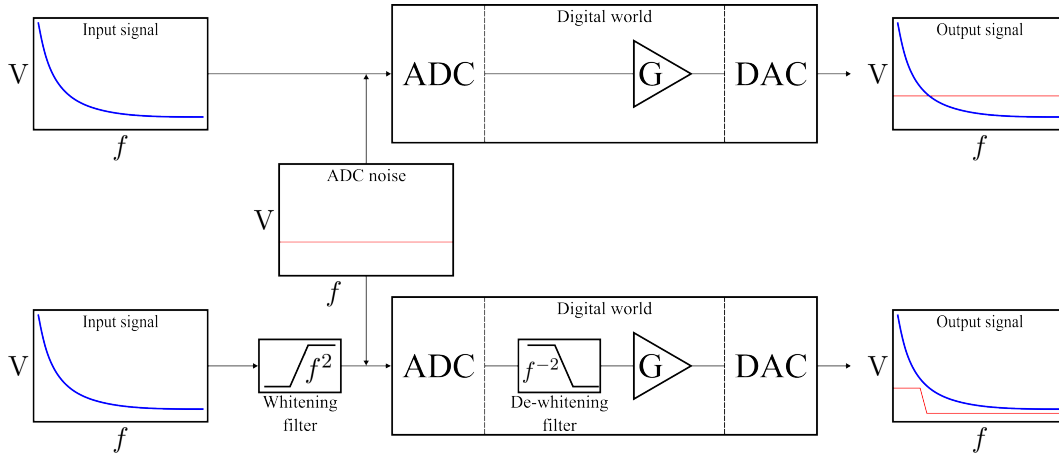


Figure A.1: Top path: Signal flow of ADC to DAC without any whitening and de-whitening filters. The analogue input gets digitised with some discrete ‘bit’ value which introduces an ADC noise. If the noise is higher than the signal then at the DAC output, some of the signal is lost to the noise.

Bottom path: Signal flow of ADC to DAC with whitening and de-whitening filters. The analogue input gets high-pass filtered prior to conversion to digital. Both the signal and the ADC noise then get low-pass filtered before reaching the DAC. This establishes a suppression of the ADC noise, and a recovery of the original signal shape, which is now not masked by the ADC noise.

The Nyquist frequency $f_{nyquist}$ is half of the analyser sampling rate f_{sample} . For our CDS, $f_{sample} = 64$ kHz. The principle of the Nyquist cut-off is in Fig. A.2. Signals above $f_{nyquist}$ are filtered out since they can introduce noise and give erroneous readings. This kind of filtering is anti-aliasing. The *opposite* of this filtering is anti-imaging and it is performed at the output.




Figure A.2: Black vertical lines represent the sample clock, where the signal intersects the line, a sample is taken. Red lines are the interpolated waveform the analyser creates from the sample measurements. When the sample frequency is greater than the signal frequency, the sample is measured accurately. When the sample and signal frequencies are equal it can be dangerous, the two extremes are shown: the first is the ideal case where the frequency and amplitude are measured correctly while the second is the worst case; a DC signal is measured. In-between these two examples, the frequency will be measured correctly but the amplitude will be an underestimate. Where the sample frequency is only just lower than the signal frequency, some amplitude beating will appear as the sample is taken at different phases of the signal. If the sample frequency is much less than the signal frequency then an incorrect frequency is measured. This effect is aliasing.

The majority of outputs the CDS drives are for the PZT actuators used for mirror displacement control. The actuators tend to have a much larger voltage range than CDS can provide. The normal DACs have an output range of ± 10 V with 2^{16} bits or counts; there is a step precision of $\sim 305.16 \mu\text{V}/\text{count}$. The high voltage DAC = $6 \times$ normal DAC + 50 V, and thus provides an output range of -10 V to $+110$ V with a step precision of ~ 1.83 mV/count; it is reserved primarily for the mirror PZTs since so far, no other device requires such voltages. Additional CDS outputs include the laser temperature controllers, and the remote shutters for the detection breadboards which operate as simple logic switches which require a DC voltage (above ~ 4 V) to remain open.

A.2 SIMULINK

Simulink is the Matlab-based software used for signal processing and designing the signal flow/logic of the control schemes. It provides the interface between the analogue real world to the digital computer/software world.

‘The model’ is the phrase used to refer to a specific control system’s Simulink model. For example, the OMC dither lock has its own model. QUEST’s length sensing and control system has its own model. QUEST’s angular control also does, but it shares some paths with the length sensing one because they use the same actuators — but in terms of a feedback loop for angular control, it is a separate model.

For each system there is a common/master model and two Michelson interferometer specific models. The master model is where the logic of the system is designed, this is where the actual workings of the loop is made. The Michelson interferometer specific models then take a ‘hyperlink’ version of that master model and tell it which physical inputs and outputs to actually use. In this way, the two Michelson interferometers are using the same processing, the same logic etc, but have distinct physical inputs and outputs. This also makes updating systems universal. If the logic of a control scheme is modified due to some new equipment or simply a different idea for a method, it occurs in the master model and is then automatically applied to the Michelson interferometer specific ones.

Fig. A.3 is AMY’s specific OMC model. It shows the selected inputs and output channels connecting to the hyperlinked master OMC model.

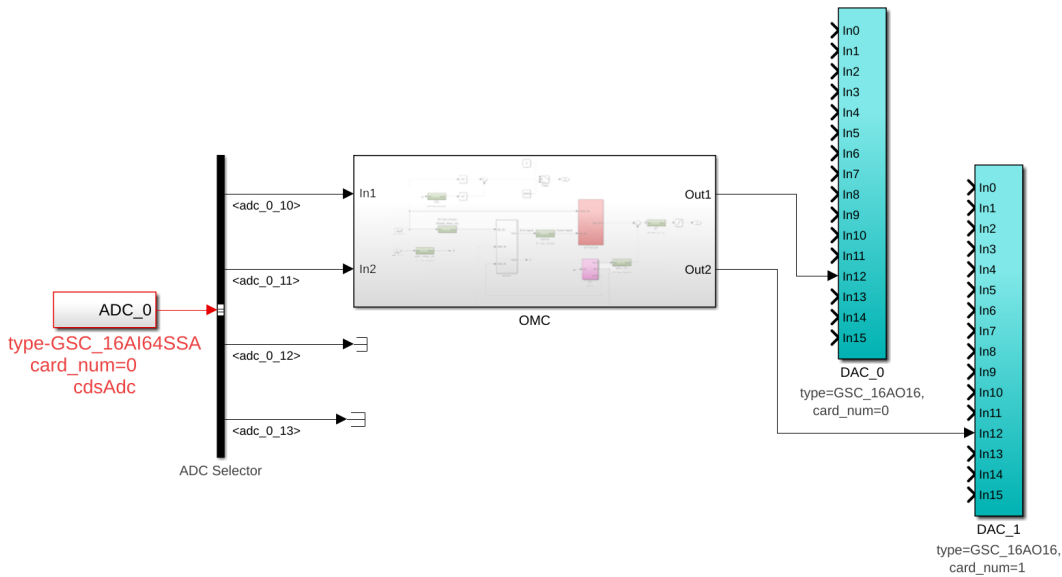


Figure A.3: Michelson interferometer 1, AMY, specific model. The central box labelled OMC is a hyperlinked insert of the common OMC model. This model serves to connect the relevant inputs and outputs of the CDS. The black bus bar on the left is the ADC selector, here the inputs labelled `<adc_0_10>` etc. identify the physical inputs of the ‘number 0’ ADC card. Two of them are terminated because they are not yet in use, these are reserved for the second transmission photodetector and a spare for any future requirements. The inputs that are connected are for the transmission photodetector (10) and the reflection photodetector (11). The teal boxes on the right are the DAC selectors. DAC_0 is the high voltage output and so the connection to that card is for the PZT. DAC_1 is the ‘normal’ voltage output, this is for the remote shutter. The BOB specific model looks identical, the only changes are the ADC and DAC port numbers.

Fig A.4 shows the OMC master model. There are currently two inputs (numbers 1 and 2 on the left side), input 1 is the transmission photodetector, input 2 is the reflection photodetector. There are two outputs, 1 is to the OMC PZT and 2 is to the remote shutter. The green boxes are *filter banks*, they generate points of interaction with the model. These are added wherever the user wants to add some digital filters or DC gain, have the ability to block or pass a signal at that point, add an offset, introduce an excitation, or simply observe the signal. Each filter bank has the input and output side. Input is the signal at the point it is input, output is the signal after whatever filters and gains installed to that filter bank are applied.

The reflection photodetector is currently only used for monitoring, hence the REFL_PDA1_DC filter bank connected to a terminator — the filter bank itself generates the path which can be monitored but it must be connected to something for the model to work, this is why signals which are not otherwise used are terminated. The transmission photodetector is used in three separate paths: monitoring for remote shutter triggering (to block the beam reaching the photodetectors if power increases above a given threshold), as a monitor for the resonance condition being held, and to complete the feedback control loop.

The OSC box is the oscillator which generates the dither signal. CLK is the *raw*

dither signal, this is given a frequency and an amplitude, it is sent directly to the PZT. The SIN and COS outputs of OSC are clones of the CLK signal (with the associated phase delay) and are fed only to the DEMOD block. This DEMOD block is a LIGO tool where the demodulation phase can be set. Within this block, the signal mixing and rotation matrix is applied to isolate and maximise the component of the signal linearly proportional to the error (see Sec 2.10.1). Out of this block comes the error signal, which can be monitored using the SERVO filter bank. SERVO is the principle controller of the loop. This is where major filters are installed and the error signal is transformed into the control signal. The other output of DEMOD is the minimised error signal, it is not used hence it is terminated.

