

School of Physics and Astronomy



Advanced Experimental Systems for Gravitational Wave Detectors and Dark Matter Searches

Alasdair L. James

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

Date: 01/08/2022

Acknowledgments

I consider myself incredibly fortunate to have been helped by so many people over the course of this PhD that it is impossible to name everybody individually without this document becoming several pages longer than it already is. As much as I want to do that anyway (because it sounds like a fun challenge), I will instead give a general thanks to everybody I have ever worked with, been helped by or annoyed over the last four years and give special appreciation to those below...

Firstly to my primary supervisor, Hartmut Grote and my secondary supervisor, Kate Dooley. I still do not fully understand why you decided to offer me a PhD placement but I will forever be grateful that you did. Hartmut, you were an incredible teacher with seemingly infinite patience, no matter how bad my soldering technique or how many circuits I blew up. I am honoured to be your first PhD student and look forward to working with you in the future.

To everyone from the Cardiff Grav Waves group: Jonathon, Ian, Dave, Ed, Eleanor, Charlie you are all amazing people and great friends. Cam, you are insane. Thank you for calming me down in my first year and for convincing me to go to Live Lounge straight after a gym session - three times. Rhys, I will always be intimidated by your physical and mental prowess. I imagine I will continue to rugby tackle you each time I see you until we are both frail old men. You were an amazing housemate and are still a truly great friend. Vassilis and Penny - the Greek god and goddess. Thank you for keeping me sane - both inside the office and out - and always making me laugh, even when I have lost my mind due to stress. I am not sure how much time I lost during my PhD due to talking to one or both of you - most likely a couple years - but it was worth it (probably).

Thank you to the notable Cardiff astrophysicists from the top floor of 53 who would offer me 'tea' every Friday afternoon and make a dependable appearance at the Flute to help de-stress each week.

Special thanks to the Cardiff experimental group: Sander, Will, Abhinav, Terri, Keiko, Lorenzo, Eyal, and Aldo. Who would've thought that this ragtag bunch of misfits would be able to accomplish so much in our basement lab! I consider you all geniuses and cannot wait to see where your work will take each of you. Good luck in detecting quantised space-time!

To Aldo especially, thank you for spending so much time helping me when I first arrived in Cardiff. I'll always be grateful and humbled by the time you took out of your own day to discuss whatever physics problem I had encountered. Whether it is on paper or not, I consider you a critical part of my supervisory team and, more importantly, a good friend.

To my friends and colleagues in the ALPS group in Germany and Florida. Working with you all was an incredible experience and allowed me to have undoubtedly the coolest commute of my entire life - riding a bike through an old particle accelerator each day. To my good friends in Hamburg: Kanioar, Rikhav, Todd, Aaron, Nina, Eldar, Ajda, Chiara and Younes. Thank you for joining me for post work weizens and showing me around the 'cool' schanze area of Hamburg.

To the many people who helped me during my whirlwind 3 week trip out to the LIGO Livingston site, especially Anamaria, Valera and Joe B. Without each of you there was no chance that this thesis would have the results it does. Also thank you to Will for being my comrade during our first New Orleans Mardi Gras experience.

To the people in my life that were sane enough not to do a PhD in physics. My family for always remaining my fiercest cheerleaders and allowing me to ramble on about whatever topic I had just learned about (even if it has exactly zero relevance to your own lives). To Brooke, for being the only person who can get as competitive as me over a game of shuffleboard (a game neither of us actually know how to play) and getting as excited as me about taking a trip to Ikea on a Sunday. To Robbie for so many years of incredible friendship, excellent advice and countless stupid stories (I look forward to eating daffodils and punching trees with you for many more years). And to the many many others, I appreciate you all and thank you for always being there to support me.

Finally I must give the obligatory - but no less important - thanks to the funding committees who have generously supported me. Thank you to the Optoelectronics Rank Prize for awarding me the funds for my trip to the US to work at the LIGO Livingston site. Thank you to Cardiff University and the Science and Technology Facilities Council (STFC) for funding my research, my 3 months in Germany and giving me enough of an income to keep a roof over my head whilst I fiddled with lasers and played with electronics. There are few people in the world that are fortunate enough to say that they are paid to do something they love and I am incredibly grateful and humbled to have counted myself amongst them for the past four years.