The AUTOSCAN block is a custom design. The interior working is explained in Fig. A.6 but basically, there is a user defined threshold for the transmitted power. This is monitored by the Check_IN input, one of the transmission photodetector paths. If this is above the threshold, the OMC is considered to be on resonance and the control signal from the SERVO is passed to the output. If the transmission photodetector is below the threshold, the OMC is considered to be away from resonance, and a slow scan is sent to the PZT instead. This slow scan pushes the concave mirror $\sim 2 \nu_{FSR}$ in order to reach a resonance. Once the threshold power is reached, the scan is blocked and the control signal is resumed to hold resonance.

The ACT filter bank of Fig. A.4 is there to facilitate an ultimate ON/OFF switch to cut signals reaching the PZT, and as a monitor of the full loop output signal. The “S” shaped block just before output 2 is a fixed limiter. It restricts the output between two values. It is there to protect the PZT since it cannot take negative voltages (see Fig. 4.14).

The remote shutter path is a simple logic gate. The Trigger box is a switch which passes either the ‘0’ or ‘-6000’ value, which one depends on the output of the sum junction just before it. The TRIG filter bank allows the user to change the trigger value, this is then subtracted from the transmission photodetector value. If the result is negative, i.e. the transmission is below the trigger value, the ‘-6000’ choice is passed to output 2 which physically outputs ~ 1.5 V to the remote shutter, holding it open. If the result of the sum junction flips to positive because the transmission exceeds the trigger, ‘0’ is passed to output 2 and 0 V is sent to the shutter, which closes it.

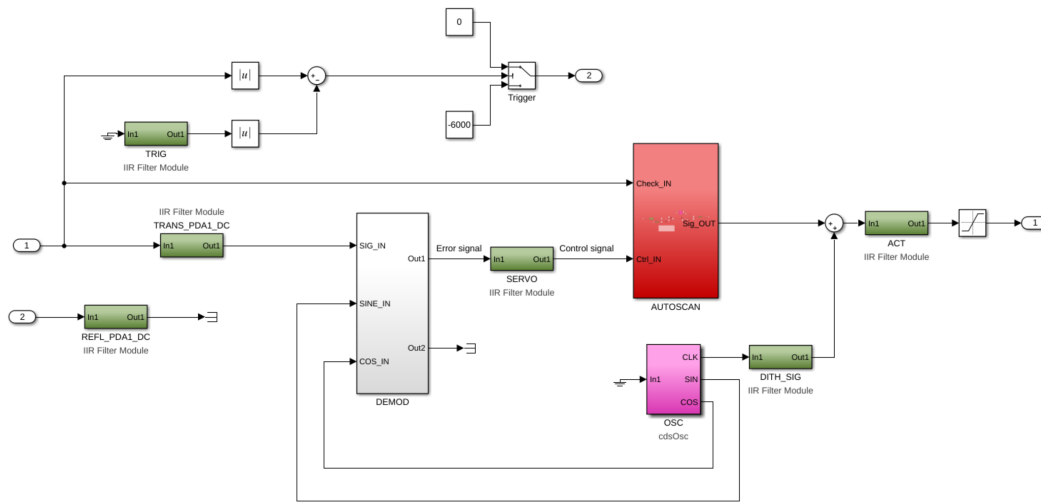


Figure A.4: Simulink model for OMC lock. Input 1 is the transmission photodetector. Input 2 is the reflection photodetector, which is only used for monitoring. Trans photodetector is split into three paths. Path one goes to the remote trigger; if power reaches a threshold, the shutter is triggered to close, blocking light from reaching the photodetector. Path two is to the AUTOSCAN box (shown in Fig. A.6) for resonance condition monitoring. Path three is for feedback control. The DEMOD box is where the oscillator’s dither signal is demodulated for error extraction. SERVO is the main loop controller. This is where the custom filters are installed for the loop to operate effectively. The output of the SERVO is the control signal, this is passed through the AUTOSCAN box in order to be blocked or passed to the PZT depending on resonance.

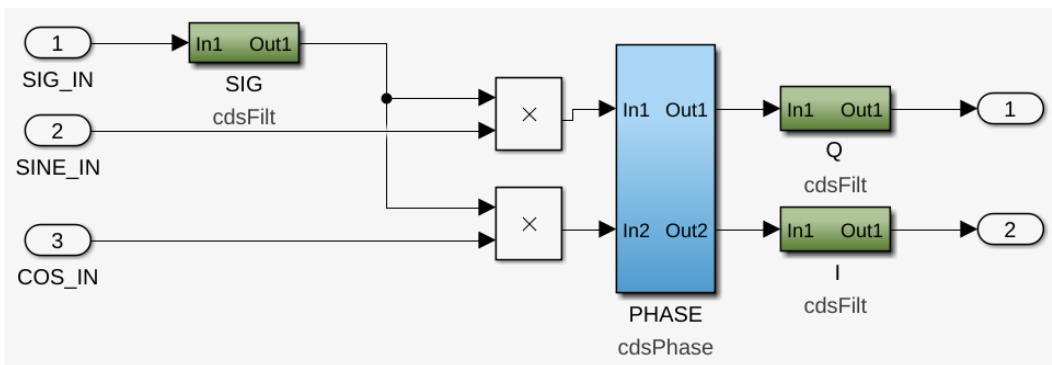


Figure A.5: The interior of the DEMOD box. Filter banks ‘Q’ and ‘I’ facilitate adding the demodulation low pass filter as well as ensuring each has a separate analysis point.

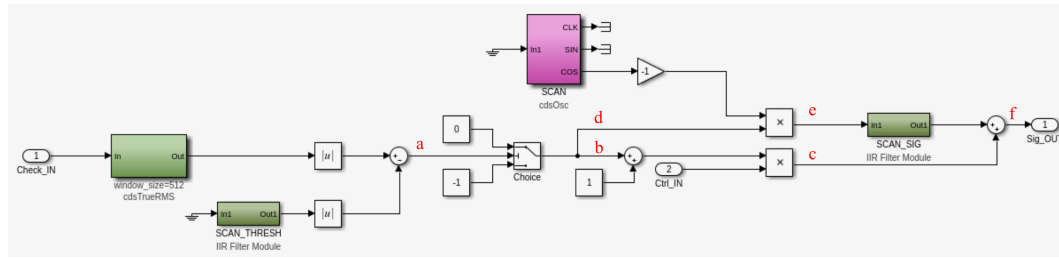


Figure A.6: The AUTOSCAN box of Fig. A.4. Check_IN is the transmission photodetector sample. SCAN_THRESH is a manually selected value. These values are summed and their difference sent to the choice switch. If (at point a) photodetector < SCAN_THRESH, 0 is passed by the choice switch, otherwise -1 is passed. If 0 is passed (following path b), 1 is multiplied by the control signal (Ctrl_IN) and it is allowed to continue to the PZT. It is also (following path d) multiplied directly against the SCAN signal, blocking it from reaching the PZT. If -1 is passed by the choice switch (following path b), the control signal is now multiplied by 0 and blocked from reaching the PZT. It is also (following path d) sent to be multiplied by the SCAN oscillator output, thus passing the signal to the PZT. The -1 multiplier for the SCAN output is simply there to correct for the -1 multiplier from the choice box.




Figure A.7: A filter bank, these are automatically generated from each of the green boxes (with cdsFilter written beneath them) of the Simulink models. The signal flows from left to right. INMON is the input value, this can be prevented from passing through using the left most ON/OFF button. EXCMON is the excitation being specifically sent to that filter bank. OFFSET is a manual value which can be turned on or off. It will be added to the INMON and EXCMON values and passed through to the filters. The FM1 - 10 central boxes are where the filters are installed. They can be individually turned on or off. The value inside the black triangle (here shown as 1.000) is the (linear, not dB) DC gain of the filter bank. Each of the FM1 - 10 filters have their own associated gain, but here the user can quickly input a gain for the entire filter bank. LIMIT restricts the output to the value set, Ramp Time is how long the filter bank will take to gradually apply changes made (such as a new DC gain). The right most ON/OFF is where the user can block all outputs from leaving the filter bank. OUTPUT is the output actually being sent out. OUTMON is the output that *would be* sent out if ON/OFF is OFF. The pale vertical arrows, IN1, IN2 and OUT are additional test points, these can be observed and analysed for various purposes.

Each system requires its own model which is designed to accomplish the needs of the loop. The models for QUEST contain similar logic and processes as shown in the models for the OMC dither lock.

The simulink models and CDS are inexorably linked. Where signal analysis or monitoring is required, the specific paths of the filter bank test points are what are used. Likewise, when injecting an excitation into a loop for transfer function measurement for example, the filter bank EXC is the point it is added.

Appendix B

Quarter waveplate argument

Derivation showing how a quarter waveplate and a mirror can be used in place of a single half waveplate.

The derivation relies on Jones calculus formulation [183] where an electric field can be represented as a 2x1 matrix which contains information about the polarisation direction only. For example, a plane EM wave of amplitude E_0 travelling in the z direction, the electric field is

$$\vec{E} = \hat{x}E_{0x}e^{i(\omega t - kz)} + \hat{y}E_{0y}e^{i(\omega t - kz)} \quad (\text{B.1})$$

where \hat{x} and \hat{y} are the amplitude of the field components aligned in the x and y directions respectively. When concerned only with polarisation direction, the Jones vector for this field is

$$\vec{E}_{J,\theta_x} = \begin{pmatrix} \cos(\theta_x) \\ \sin(\theta_x) \end{pmatrix} \quad (\text{B.2})$$

where θ_x is the polarisation angle with respect to the x axis. We can do this because the information that the field is travelling is not important, only the relative polarisation direction.

Waveplates work by using the property of birefringence, which means a substance (normally a crystal) has a particular refractive index in one direction or axis, and a different refractive index in the perpendicular axis. With these axes, a new frame can be defined i.e. one axis is the waveplate's x-axis and the other is the y-axis and due to the birefringence, $n_x \neq n_y$. The definition of refractive index is $n = c/v$ and so this birefringence results in a difference in the speed of the light through the axes of the waveplate. If $n_x < n_y$, then the x-axis is the 'fast' axis and the y-axis is the 'slow' axis. The field in the fast axis will have an increased wavelength with respect to the field in the slow axis; or equivalently, a phase delay is introduced between the perpendicularly polarised fields.

Half waveplates introduce a phase delay of half a wavelength between the fields. With linearly polarised light going in, this has the effect of rotating the field's axis of polarisation by flipping about the waveplate's fast axis. For a quarter waveplate a quarter wavelength phase delay is introduced which results in linear polarised light becoming some form of elliptically polarised light — in the right circumstances, circularly polarised or linearly polarised light can be output but these are just examples of the extreme ends of the ellipse geometry.

For the half waveplate and quarter waveplate, the transformation matrices are J_{HWP} and J_{QWP} , respectively [183]

$$J_{HWP} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (\text{B.3})$$

$$J_{QWP} = \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix} \quad (\text{B.4})$$

When this field passes through a waveplate, the new polarisation is calculated as

$$J_{WP} \cdot \vec{E}_J \quad (\text{B.5})$$

however, in order to account for the angle of the waveplate's fast axis relative to the frame in which the field is measured against, the rotation matrix J_{Rot}

$$J_{Rot}(\theta) = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \quad (\text{B.6})$$

must be used in sequence [183]

$$J_{WP}(\theta) = J_{Rot}(-\theta) \cdot J_{WP} \cdot J_{Rot}(\theta) \quad (\text{B.7})$$

The practice of rotating the frame to the new axis and back to the original axis is quite common. $J_{WP}(\theta)$ first calculates how much of the input field aligns with the waveplate's axes, and then $J_{Rot}(-\theta)$ calculates how much of that remaining field now aligns back with the original axis.

This results in the Jones vectors for waveplates at some angle θ [183]

$$J_{HWP}(\theta) = \begin{pmatrix} \cos(2\theta) & \sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{pmatrix} \quad (\text{B.8})$$

$$J_{QWP}(\theta) = \begin{pmatrix} \cos^2(\theta) - i \sin^2(\theta) & \cos(\theta) \sin(\theta) + i \sin(\theta) \cos(\theta) \\ \sin(\theta) \cos(\theta) + i \cos(\theta) \sin(\theta) & \sin^2(\theta) - i \cos^2(\theta) \end{pmatrix} \quad (\text{B.9})$$

The next component for this proof is the mirror, which has a simple Jones vector

[183]

$$J_M = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \quad (\text{B.10})$$

The argument is whether a linearly polarised field can pass through the quarter waveplate, get reflected and then pass again through the same quarter waveplate and be output at linearly polarised light at some new angle relative to the input, as if it had passed once through a half waveplate. The proof simply requires multiplying the relevant Jones vectors in sequence, as $J_{QWP}(-\theta) \cdot J_M \cdot J_{QWP}(\theta)$ where the $-\theta$ is because the field is now coming at it from the other direction.

$$\begin{aligned} & \begin{pmatrix} \cos^2(-\theta) - i \sin^2(-\theta) & \cos(-\theta) \sin(-\theta) + i \sin(-\theta) \cos(-\theta) \\ \sin(-\theta) \cos(-\theta) + i \cos(-\theta) \sin(-\theta) & \sin^2(-\theta) - i \cos^2(-\theta) \end{pmatrix} \\ & \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \cos^2(\theta) - i \sin^2(\theta) & \cos(\theta) \sin(\theta) + i \sin(\theta) \cos(\theta) \\ \sin(\theta) \cos(\theta) + i \cos(\theta) \sin(\theta) & \sin^2(\theta) - i \cos^2(\theta) \end{pmatrix} \\ & = \begin{pmatrix} -\cos(2\theta) & -\sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{pmatrix} = J_{QWP_{DP}}(\theta) \quad (\text{B.11}) \end{aligned}$$

where $J_{QWP_{DP}}(\theta)$ refers to the quarter waveplate ‘double pass’. The magic of Matlab’s simplify function...

Comparing the same field passing through the half waveplate and the double pass quarter waveplate

$$J_{HWP}(\theta) \cdot \vec{E}_{J,\psi} = \begin{pmatrix} \cos(2\theta) & \sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{pmatrix} \begin{pmatrix} \cos(\psi) \\ \sin(\psi) \end{pmatrix} = \begin{pmatrix} \cos(2\theta - \psi) \\ \sin(2\theta - \psi) \end{pmatrix} \quad (\text{B.12})$$

$$J_{QWP_{DP}}(\theta) \cdot \vec{E}_{J,\psi} = \begin{pmatrix} -\cos(2\theta) & -\sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{pmatrix} \begin{pmatrix} \cos(\psi) \\ \sin(\psi) \end{pmatrix} = \begin{pmatrix} -\cos(2\theta - \psi) \\ \sin(2\theta - \psi) \end{pmatrix} \quad (\text{B.13})$$

So, passing linear light through a quarter waveplate twice via reflection is the same result as passing it through a half waveplate once with the difference being an additional flip about the y-axis. The main result is that the output is linear and at a different angle as the input; proving the validity of using a quarter waveplate and a mirror to control the angle of linearly polarised light.

Appendix C

Mechanical dimensions

The following sections give the dimensions for the 3D CAD items designed during this PhD.

C.1 Beam dumps

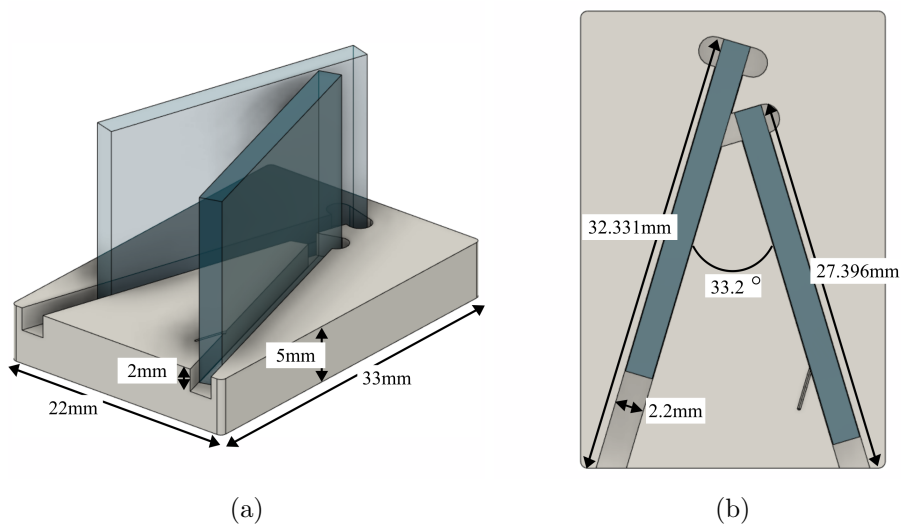


Figure C.1: The beam dump platform dimensions. The 2 mm thick glass itself was bought as 50x50 mm squares and cut into 2.5x1.66 mm peices using a diamond tipped, water lubricated saw.

C.2 Gluing jig




Figure C.2: The gluing jig dimensions. This keeps the spacer, PZT and concave mirror centralised while gluing. The arms are designed to be thin and spaced such that the glue does not touch the jig itself; so it can be left in place as the glue cures.

C.3 OMC housing

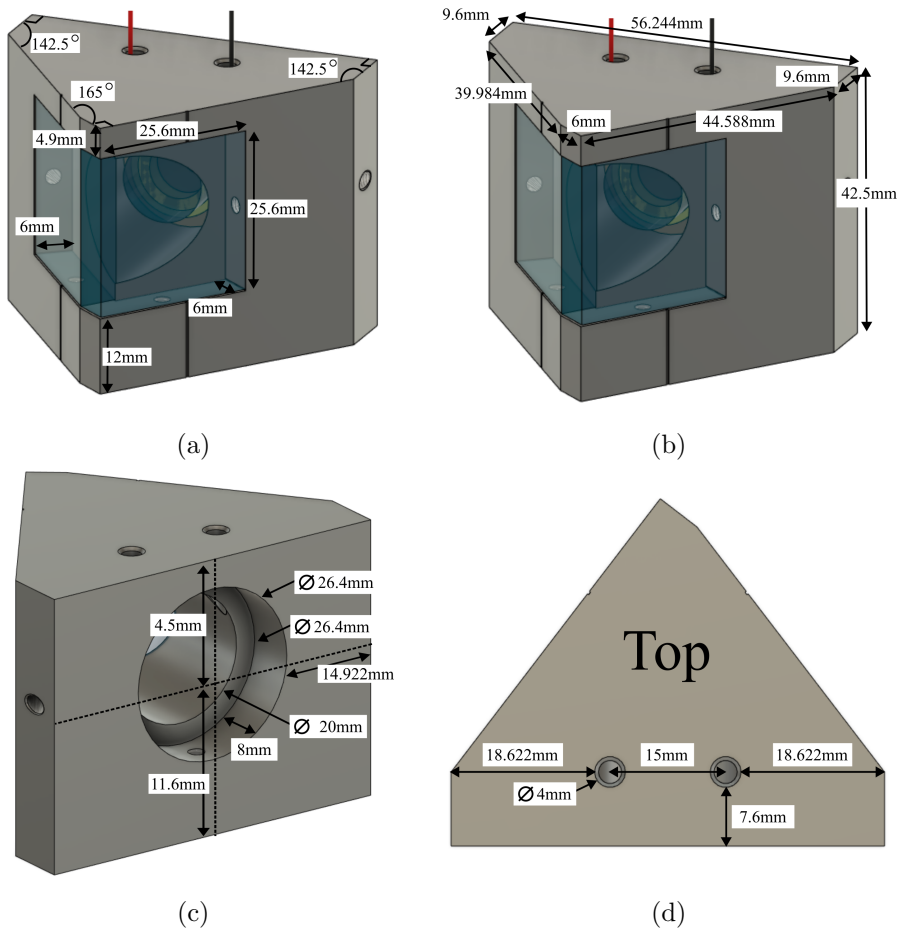


Figure C.3: The OMC housing dimensions.