Summary of thesis

This thesis outlines the work I undertook at Cardiff University as part of my degree for doctor of philosophy. The document is split into two distinct parts.

The first section covers the development of new photodetector electronics that successfully reduced the dark noise of the LIGO experiment by a factor of 6.5 at 2 kHz. This reduction in the dark noise is mainly achieved by filtering the majority of the DC photocurrent from the audio readout path and increasing the transimpedance resistor by a factor of 100. The new readouts are projected to improve the overall DARM sensitivity of LIGO by 10 % at 2 kHz in the next observing run.

The second section covers the implementation of a laser auto-alignment system using differential wavefront sensing for the ALPS II dark matter search. The auto-alignment system is currently undergoing further commissioning but has already shown an improvement in low frequency cavity noise of roughly an order of magnitude which will help the ALPS II detector reach the target design sensitivity in its search for axion-like particles in the upcoming science run at the end of 2022.

Both projects involved the development of advanced techniques for ultra-precise interferometry experiments but the goals and techniques implemented were vastly different. This document outlines the approach in tackling these projects and will hopefully prove a useful guide for others in the future.

Contents

1	Introduction	1
2	LIGO	4
2.1	LIGO Theory	5
2.1.1	Riding the Wave	5
2.1.2	Astronomical Sources	8
2.1.3	From Source to Detection	12
2.1.4	Readout Schemes	18
2.1.5	Noise in the LIGO Experiment	27
2.1.6	The Dark Noise Dilemma	29
2.1.7	Photodetector Model and Other Solutions	33
2.1.8	Old Readout Design	36
2.2	New Readout Design	39
2.2.1	Transfer Function of Audio Output	40
2.2.2	Noise Model of Audio Output	43
2.2.3	Inductor Theory	44
2.2.4	Inductor Design	46
2.2.5	RF Readout Design	48
2.2.6	Bias Readout Design	51
2.2.7	Vacuum Case and Assembly	53
2.2.8	Heat considerations	56
2.3	Measurements and Results	56
2.3.1	Inductor Characterisation	56
2.3.2	Audio Readout	64
2.3.3	Bias Readout	68
2.3.4	RF Readout	69
2.3.5	Op Amp Selection	70
2.3.6	AC Signal Saturation	73
2.3.7	DC Signal Saturation	75
2.3.8	LIGO Results	81
2.3.9	VIRGO Prototype	87
2.4	New LIGO Sensitivity Projection	91
2.5	Future Work	94
2.5.1	Test Relay Improvements	94
2.5.2	DC Signal Injection	95
2.5.3	Inductor Changes	95
2.5.4	Moving the Electronics and Introducing Cryogenics	95
2.6	LIGO Concluding Remarks	97

3	ALPS	98
3.1	ALPS Theory	98
3.1.1	The Standard Model and Dark Matter	98
3.1.2	Axion-like Particles	100
3.2	ALPS Experiment Design	106
3.2.1	ALPS IIc Detection Schemes	107
3.2.2	Optical Table Design	108
3.2.3	Differential Wavefront Sensing Theory	112
3.2.4	Differential Wavefront Sensing Design	115
3.2.5	Signal Processing	120
3.2.6	Mirror Actuators	122
3.3	250 m Cavity Results	124
3.3.1	Two minute Lock	124
3.3.2	Two Hour Lock	129
3.4	Future Work	134
3.4.1	Higher Range Mirror Actuators	134
3.4.2	Linear Combination Matrix	134
3.4.3	Rotation into PZT Matrix	135
3.4.4	Filter Stages	135
3.4.5	Spot Position Control	135
3.4.6	Digital to Analogue Converter and High Voltage Amplifiers	136
3.4.7	General Debugging	136
3.4.8	NL Setup	137
3.5	ALPS Concluding Remarks	138
4	Conclusion	140
5	Appendix	141
5.1	Gaussian Beams	141
5.2	Mode Matching a Fabry-Perot Cavity	142
5.3	DCPD pin layout	145
5.4	PCB Photos	146
5.5	VIRGO-Style Inductor	147
5.6	Variations of the transfer function with DC Current	148
5.7	Full Circuit Schematic	149
5.8	Inductor Mounting	151
5.9	Complex zero/pole Formulation	152
5.10	Auto-alignment Ringing Error Signals	154
5.11	Auto-Alignment Sampling Frequency	154
5.12	Auto-Alignment Two minute time-series	155
5.13	Auto-alignment Two hour control signals	157
5.14	ALPS Photos	157
5.15	LIGO Filtering Stages	159

List of Figures

2.1	Visual representation of the effect of the two polarisations of a gravitational wave	7
2.2	Characteristic strain from a compact binary merger	10
2.3	Basic Michelson interferometer	13
2.4	Michelson interferometer with Fabry-Perot arm cavities	17
2.5	Michelson interferometer with a power recycling cavity	18
2.6	Heterodyne detection scheme for LIGO	19
2.7	DC detection scheme for LIGO	21
2.8	Shot noise levels vs LO power	24
2.9	Balanced homodyne detection scheme for LIGO	25
2.10	DARM sensitivity curve for the Hanford detector during O3	27
2.11	Ratio of shot noise to dark noise vs squeezing	30
2.12	DARM sensitivity projection at 1 kHz vs squeezing	30
2.13	Reduction factor of LIGO noise vs reduction factor of Johnson noise	31
2.14	Factor improvement to LIGO's observable volume vs reduction factor of Johnson noise	32
2.15	Cost per merger vs reduction factor of Johnson noise	33
2.16	Simple photodiode schematic	34
2.17	Schematic for the DCPD readout used throughout O3	36
2.18	Transfer function of the audio path of the O3 DCPD readout	37
2.19	Voltage noise at the output of the O3 DCPD readout	37
2.20	Input referred noise of the DCPD readout used throughout O3	38
2.21	Input referred noise of the DCPD readout used throughout O3 - zoomed to 20 Hz - 10 kHz	38
2.22	Schematic for a simple frequency dependent readout	39
2.23	Simplified schematic of the new DCPD readout	40
2.24	LISO simulation of the audio path transfer function from the photodiode input to the single ended output	42
2.25	LISO simulation of the voltage noise of the DCPD readout at the single ended output on the audio path	43
2.26	Comparing the input referred noise for a 1 k Ω and 10 k Ω transimpedance resistor	44
2.27	Dimensions of the bulk inductor used in the DCPD readout	46
2.28	Dimensions of the inductor sheets	46
2.29	Pictures showing the disassembly of the inductor	47
2.30	Schematic of the DCPD showing the cross-over of the RF and audio path	49
2.31	Fraction of the photocurrent sent down each path of the DCPD	50