Bibliography

- [1] Thomas Young. I. The Bakerian Lecture. Experiments and calculations relative to physical optics. *Philosophical Transactions of the Royal Society of London*, 94:1–16, January 1997. doi: 10.1098/rstl.1804.0001. URL <https://royalsocietypublishing.org/doi/10.1098/rstl.1804.0001>. Publisher: Royal Society.
- [2] A. J. Dempster and H. F. Batho. Light Quanta and Interference. *Physical Review*, 30(5):644–648, November 1927. doi: 10.1103/PhysRev.30.644. URL <https://link.aps.org/doi/10.1103/PhysRev.30.644>. Publisher: American Physical Society.
- [3] Claus Jönsson. Electron Diffraction at Multiple Slits. *American Journal of Physics*, 42(1):4–11, January 1974. ISSN 0002-9505. doi: 10.1119/1.1987592. URL <https://aapt.scitacion.org/doi/10.1119/1.1987592>. Publisher: American Association of Physics Teachers.
- [4] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354:769–822, January 1916. ISSN 0003-3804. doi: 10.1002/andp.19163540702. URL <https://ui.adsabs.harvard.edu/abs/1916AnP...354..769E>. ADS Bibcode: 1916AnP...354..769E.
- [5] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. OUP Oxford, Oxford, illustrated edition edition, October 2007. ISBN 978-0-19-857074-5.
- [6] Pierre Giacomo. The Michelson interferometer. *Microchimica Acta*, 93(1):19–31, January 1987. ISSN 1436-5073. doi: 10.1007/BF01201680. URL <https://doi.org/10.1007/BF01201680>.
- [7] LIGO Scientific Collaboration and Virgo Collaboration. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6):061102, February 2016. doi: 10.1103/PhysRevLett.116.061102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102>. Publisher: American Physical Society.

-
- [8] The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run, November 2021. URL <http://arxiv.org/abs/2111.03606>. arXiv:2111.03606 [astro-ph, physics:gr-qc].
- [9] The LIGO Scientific Collaboration and the Virgo Collaboration. GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, May 2022. URL <http://arxiv.org/abs/2108.01045>. arXiv:2108.01045 [gr-qc].
- [10] The LIGO Scientific Collaboration and the Virgo Collaboration. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *Physical Review X*, 11(2):021053, June 2021. ISSN 2160-3308. doi: 10.1103/PhysRevX.11.021053. URL <http://arxiv.org/abs/2010.14527>. arXiv:2010.14527 [astro-ph, physics:gr-qc].
- [11] The LIGO Scientific Collaboration and the Virgo Collaboration. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Physical Review X*, 9(3):031040, September 2019. ISSN 2160-3308. doi: 10.1103/PhysRevX.9.031040. URL <http://arxiv.org/abs/1811.12907>. arXiv:1811.12907 [astro-ph, physics:gr-qc].
- [12] Davide Castelvecchi and Alexandra Witze. Einstein’s gravitational waves found at last. *Nature*, February 2016. ISSN 1476-4687. doi: 10.1038/nature.2016.19361. URL <https://www.nature.com/articles/nature.2016.19361>. Publisher: Nature Publishing Group.
- [13] B. P. Abbott et al. ASTROPHYSICAL IMPLICATIONS OF THE BINARY BLACK HOLE MERGER GW150914. *The Astrophysical Journal Letters*, 818(2):L22, February 2016. ISSN 2041-8205. doi: 10.3847/2041-8205/818/2/L22. URL <https://dx.doi.org/10.3847/2041-8205/818/2/L22>. Publisher: The American Astronomical Society.
- [14] Abhay Ashtekar. Introduction to Loop Quantum Gravity. *arXiv:1201.4598 [gr-qc, physics:hep-th, physics:math-ph]*, 863:31–56, 2013. doi: 10.1007/978-3-642-33036-0_2. URL <http://arxiv.org/abs/1201.4598>.
- [15] E. Kiritsis. Introduction to Superstring Theory. *arXiv:hep-th/9709062*, March 1998. URL <http://arxiv.org/abs/hep-th/9709062>.
- [16] Sabine Hossenfelder. Minimal Length Scale Scenarios for Quantum Gravity. *Living Reviews in Relativity*, 16(1):2, December 2013. ISSN 2367-3613, 1433-8351. doi: 10.12942/lrr-2013-2. URL <http://arxiv.org/abs/1203.6191>.
-

- [17] S.W. Hawking. Spacetime foam. *Nuclear Physics B*, 144(2-3):349–362, November 1978. ISSN 05503213. doi: 10.1016/0550-3213(78)90375-9. URL <https://linkinghub.elsevier.com/retrieve/pii/0550321378903759>.
- [18] Giovanni Amelino-Camelia. A phenomenological description of quantum-gravity-induced space-time noise. *Nature*, 410(6832):1065–1067, April 2001. ISSN 0028-0836, 1476-4687. doi: 10.1038/35074035. URL <http://arxiv.org/abs/gr-qc/0104086>.
- [19] Giovanni Amelino-Camelia. Quantum Spacetime Phenomenology. *Living Reviews in Relativity*, 16(1):5, December 2013. ISSN 2367-3613, 1433-8351. doi: 10.12942/lrr-2013-5. URL <http://arxiv.org/abs/0806.0339>.
- [20] Y. Jack Ng and H. van Dam. Measuring the foaminess of space-time with gravity-wave interferometers. *Foundations of Physics*, 30(5):795–805, 2000. ISSN 00159018. doi: 10.1023/A:1003745212871. URL <http://arxiv.org/abs/gr-qc/9906003>.
- [21] Y. Jack Ng and H. Van Dam. LIMIT TO SPACE-TIME MEASUREMENT. *Modern Physics Letters A*, 09(04):335–340, February 1994. ISSN 0217-7323, 1793-6632. doi: 10.1142/S0217732394000356. URL <https://www.worldscientific.com/doi/abs/10.1142/S0217732394000356>.
- [22] C. J. Hogan. Measurement of Quantum Fluctuations in Geometry. *Physical Review D*, 77(10):104031, May 2008. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.77.104031. URL <http://arxiv.org/abs/0712.3419>.
- [23] S. Deser and P. van Nieuwenhuizen. Nonrenormalizability of the quantized Dirac-Einstein system. *Physical Review D*, 10(2):411–420, July 1974. doi: 10.1103/PhysRevD.10.411. URL <https://link.aps.org/doi/10.1103/PhysRevD.10.411>.
- [24] Tao Han and Scott Willenbrock. Scale of Quantum Gravity. *Physics Letters B*, 616(3-4):215–220, June 2005. ISSN 03702693. doi: 10.1016/j.physletb.2005.04.040. URL <http://arxiv.org/abs/hep-ph/0404182>.
- [25] C. P. Burgess. Quantum Gravity in Everyday Life: General Relativity as an Effective Field Theory. *Living Reviews in Relativity*, 7(1):5, December 2004. ISSN 2367-3613, 1433-8351. doi: 10.12942/lrr-2004-5. URL <http://arxiv.org/abs/gr-qc/0311082>.
- [26] R. Bousso. The holographic principle. *Reviews of Modern Physics*, 74(3):825–874, August 2002. ISSN 0034-6861, 1539-0756. doi: 10.1103/RevModPhys.74.825. URL <http://arxiv.org/abs/hep-th/0203101>.

-
- [27] Erik P. Verlinde and Kathryn M. Zurek. Observational Signatures of Quantum Gravity in Interferometers. *arXiv:1902.08207 [gr-qc, physics:hep-ph, physics:hep-th]*, February 2019. URL <http://arxiv.org/abs/1902.08207>.
- [28] Kathryn M. Zurek. Snowmass 2021 White Paper: Observational Signatures of Quantum Gravity, May 2022. URL <http://arxiv.org/abs/2205.01799>. arXiv:2205.01799 [gr-qc, physics:hep-ph, physics:hep-th].
- [29] S. W. Hawking. Breakdown of predictability in gravitational collapse. *Physical Review D*, 14(10):2460–2473, November 1976. doi: 10.1103/PhysRevD.14.2460. URL <https://link.aps.org/doi/10.1103/PhysRevD.14.2460>. Publisher: American Physical Society.
- [30] Charles W. Misner, K. S. Thorne, and J.A. Wheeler. *Gravitation*. Princeton University Press, San Francisco, 1973. ISBN 978-0-7167-0344-0 978-0-691-17779-3.
- [31] Ina Fourie. Entropy and Information Theory (2nd ed.). *Online Information Review*, 36(3):481–482, January 2012. ISSN 1468-4527. doi: 10.1108/14684521211241477. URL <https://doi.org/10.1108/14684521211241477>. Publisher: Emerald Group Publishing Limited.
- [32] Shinsei Ryu and Tadashi Takayanagi. Aspects of holographic entanglement entropy. *Journal of High Energy Physics*, 2006(08):045, August 2006. ISSN 1126-6708. doi: 10.1088/1126-6708/2006/08/045. URL <https://dx.doi.org/10.1088/1126-6708/2006/08/045>.
- [33] Mark Srednicki. Entropy and area. *Physical Review Letters*, 71(5):666–669, August 1993. doi: 10.1103/PhysRevLett.71.666. URL <https://link.aps.org/doi/10.1103/PhysRevLett.71.666>. Publisher: American Physical Society.
- [34] Horacio Casini, Marina Huerta, and Robert C. Myers. Towards a derivation of holographic entanglement entropy. *Journal of High Energy Physics*, 2011(5):36, May 2011. ISSN 1029-8479. doi: 10.1007/JHEP05(2011)036. URL [http://link.springer.com/10.1007/JHEP05\(2011\)036](http://link.springer.com/10.1007/JHEP05(2011)036).
- [35] Tom Banks, Patrick Draper, and Szilard Farkas. Path integrals for causal diamonds and the covariant entropy principle. *Physical Review D*, 103(10):106022, May 2021. doi: 10.1103/PhysRevD.103.106022. URL <https://link.aps.org/doi/10.1103/PhysRevD.103.106022>. Publisher: American Physical Society.
- [36] Joshua H. Cooperman and Markus A. Luty. Renormalization of entanglement entropy and the gravitational effective action. *Journal of High Energy Physics*,
-