2.32	Impedance as seen by the PD	50
2.33	I-V curve of a typical photodiode	51
2.34	Schematic for the bias monitoring path	53
2.35	LISO model of the transfer function of the bias monitoring path	53
2.36	Assembly diagram of the DCPD vacuum enclosure	54
2.37	Picture of the inductor wires soldered to the PCB	55
2.38	Assembly of the DCPD vacuum case	55
2.39	Thermal model of the DCPD vacuum case	56
2.40	Schematic used to test the DC saturation of the inductors	57
2.41	Inductance vs DC current	58
2.42	Inductance vs DC current with changing temperature	59
2.43	Inductance vs temperature	59
2.44	Schematic of the parasitic capacitance measurement	60
2.45	Phase measurement of the parasitic capacitance test	61
2.46	Picture of the prototype inductor and its mu-metal shield	61
2.47	Transfer function of the induced voltage across the inductor for an injected magnetic field	62
2.48	Measurement of the magnetic fields outside the LIGO HAM-6 chamber	63
2.49	Estimated magnetic pickup across the inductor once installed at LIGO	63
2.50	Transfer function from the photodiode input pin to audio single ended output	64
2.51	Voltage noise floor measured at the output of the audio path	66
2.52	Input referred noise on the audio path of the DCPD readout	67
2.53	Bias monitoring readout transfer function	68
2.54	Voltage noise floor at the output of the bias monitoring output	69
2.55	Transfer function of the photodiode input pin to the RF output	70
2.56	Input referred noise of the audio path with LT1028 as the transimpedance resistor	71
2.57	LISO simulation of the voltage noise for the LT1028 as the op amp in the first transimpedance stage	71
2.58	Input referred noise for the various op amps tested in the transimpedance stage	73
2.59	Photocurrent noise from the Livingston detector during O3	74
2.60	AC signal injection saturation measurement	74
2.61	Schematic for the low pass filter used in the DC signal injection	75
2.62	Passive low-pass filter schematic	75
2.63	Noise floor measurements with DC signal injection	76
2.64	Average DC noise increase	77
2.65	Schematic for inductor noise measurement with DC signal injection	78
2.66	Inductor noise measurement with DC signal injection	78
2.67	Current noise from the inductor for 10 mA DC injection	79
2.68	Input referred noise with the inductor noise included in quadrature	79
2.69	Inductor noise with a magnetic field injection	80
2.70	Picture from the installation of the DCPDs in the HAM 6 vacuum chamber	81
2.71	Transfer function discrepancies from November to February	82
2.72	Noise floor measurements over the course of 3 days	83
2.73	Transfer function discrepancies from February to April	84

2.74	Observed change in the transfer function alongside thermal model of the changes in the transfer function	85
2.75	Input referred noise measured at LIGO Livingston using a signal analyser and breakout box	86
2.76	Input referred noise measured using a breakout box vs the same measurement using the CDS	86
2.77	Transfer function of the VIRGO-style prototype	88
2.78	Voltage noise of the VIRGO-style readout	89
2.79	LISO noise of the VIRGO-style readout	89
2.80	Input referred noise of the VIRGO-style readout	90
2.81	DARM sensitivity projection using the new readouts	92
2.82	DARM sensitivity projection using increased squeezing with the new readouts	93
2.83	Ratio of the change in DARM sensitivity for various squeezing levels using the new readouts	93
3.1	Simplified diagrams of the helioscope experiments	102
3.2	Simplified diagrams of the haloscope experiments	103
3.3	Simplified diagram of an LSW experiment	104
3.4	Axion-like particles coupling constant vs particle mass	105
3.5	Diagram of the ALPS IIc experiment	106
3.6	Schematic for the heterodyne detection scheme	107
3.7	Schematic for the TES detection scheme and plot of resistance vs temperature for the TES electronics	108
3.8	ALPS NL optical table layout	110
3.9	Differential wavefront sensing diagram	112
3.10	Mode matching diagram for QPD 1	116
3.11	Mode matching diagram for QPD 2	116
3.12	Mode matching diagram for auto-alignment steering mirrors	117
3.13	Auto-alignment ‘aligned’ case	117
3.14	Diagram of a lateral shift of the COB	118
3.15	Diagram of adjustment of the curved input mirror to correct a lateral shift of the COB	118
3.16	Diagram of adjustment of auto-alignment steering mirrors to correct for a lateral shift of the COB	118
3.17	Diagram of an angular tilt of the COB	119
3.18	Diagram of the adjustment of the auto-alignment steering mirrors to correct for an angular tilt of the COB	119
3.19	Auto-alignment signal processing flow chart	120
3.20	Picture and PZT schematic for the Thorlabs auto-alignment steering mirrors	123
3.21	Picture and PZT schematic for the PI auto-alignment steering mirrors	123
3.22	Time series for the error signal of QPD 1 X during a two minute lock	125
3.23	ASD for the error signal of QPD 1 X during a two minute lock . . .	126
3.24	ASD for the error signal of QPD 1 Y during a two minute lock . . .	126
3.25	ASD for the error signal of QPD 2 X during a two minute lock . . .	127
3.26	ASD for the error signal of QPD 2 Y during a two minute lock . . .	127
3.27	Time series of the power in transmission of the 250 m cavity during a 2 minute lock	128