- 2014(12):45, December 2014. ISSN 1029-8479. doi: 10.1007/JHEP12(2014)045. URL [https://doi.org/10.1007/JHEP12\(2014\)045](https://doi.org/10.1007/JHEP12(2014)045).
- [37] Curtis Callan and Frank Wilczek. On geometric entropy. *Physics Letters B*, 333(1):55–61, July 1994. ISSN 0370-2693. doi: 10.1016/0370-2693(94)91007-3. URL <https://www.sciencedirect.com/science/article/pii/0370269394910073>.
- [38] Leonard Susskind and John Uglum. Black hole entropy in canonical quantum gravity and superstring theory. *Physical Review. D, Particles and Fields*, 50(4):2700–2711, August 1994. ISSN 0556-2821. doi: 10.1103/physrevd.50.2700.
- [39] Karl Pearson. The Problem of the Random Walk. *Nature*, 72(1865):294–294, July 1905. ISSN 1476-4687. doi: 10.1038/072294b0. URL <https://www.nature.com/articles/072294b0>. Number: 1865 Publisher: Nature Publishing Group.
- [40] Peter Fritschel. Second generation instruments for the Laser Interferometer Gravitational Wave Observatory (LIGO). page 282, February 2003. doi: 10.1117/12.459090. URL <http://arxiv.org/abs/gr-qc/0308090>. arXiv:gr-qc/0308090.
- [41] Albert Einstein. On the Electrodynamics of Moving Bodies. *Annalen der Physik*, June 1905. URL <http://archive.org/details/einstein-1905-relativity>.
- [42] H. Minkowski. Die Grundgleichungen für die elektromagnetischen Vorgänge in bewegten Körpern. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1908:53–111, 1908. URL <https://eudml.org/doc/58707>.
- [43] John Archibald Wheeler and Kenneth Ford. Geons, Black Holes and Quantum Foam: A Life in Physics. *American Journal of Physics*, 68, June 2000. doi: 10.1119/1.19497. URL https://www.researchgate.net/publication/253780184_Geons_Black_Holes_and_Quantum_Foam_A_Life_in_Physics.
- [44] Lorenzo Aiello. *Development of new approaches for optical aberration control in gravitational wave interferometers*. PhD thesis, Gran Sasso Science Institute, April 2019. URL <https://iris.gssi.it/handle/20.500.12571/9702>.
- [45] Scott A Hughes. Black Hole Physics via Gravitational Waves, August 2017. URL <https://bhi.fas.harvard.edu/files/bhi/files/hughes.pdf>.
- [46] Eanna E. Flanagan and Scott A. Hughes. Measuring gravitational waves from binary black hole coalescences: I. Signal to noise for inspiral, merger, and ringdown. *Physical Review D*, 57(8):4535–4565, April 1998. ISSN 0556-2821,

-
- 1089-4918. doi: 10.1103/PhysRevD.57.4535. URL <http://arxiv.org/abs/gr-qc/9701039>. arXiv:gr-qc/9701039.
- [47] Scott A. Hughes. Gravitational waves from merging compact binaries. *Annual Review of Astronomy and Astrophysics*, 47(1):107–157, September 2009. ISSN 0066-4146, 1545-4282. doi: 10.1146/annurev-astro-082708-101711. URL <http://arxiv.org/abs/0903.4877>. arXiv:0903.4877 [astro-ph, physics:gr-qc].
- [48] S. M. Vermeulen, L. Aiello, A. Ejlli, and W. L. Griffiths et al. An experiment for observing quantum gravity phenomena using twin table-top 3D interferometers. *Classical and Quantum Gravity*, 38(8):085008, March 2021. ISSN 0264-9381. doi: 10.1088/1361-6382/abe757. URL <https://iopscience.iop.org/article/10.1088/1361-6382/abe757/meta>. Publisher: IOP Publishing.
- [49] Aaron S. Chou et al. MHz Gravitational Wave Constraints with Decameter Michelson Interferometers. *Physical Review D*, 95(6):063002, March 2017. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.95.063002. URL <http://arxiv.org/abs/1611.05560>.
- [50] B. J. Carr and S. W. Hawking. Black Holes in the Early Universe. *Monthly Notices of the Royal Astronomical Society*, 168(2):399–415, August 1974. ISSN 0035-8711. doi: 10.1093/mnras/168.2.399. URL <https://doi.org/10.1093/mnras/168.2.399>.
- [51] F. Zwicky. Republication of: The redshift of extragalactic nebulae. *General Relativity and Gravitation*, 41:207–224, January 2009. ISSN 0001-7701. doi: 10.1007/s10714-008-0707-4. URL <https://ui.adsabs.harvard.edu/abs/2009GRGr..41..207Z>. ADS Bibcode: 2009GRGr..41..207Z.
- [52] J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bulletin of the Astronomical Institutes of the Netherlands*, 6:249, August 1932. ISSN 0365-8910. URL <https://ui.adsabs.harvard.edu/abs/1932BAN.....6..249O>. ADS Bibcode: 1932BAN.....6..249O.
- [53] S. M. Faber and J. S. Gallagher. Masses and Mass-To-Light Ratios of Galaxies. *Annual Review of Astronomy and Astrophysics*, 17(1):135–187, 1979. doi: 10.1146/annurev.aa.17.090179.001031. URL <https://doi.org/10.1146/annurev.aa.17.090179.001031>. _eprint: <https://doi.org/10.1146/annurev.aa.17.090179.001031>.
- [54] Edvige Corbelli and Paolo Salucci. The extended rotation curve and the dark matter halo of M33. *Monthly Notices of the Royal Astronomical Society*, 311(2):441–447, January 2000. ISSN 0035-8711. doi: 10.1046/j.1365-8711.2000.03075.x. URL <https://doi.org/10.1046/j.1365-8711.2000.03075.x>.
-

- [55] James Binney and Michael Merrifield. *Galactic Astronomy*. Princeton University Press, September 1998. ISBN 978-0-691-02565-0. URL <https://press.princeton.edu/books/paperback/9780691025650/galactic-astronomy>.
- [56] Steven W. Allen, August E. Evrard, and Adam B. Mantz. Cosmological Parameters from Observations of Galaxy Clusters. *Annual Review of Astronomy and Astrophysics*, 49(1):409–470, September 2011. ISSN 0066-4146, 1545-4282. doi: 10.1146/annurev-astro-081710-102514. URL <http://arxiv.org/abs/1103.4829>. arXiv:1103.4829 [astro-ph].
- [57] Priyamvada Natarajan et al. Mapping substructure in the HST Frontier Fields cluster lenses and in cosmological simulations. *Monthly Notices of the Royal Astronomical Society*, 468(2):1962–1980, June 2017. ISSN 0035-8711, 1365-2966. doi: 10.1093/mnras/stw3385. URL <http://arxiv.org/abs/1702.04348>. arXiv:1702.04348 [astro-ph].
- [58] Planck Collaboration. Planck 2015 results. XIII. Cosmological parameters. *Astronomy & Astrophysics*, 594:A13, October 2016. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361/201525830. URL <http://arxiv.org/abs/1502.01589>. arXiv:1502.01589 [astro-ph].
- [59] Lorenzo Aiello et al. Constraints on Scalar Field Dark Matter from Colocated Michelson Interferometers. *Physical Review Letters*, 128(12):121101, March 2022. doi: 10.1103/PhysRevLett.128.121101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.128.121101>. Publisher: American Physical Society.
- [60] Sander M. Vermeulen and Philip Relton et al. Direct limits for scalar field dark matter from a gravitational-wave detector. *Nature*, 600(7889):424–428, December 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-04031-y. URL <https://www.nature.com/articles/s41586-021-04031-y>. Number: 7889. Publisher: Nature Publishing Group.
- [61] James Lough et al. First Demonstration of 6 dB Quantum Noise Reduction in a Kilometer Scale Gravitational Wave Observatory. *Physical Review Letters*, 126(4):041102, January 2021. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.126.041102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.126.041102>.
- [62] Aaron Chou et al. The Holometer: An Instrument to Probe Planckian Quantum Geometry. *Classical and Quantum Gravity*, 34(6):065005, March 2017. ISSN 0264-9381, 1361-6382. doi: 10.1088/1361-6382/aa5e5c. URL <http://arxiv.org/abs/1611.08265>.

-
- [63] J. M. Vaughan. *The Fabry–Perot Interferometer: History, Theory, Practice and Applications*. Routledge, New York, October 2017. ISBN 978-0-203-73671-5. doi: 10.1201/9780203736715.
- [64] J. B. Hawkes and R. W. Astheimer. An Improved Jamin Interferometer. *JOSA*, 38(7):617–622, July 1948. doi: 10.1364/JOSA.38.000617. URL <https://opg.optica.org/josa/abstract.cfm?uri=josa-38-7-617>. Publisher: Optica Publishing Group.
- [65] Daniel Malacara. *Optical Shop Testing*. John Wiley & Sons, July 2007. ISBN 978-0-471-48404-2. Google-Books-ID: 8sJcDwAAQBAJ.
- [66] Hamed Saghaei, Payam Elyasi, and Rouhollah Karimzadeh. Design, fabrication, and characterization of Mach–Zehnder interferometers. *Photonics and Nanostructures - Fundamentals and Applications*, 37:100733, December 2019. ISSN 1569-4410. doi: 10.1016/j.photonics.2019.100733. URL <https://www.sciencedirect.com/science/article/pii/S156944101930149X>.
- [67] Charlotte Bond, Daniel Brown, Andreas Freise, and Kenneth A. Strain. Interferometer techniques for gravitational-wave detection. *Living Reviews in Relativity*, 19(1):3, December 2016. ISSN 2367-3613, 1433-8351. doi: 10.1007/s41114-016-0002-8. URL <http://link.springer.com/10.1007/s41114-016-0002-8>.
- [68] Walid A. Zgallai et al. *Biomedical Signal Processing and Artificial Intelligence in Healthcare - 1st Edition*, January 2023. URL <https://www.elsevier.com/books/T/A/9780128189467>.
- [69] Fabrice Matichard et al. LIGO-P1200040-v55: Seismic Isolation of Advanced LIGO: Review of Strategy, Instrumentation, and Performance (CQG 2015), March 2019. URL <https://dcc.ligo.org/LIGO-P1200040>.
- [70] Avraham Gover, Ariel Nause, Egor Dyunin, and Mikhail Fedurin. Beating the shot-noise limit. *Nature Physics*, 8(12):877–880, December 2012. ISSN 1745-2481. doi: 10.1038/nphys2443. URL <https://www.nature.com/articles/nphys2443>. Number: 12 Publisher: Nature Publishing Group.
- [71] David Reitze, Peter Saulson, and Hartmut Grote. *Advanced Interferometric Gravitational-Wave Detectors: (In 2 Volumes) Volume I: Essentials of Gravitational-Wave Detectors Volume II: Advanced LIGO, Advanced Virgo and Beyond*, volume 05 of *100 Years of General Relativity*. World Scientific, May 2019. ISBN 978-981-314-607-5 978-981-314-608-2. doi: 10.1142/10181. URL <https://www.worldscientific.com/worldscibooks/10.1142/10181>.
-

- [72] Anna Catriona Green. *When light gets pushy: radiation pressure effects in interferometric gravitational wave detectors*. d_ph, University of Birmingham, December 2018. URL <https://etheses.bham.ac.uk/id/eprint/8512/>.
- [73] Carlton M. Caves. Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer. *Physical Review Letters*, 45(2):75–79, July 1980. doi: 10.1103/PhysRevLett.45.75. URL <https://link.aps.org/doi/10.1103/PhysRevLett.45.75>. Publisher: American Physical Society.
- [74] Craig Cahillane and Georgia Mansell. Review of the Advanced LIGO gravitational wave observatories leading to observing run four. *Galaxies*, 10(1): 36, February 2022. ISSN 2075-4434. doi: 10.3390/galaxies10010036. URL <http://arxiv.org/abs/2202.00847>. arXiv:2202.00847 [astro-ph, physics:gr-qc, physics:physics].
- [75] The Virgo Collaboration. Quantum Backaction on kg-Scale Mirrors: Observation of Radiation Pressure Noise in the Advanced Virgo Detector. *Physical Review Letters*, 125(13):131101, September 2020. doi: 10.1103/PhysRevLett.125.131101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.125.131101>. Publisher: American Physical Society.
- [76] Peter R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing Company, 2 edition, April 2017. ISBN 978-981-314-307-4. URL <https://www.worldscientific.com/worldscibooks/10.1142/10116#t=aboutBook>.
- [77] Eugene Hecht. *Optics*. Pearson Education, Inc, Boston, 5 ed edition, 2017. ISBN 978-0-13-397722-6.
- [78] A. E. Siegman. *Lasers*. University Science Books, 1986. ISBN 978-0-935702-11-8. Google-Books-ID: 1BZVwUZLTkAC.
- [79] Alex Abramovici and Jake Chapsky. *Feedback Control Systems: A Fast-Track Guide for Scientists and Engineers*. Springer Science & Business Media, September 2000. ISBN 978-0-7923-7935-5.
- [80] Eugene Hecht. *Optics*. Addison Wesley, Reading, Mass, 2nd edition edition, January 1987. ISBN 978-0-201-11609-0.
- [81] OPI. Hermite-Gaussian Modes, December 2010. URL http://www.optique-ingenieur.org/en/courses/OPI_ang_M01_C03/co/Contenu_13.html.
- [82] OPI. High order modes, December 2010. URL http://www.optique-ingenieur.org/en/courses/OPI_ang_M01_C03/co/Grain_OPI_ang_M01_C03_5.html.

-
- [83] Miguel A. Bandres and Julio C. Gutiérrez-Vega. Ince–Gaussian beams. *Optics Letters*, 29(2):144–146, January 2004. ISSN 1539-4794. doi: 10.1364/OL.29.000144. URL <https://opg.optica.org/ol/abstract.cfm?uri=ol-29-2-144>. Publisher: Optica Publishing Group.
- [84] Ebrahim Karimi et al. Hypergeometric-Gaussian modes. *Optics Letters*, 32(21):3053–3055, November 2007. ISSN 1539-4794. doi: 10.1364/OL.32.003053. URL <https://opg.optica.org/ol/abstract.cfm?uri=ol-32-21-3053>. Publisher: Optica Publishing Group.
- [85] Simin Feng and Herbert G. Winful. Physical origin of the Gouy phase shift. *Optics Letters*, 26(8):485, April 2001. ISSN 0146-9592, 1539-4794. doi: 10.1364/OL.26.000485. URL <https://opg.optica.org/abstract.cfm?URI=ol-26-8-485>.
- [86] Hartmut Grote. *Advanced Interferometric Gravitational-wave Detectors (In 2 Volumes)*. World Scientific, March 2019. ISBN 978-981-314-609-9.
- [87] The LIGO Scientific Collaboration. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, April 2015. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/32/7/074001. URL <http://arxiv.org/abs/1411.4547>. arXiv:1411.4547 [astro-ph, physics:gr-qc, physics:physics].
- [88] P Fritschel, V Frolov, and D Reitze. LIGO-T1100201-v4: Commissioning the Advanced LIGO L1 Input Mode Cleaner, March 2012. URL <https://dcc.ligo.org/LIGO-T1100201/public>.
- [89] LIGO-T1000276-v5: Output Mode Cleaner (OMC) Design, . URL <https://dcc.ligo.org/LIGO-T1000276>.
- [90] P. Barriga et al. Optical design of a high power mode-cleaner for AIGO. *General Relativity and Gravitation*, 37(9):1609–1619, September 2005. ISSN 0001-7701, 1572-9532. doi: 10.1007/s10714-005-0146-4. URL <http://link.springer.com/10.1007/s10714-005-0146-4>.
- [91] N. Uehara and K. Ueda. Accurate measurement of ultralow loss in a high-finesse Fabry-Perot interferometer using the frequency response functions. *Applied Physics B*, 61(1):9–15, July 1995. ISSN 1432-0649. doi: 10.1007/BF01090966. URL <https://doi.org/10.1007/BF01090966>.
- [92] Dana Z. Anderson. Alignment of resonant optical cavities. *Applied Optics*, 23(17):2944–2949, September 1984. ISSN 2155-3165. doi: 10.1364/AO.23.002944. URL <https://opg.optica.org/ao/abstract.cfm?uri=ao-23-17-2944>. Publisher: Optica Publishing Group.
-

- [93] J. Magnes et al. Quantitative and Qualitative Study of Gaussian Beam Visualization Techniques. *arXiv:physics/0605102*, May 2006. URL <http://arxiv.org/abs/physics/0605102>. arXiv: physics/0605102.
- [94] Hitay Ozbay. *Introduction to Feedback Control Theory*. CRC Press, Boca Raton, 1st edition edition, July 1999. ISBN 978-0-8493-1867-2.
- [95] Moku:Lab, January 2023. URL <https://www.liquidinstruments.com/products/hardware-platforms/mokulab/>.
- [96] Mephisto/Mephisto S. page 4, . URL <https://www.coherent.com/content/dam/coherent/site/en/resources/datasheet/lasers/mephisto-mephisto-s-ds.pdf>.
- [97] SRS SR560 Low-Noise Voltage Preamplifier, . URL <https://www.lambdaphoto.co.uk/sr560-low-noise-voltage-preamplifier.html>.
- [98] Thorlabs . PDA20CS2 InGaAs Switchable Gain Amplified Detecto, June 2023. URL <https://www.thorlabs.com>.
- [99] S Hild et al. DC-readout of a signal-recycled gravitational wave detector. *Classical and Quantum Gravity*, 26(5):055012, March 2009. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/26/5/055012. URL <https://iopscience.iop.org/article/10.1088/0264-9381/26/5/055012>.
- [100] Steven W. Smith. *The Scientist & Engineer's Guide to Digital Signal Processing*. California Technical Pub, San Diego, Calif, March 1998. ISBN 978-0-9660176-3-2.
- [101] Peter Fritschel, Matthew Evans, and Valery Frolov. Balanced homodyne readout for quantum limited gravitational wave detectors. *Optics express*, 22:4224–34, February 2014. doi: 10.1364/OE.22.004224.
- [102] E. D. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, January 2001. ISSN 0002-9505, 1943-2909. doi: 10.1119/1.1286663. URL <http://aapt.scitation.org/doi/10.1119/1.1286663>.
- [103] Adam Joseph Mullavey. *Arm length stabilisation for advanced gravitational wave detectors*. PhD thesis, 2012. URL <https://openresearch-repository.anu.edu.au/handle/1885/155176>.
- [104] Beam Blocks and Traps, . URL https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=1449.
- [105] David Brewster. On the Laws Which Regulate the Polarisation of Light by Reflexion from Transparent Bodies. *Philosophical Transactions of the Royal*