3.28	ASD of the power in transmission of the 250 m cavity during a 2 minute lock	129
3.29	Time series of the power in transmission of the 250 m cavity during a 2 hour lock	130
3.30	ASD of the power in transmission of the 250 m cavity during a 2 hour lock	130
3.31	Control voltage for SM 2 Y during the 2 hour lock	131
3.32	Adjusting the DC control voltages using additional pico-motor steering mirrors	132
3.33	Time series of the demodulated signals on each quadrant of QPD 1 during the 2 minute lock	137
3.34	Scaled the x-axis of the time series to see the brief glitches	138
5.1	Gaussian beam propagation diagram	141
5.2	Mode matching diagram	144
5.3	Pin layout for the D-sub 9 and D-sub 25 connectors	145
5.4	Top view of the DCPD PCB	146
5.5	Bottom view of the DCPD PCB	147
5.6	Photo of the inductor used in the VIRGO-style readout	147
5.7	The simulated change in transfer function for various DC photocurrents	148
5.8	First page of the electronics schematic for the DCPD	149
5.9	Second page of the electronics schematic for the DCPD	150
5.10	Assembly of the inductors in their mu-metal shield	151
5.11	Inductor shield with inductors mounted	151
5.12	Alternate representation of the DCPD audio path	152
5.13	Analytic model of the transfer function from photodiode input to the output of the transimpedance stage	153
5.14	Ringling error signals from QPD1 and 2 in X and Y during a 2 minute test of the AA loop	154
5.15	Time series of data outputs while the AA loop is engaged	155
5.16	Times series for the error signal of QPD 1 Y during the 2 minute lock	155
5.17	Times series for the error signal of QPD 2 X during the 2 minute lock	156
5.18	Times series for the error signal of QPD 2 Y during the 2 minute lock	156
5.19	Control voltages during the 2 hour lock	157
5.20	Picture of the ALPS NR optical table	157
5.21	Hard at work on the PC optical table.	158
5.22	A successful day! :)	159
5.23	The inverse DCPD filter	160
5.24	The inverse whitening filter	160

List of Tables

2.1	Parameters from a LISO fit on the audio path transfer function . . .	65
2.2	Key values for the various op amps tested in the readout	72
2.3	The factor improvement in DARM sensitivity with the new DCPD readouts	94
3.1	ALPS Laser specifications	109
3.2	Auto-alignment correction for Pico-motor adjustment	133
5.1	D-sub 25 test voltages	145
5.2	Test voltages for the D-sub 9	146

Acronyms

AA Auto-Alignment.

ADC Analogue to Digital Converter.

ALP Axion-Like Particle.

ALPS Any Light Particle Search.

AS Anti-Symmetric.

ASD Amplitude Spectral Density.

BBH Binary Black Hole.

BHD Balanced Homodyne Detection.

BNS Binary Neutron Star.

CMB Cosmic Microwave Background.

COB Central Optical Bench.

CP Charge-Parity.

DAC Digital to Analogue Converter.

DARM Differential Arm Motion.

DCPD Direct Current PhotoDiode.

DM Dark Matter.

DWS Differential Wavefront Sensing.

EOM Electro-Optic Modulator.

ETM End Test Mass.

FSR Free Spectral Range.

GR General Relativity.

GW Gravitational Wave.

HPL High Power Laser.

HVA High Voltage Amplifier.

IR Infrared.

IRN Input Referred Noise.

ITM Input Test Mass.

LIGO Laser Interferometer Gravitational-Wave Observatory.

LO Local Oscillator.

LSW Light Shining through Walls.

NSBH Neutron Star Black Hole.

OMC Output Mode Cleaner.

PC Production Cavity.

PDH Pound-Drever-Hall.

PLL Phase Locked Loop.

PMC Pre-Mode Cleaner.

PRM Power Recycling Mirror.

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

QCD Quantum Chromodynamics.

QPD Quadrant PhotoDiode.

RC Regeneration Cavity.