-
- Society of London*, 105:125–159, 1815. ISSN 0261-0523. URL <http://www.jstor.org/stable/107362>. Publisher: The Royal Society.
- [106] Ramón Paniagua-Domínguez et al. Generalized Brewster effect in dielectric metasurfaces. *Nature Communications*, 7:10362, January 2016. ISSN 2041-1723. doi: 10.1038/ncomms10362. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4735648/>.
- [107] LIGO-T1300867-v1: Control loop design for the aLIGO second loop power stabilization, . URL <https://dcc.ligo.org/LIGO-T1300867>.
- [108] LIGO-G1601674-v1: Laser Power Stabilization for the AEI 10m Prototype, . URL <https://dcc.ligo.org/LIGO-G1601674>.
- [109] LIGO-P1400245-v1: Stabilized High Power Lasers and Spatial Mode Conversion, . URL <https://dcc.ligo.org/LIGO-P1400245>.
- [110] LIGO-P1100192-v1: Stabilized high-power laser system for the gravitational wave detector Advanced LIGO, . URL <https://dcc.ligo.org/LIGO-P1100192>.
- [111] LIGO-G1000320-v1: Power stabilization of the Advanced LIGO laser, . URL <https://dcc.ligo.org/LIGO-G1000320>.
- [112] BG39 | SCHOTT Advanced Optics, . URL <https://www.schott.com/shop/advanced-optics/en/Matt-Filter-Plates/BG39/c/glass-BG39>.
- [113] Catalogs and datasheets for Optical Filter Glass | SCHOTT, . URL <https://www.schott.com/en-gb/products/optical-filter-glass-p1000266/downloads>.
- [114] Carlton M. Caves. Quantum-mechanical noise in an interferometer. *Physical Review D*, 23(8):1693–1708, April 1981. ISSN 0556-2821. doi: 10.1103/PhysRevD.23.1693. URL <https://link.aps.org/doi/10.1103/PhysRevD.23.1693>.
- [115] Henning Vahlbruch et al. The GEO600 squeezed light source. *Classical and Quantum Gravity*, 27(8):084027, April 2010. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/27/8/084027. URL <http://arxiv.org/abs/1004.4975>. arXiv:1004.4975 [quant-ph].
- [116] J. Aasi et al. Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photonics*, 7(8):613–619, August 2013. ISSN 1749-4893. doi: 10.1038/nphoton.2013.177. URL <https://www.nature.com/articles/nphoton.2013.177>. Number: 8 Publisher: Nature Publishing Group.
-

- [117] Moritz Mehmet and Henning Vahlbruch. The Squeezed Light Source for the Advanced Virgo Detector in the Observation Run O3. *Galaxies*, 8(4): 79, December 2020. ISSN 2075-4434. doi: 10.3390/galaxies8040079. URL <https://www.mdpi.com/2075-4434/8/4/79>. Number: 4 Publisher: Multi-disciplinary Digital Publishing Institute.
- [118] Gerard t Hooft. The Quantum Black Hole as a Hydrogen Atom: Microstates Without Strings Attached. *arXiv:1605.05119 [gr-qc, physics:hep-th, physics:quant-ph]*, May 2016. URL <http://arxiv.org/abs/1605.05119>.
- [119] Craig Hogan. Nonlocal entanglement and directional correlations of primordial perturbations on the inflationary horizon. *Physical Review D*, 99(6):063531, March 2019. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.99.063531. URL <https://link.aps.org/doi/10.1103/PhysRevD.99.063531>.
- [120] Muzammil A. Arain and Guido Mueller. Design of the Advanced LIGO recycling cavities. *Optics Express*, 16(14):10018, July 2008. ISSN 1094-4087. doi: 10.1364/OE.16.010018. URL <https://opg.optica.org/oe/abstract.cfm?uri=oe-16-14-10018>.
- [121] Kiwamu Izumi and Daniel Sigg. Advanced LIGO: length sensing and control in a dual recycled interferometric gravitational wave antenna. *Classical and Quantum Gravity*, 34(1):015001, January 2017. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/34/1/015001. URL <https://iopscience.iop.org/article/10.1088/0264-9381/34/1/015001>.
- [122] LIGO-P1400105-v5: Achieving Resonance in the Advanced LIGO Gravitational-Wave Interferometer, . URL <https://dcc.ligo.org/LIGO-P1400105/public>.
- [123] LIGO-T070303-x0: Arm Cavity Finesse for Advanced LIGO, . URL <https://dcc.ligo.org/LIGO-T070303/public>.
- [124] Ohkyung Kwon and Craig J. Hogan. Interferometric Tests of Planckian Quantum Geometry Models. *Classical and Quantum Gravity*, 33(10):105004, May 2016. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/33/10/105004. URL <http://arxiv.org/abs/1410.8197>. arXiv:1410.8197 [gr-qc, physics:hep-th, physics:quant-ph].
- [125] Dongjun Li et al. Interferometer Response to Geotropic Fluctuations, September 2022. URL <http://arxiv.org/abs/2209.07543>. arXiv:2209.07543 [gr-qc, physics:hep-ph, physics:hep-th].
- [126] LIGO-T1800042-v5: The A+ design curve, . URL <https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?.submit=Number&docid=T1800042&version=>.

-
- [127] A. Chou et al. Interferometric Constraints on Quantum Geometrical Shear Noise Correlations. *Classical and Quantum Gravity*, 34(16):165005, August 2017. ISSN 0264-9381, 1361-6382. doi: 10.1088/1361-6382/aa7bd3. URL <http://arxiv.org/abs/1703.08503>.
- [128] Tomotada Akutsu, et al. Search for a Stochastic Background of 100-MHz Gravitational Waves with Laser Interferometers. *Physical Review Letters*, 101(10):101101, September 2008. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.101.101101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.101.101101>.
- [129] Atsushi Nishizawa. Laser-interferometric detectors for gravitational wave backgrounds at 100 MHz: Detector design and sensitivity. *Physical Review D*, 77(2):022002, January 2008. ISSN 1550-7998, 1550-2368. doi: 10.1103/PhysRevD.77.022002. URL <https://link.aps.org/doi/10.1103/PhysRevD.77.022002>.
- [130] Ph Bernard, G. Gemme, R. Parodi, and E. Picasso. A detector of small harmonic displacements based on two coupled microwave cavities. *Review of Scientific Instruments*, 72(5):2428–2437, May 2001. ISSN 0034-6748, 1089-7623. doi: 10.1063/1.1366636. URL <http://arxiv.org/abs/gr-qc/0103006>. arXiv:gr-qc/0103006.
- [131] A M Cruise and R M J Ingley. A prototype gravitational wave detector for 100 MHz. *Classical and Quantum Gravity*, 23(22):6185–6193, November 2006. ISSN 0264-9381, 1361-6382. doi: 10.1088/0264-9381/23/22/007. URL <https://iopscience.iop.org/article/10.1088/0264-9381/23/22/007>.
- [132] Gaussian Beam Optics, . URL <http://experimentationlab.berkeley.edu/sites/default/files/MOT/Gaussian-Beam-Optics.pdf>.
- [133] BP209-VIS/M Dual Scanning Slit Beam Profiler, 200 - 1100 nm, Ø2.5 µm - Ø9 mm, Metric, . URL <https://www.thorlabs.com/thorproduct.cfm?partnumber=BP209-VIS/M>.
- [134] Nur Ismail, Cristine Calil Kores, Dimitri Geskus, and Markus Pollnau. Fabry-Pérot resonator: spectral line shapes, generic and related Airy distributions, linewidths, finesses, and performance at low or frequency-dependent reflectivity. *Optics Express*, 24(15):16366–16389, July 2016. ISSN 1094-4087. doi: 10.1364/OE.24.016366. URL <https://opg.optica.org/oe/abstract.cfm?uri=oe-24-15-16366>. Publisher: Optica Publishing Group.
- [135] Laser Lock Box (Moku:Lab), January 2023. URL <https://www.liquidinstruments.com/products/integrated-instruments/laser-lock-box-mokulab/>.
-

- [136] Digital Filter Box (Moku:Lab), . URL <https://www.liquidinstruments.com/products/integrated-instruments/digital-filter-box-mokulab/>.
- [137] neolase. Laser Amplifiers, June 2020. URL <https://neolase.com/produkte/laser-verstaerker/?lang=en>.
- [138] A. Thüring and N. Lastzka. JamMT. URL [http://www.sr.bham.ac.uk/dokuwiki/doku.php?id=geosim:jammt&s\[\]=jammt](http://www.sr.bham.ac.uk/dokuwiki/doku.php?id=geosim:jammt&s[]=jammt).
- [139] 8823-UHV Piezo Mirror Mount, . URL <https://www.newport.com/p/8823-UHV>.
- [140] S-316 Piezo Z and Tip/Tilt Scanner, January 2023. URL <https://www.pi-usa.us/en/products/fast-steering-mirrors-fsm-and-piezo-tiptilt-platforms-for-active-optics/s-310-s-316-piezo-z-and-tiptilt-scanner-300600>.
- [141] S425C-L Thermal Power Sensor Head, Surface Absorber, 0.19 - 20 μm , 50 W, $\varnothing 25.4$ mm, . URL <https://www.thorlabs.com/thorproduct.cfm?partnumber=S425C-L>.
- [142] S132C Slim Photodiode Power Sensor, Ge, 700 - 1800 nm, 500 mW, February 2023. URL <https://www.thorlabs.com/thorproduct.cfm?partnumber=S132C>.
- [143] 1801-FS Optical Receiver, . URL <https://www.newport.com/p/1801-FS>.
- [144] S. T. Pradyumna et al. Twin beam quantum-enhanced correlated interferometry for testing fundamental physics. *Communications Physics*, 3(1):104, December 2020. ISSN 2399-3650. doi: 10.1038/s42005-020-0368-5. URL <http://www.nature.com/articles/s42005-020-0368-5>.
- [145] I. Ruo Berchera et al. One- and two-mode squeezed vacuum and squeezed light in correlated interferometry. *Physical Review A*, 92(5):053821, November 2015. ISSN 1050-2947, 1094-1622. doi: 10.1103/PhysRevA.92.053821. URL <http://arxiv.org/abs/1501.07516>.
- [146] I. Ruo Berchera, I. P. Degiovanni, S. Olivares, and M. Genovese. Quantum Light in Coupled Interferometers for Quantum Gravity Tests. *Physical Review Letters*, 110(21):213601, May 2013. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.110.213601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.110.213601>.
- [147] Simon Chelkowski. Squeezed Light and Laser Interferometric Gravitational Wave Detectors. page 274.

-
- [148] Sheila E Dwyer. Quantum noise reduction using squeezed states in LIGO. page 223.
- [149] Mageswaran, Mohana. LIGO-D1100290-v1: aLigo Optical Lever QPD Amplifier. URL <https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?.submit=Identifier&docid=D1100290&version=1>.
- [150] F. Acernese et al. Status of VIRGO. volume 5500. Spie Optical Engineering Press, June 2004. URL <https://hal.in2p3.fr/in2p3-00025835>. Issue: 5500.
- [151] M. Prijatelj et al. The output mode cleaner of GEO 600. *Classical and Quantum Gravity*, 29(5):055009, February 2012. ISSN 0264-9381. doi: 10.1088/0264-9381/29/5/055009. URL <https://iopscience.iop.org/article/10.1088/0264-9381/29/5/055009/meta>. Publisher: IOP Publishing.
- [152] Junko Kasuya, John Winterflood, Ju Li, and Kentaro Somiya. Optical design and suspension system of the KAGRA output mode-cleaner. *Journal of Physics: Conference Series*, 957(1):012009, February 2018. ISSN 1742-6596. doi: 10.1088/1742-6596/957/1/012009. URL <https://dx.doi.org/10.1088/1742-6596/957/1/012009>. Publisher: IOP Publishing.
- [153] C Mathis et al. Resonances and instabilities in a bidirectional ring laser. *Physica D: Nonlinear Phenomena*, 96(1-4):242–250, September 1996. ISSN 01672789. doi: 10.1016/0167-2789(96)00024-3. URL <https://linkinghub.elsevier.com/retrieve/pii/0167278996000243>.
- [154] Edmund Optics. Understanding Surface Quality Specifications | Edmund Optics, January 2023. URL <https://www.edmundoptics.co.uk/knowledge-center/application-notes/lasers/understanding-surface-quality-specifications/>.
- [155] Laseroptik GmbH, 2023. URL <https://www.laseroptik.com/>.
- [156] Alena Ananyeva and Calum Torrie. LIGO-T1700128-v9: Infrared Reflectance & Scatter Tests for Potential Scattered Light Mitigation Coatings (R² vs. BRDF), January 2023. URL <https://dcc.ligo.org/LIGO-T1700128>.
- [157] LIGO-G1700379-v1: Infrared and Ultraviolet Reflectance & Scatter Tests for Potential Scattered Light Mitigation Coatings, March 2017. URL <https://dcc.ligo.org/LIGO-G1700379>.
- [158] Ultra-Low Loss SuperMirrors™ - MICRO-CONTROLE / Spectra-Physics - PDF Catalogs | Technical Documentation | Brochure, . URL <https://pdf.directindustry.com/pdf/micro-controle-spectra-physics/ultra-low-loss-supermirrors/7436-138515.html>.
-

- [159] P Fritschel. Output Mode Cleaner Design. *LIGO Laboratory / LIGO Scientific Collaboration*, page 15, February 2004. URL <https://labcit.ligo.caltech.edu/~pking/omc/docs/notes/omc-design.pdf>.
- [160] WebPlotDigitizer - Copyright 2010-2022 Ankit Rohatgi, . URL <https://apps.automeris.io/wpd/>.
- [161] K. P. Birch and M. J. Downs. Correction to the Updated Edlén Equation for the Refractive Index of Air. *Metrologia*, 31(4):315–316, January 1994. ISSN 0026-1394. doi: 10.1088/0026-1394/31/4/006. URL <https://doi.org/10.1088/0026-1394/31/4/006>. Publisher: IOP Publishing.
- [162] Air Refractive Index, . URL <https://www.mathworks.com/matlabcentral/fileexchange/31240-air-refractive-index>.
- [163] dB Conversion of sound pressure to sound intensity, August 2022. URL <http://www.sengpielaudio.com/calculator-soundlevel.htm>.
- [164] PA44M3KW Piezo Ring Chip, February 2023. URL <https://www.thorlabs.com/thorproduct.cfm?partnumber=PA44M3KW>.
- [165] First Contact Polymer Kits, . URL <https://www.photoniccleaning.com/Kits-s/112.htm>.
- [166] James M. Fiore. *AC Electrical Circuit Analysis: A Practical Approach*. Independently published, April 2020. ISBN 9798605022282.
- [167] Eric W. Weisstein. Lorentzian Function. URL <https://mathworld.wolfram.com/>. Publisher: Wolfram Research, Inc.
- [168] Alexa Nitzan Staley. *Locking the Advanced LIGO Gravitational Wave Detector: with a focus on the Arm Length Stabilization Technique*. PhD thesis, Columbia University, 2015. URL <https://doi.org/10.7916/D8X34WQ4>.
- [169] Fundamentals of Phase Locked Loops (PLLs), 2016. URL <https://www.analog.com/media/en/training-seminars/tutorials/MT-086.pdf>.
- [170] Mueller Guido et al. LIGO-T1000298-v2: Advanced LIGO Length Sensing and Control Final Design, February 2010. URL <https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?.submit=Identifier&docid=T1000298&version=>.
- [171] Denis V. Martynov. *Lock Acquisition and Sensitivity Analysis of Advanced LIGO Interferometers*. phd, California Institute of Technology, 2015. URL <https://resolver.caltech.edu/CaltechTHESIS:05282015-142013480>.

-
- [172] Adam J. Mullavey et al. Stable transfer of an optical frequency standard via a 4.6 km optical fiber. *Optics Express*, 18(5):5213–5220, March 2010. ISSN 1094-4087. doi: 10.1364/OE.18.005213. URL <https://opg.optica.org/oe/abstract.cfm?uri=oe-18-5-5213>. Publisher: Optica Publishing Group.
- [173] J. H. Poeld. LIGO-T1400015-v1: Reference cavity temperature control, January 2014. URL <https://dcc.ligo.org/LIGO-T1400015>.
- [174] J. H. Poeld. LIGO-D1300949-v2: aLIGO PSL reference cavity heater/fieldbox, January 2014. URL <https://dcc.ligo.org/D1300949>.
- [175] Peter Fritschel. LIGO-T1300090-v2: IO Mode Cleaner, September 2004. URL <https://dcc.ligo.org/LIGO-T1300090>.
- [176] Bram Slagmolen and John Miller. LIGO-T1000555-v4: Adv. LIGO Arm Length Stabilisation - Vertex Layout Overview, May 2015. URL <https://dcc.ligo.org/LIGO-T1000555>.
- [177] Yutaro Enomoto. *JGW-P1911103-v12: [Thesis] Interferometer Locking Scheme for Advanced Gravitational-Wave Detectors and Beyond*. PhD thesis, The University of Tokyo, December 2019. URL <https://gwdoc.icrr.u-tokyo.ac.jp/cgi-bin/DocDB/ShowDocument?docid=11103>.
- [178] Ryosuke Sugimoto and U Toyama. Development of Auxiliary Locking System in Gravitational Wave Telescope KAGRA II, March 2020.
- [179] 1064 nm 2x2 Polarization-Maintaining Fiber Optic Couplers / Taps, January 2023. URL https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=10248.
- [180] Thorlabs. Manual Fiber Polarization Controllers, February 2023. URL <https://www.thorlabs.com>.
- [181] FPC020 Fiber Polarization Controller, 2 Ø18 mm Paddles, No Fiber, June 2023. URL <https://www.thorlabs.com/thorproduct.cfm?partnumber=FPC020>.
- [182] Lumped LC Low Pass Filter, DC - 1.9 MHz, 50Ohm | SLP-1.9+ | Mini-Circuits, January 2023. URL <https://www.minicircuits.com/WebStore/dashboard.html?model=SLP-1.9%2B>.
- [183] Edward Collett. *Field Guide to Polarization*. SPIE, 1000 20th Street, Bellingham, WA 98227-0010 USA, September 2005. ISBN 978-0-8194-5868-1. doi: 10.1117/3.626141. URL <http://link.aip.org/link/doi/10.1117/3.626141>.
-