RF Radio Frequency.

RL Reference Laser.

RMS Root Mean Square.

ROC Radius of Curvature.

SM Steering Mirror.

SNR Signal to Noise Ratio.

SRM Signal Recycling Mirror.

TES Transition Edge Sensor.

...Physics at its most fundamental, the very fabric of our universe, results directly from the interaction of certain fairly simple rules, and chance; the same description may be applied to the best, most elegant and both intellectually and aesthetically satisfying games.

Iain M. Banks

Chapter 1

Introduction

Interferometers are devices that extract information from interference. Typically this interference comes from electromagnetic radiation, however that need not always be the case. One of the cornerstones of quantum mechanics is the double slit experiment that produces the same interference pattern using optical light [1] as it does with electron beams [2]. Another common misconception is that the term ‘interferometer’ is often used to refer specifically to the Michelson interferometer, however this is only a subset of interferometric devices, albeit an extremely useful one. The name originated from Albert Michelson in 1880 who used it to measure the speed of light [3] to be $3.001 \pm 0.005 \times 10^8$ m/s. He later proposed one could use the instrument detect the mysterious ‘ether’ (also commonly spelled ‘aether’) by measuring the interference pattern of two perpendicular beams of light. In 1887 he worked alongside Edward Morley in the famous Michelson-Morley experiment [4] and determined that there was no evidence of an ether-like substance. Despite the null result this experiment led to a chain reaction of other scientific breakthroughs including the development of the Lorentz transformation [5] and of course the gradual refinement of their experimental device that became known as the Michelson interferometer.

Little more than 10 years after the Michelson-Morley experiment, Charles Fabry and Alfred Perot were utilising a different technique looking at the interference pattern produced within a thin etalon [6]. Using this technique they successfully measured displacements as small as $1/2 \mu\text{m}$ and introduced the now well known Fabry-Perot cavity.

Both the Michelson interferometer and the Fabry-Perot cavity were limited in their precision based on the available light sources of their time, with only incoherent and divergent sources available to them. With the advent of the laser in 1960, a spatially and temporally coherent source of light became available for optical physicists and it enabled the Fabry-Perot resonator to resonantly amplify a light field within its cavity mirrors to unprecedented levels. One of the most notable Fabry-Perot cavities is the

PVLAS experiment [7] that achieves a cavity finesse of over 700,000 and previously held the record for the longest single cavity storage time at 2.7 ms [8]. That record has now been beaten by the Any-Light Particle Search (ALPS) although the result is yet to be published¹. The advent of the laser also led to improvements in the Michelson interferometer. As we will see in this thesis, the two optical techniques have been combined over the last 50 years to produce the LIGO detector [9], the most sensitive metrological device in existence.

Precision measurement using optical interferometry is constantly in a state of improvement and refinement. Less than 10 years ago the LIGO interferometer reached sensitivities high enough that it became the first device to successfully detect gravitational waves. In the coming decades, next generation detectors are expected to exceed the sensitivity of the LIGO detector and probe new frequency regions that were previously inaccessible. The ALPS experiment too, is also utilising new optical techniques that build off the cumulative decades of experience in Fabry-Perot cavity operation and plans to probe regions of parameter space for axion-like particles that have never been tested before. The contribution made by an individual to an experiment is often times small and may appear inconsequential in isolation, but the collaborative nature of modern day experiments means that individual contributions are summed together and the collaboration is able to produce the most advanced devices ever seen.

This thesis outlines my individual contribution in improving both; the ALPS dark matter detector and the LIGO gravitational wave interferometers. As mentioned above, my work in each experiment is only one part of a much bigger project so I will endeavour to go into detail on the work I undertook and also explain how the work affects the experiments as a whole. The structure of the thesis is as follows:

Chapter 2 covers my work within the LIGO experiment. The LIGO experiment is a pair of gravitational wave interferometers built in the US that detects distortions in space-time emitted by dense astrophysical binary systems. My work involved the design and implementation of new photodetector readout electronics that successfully reduced the dark noise in the detector by a factor of 6.5 at 1 kHz. This included the tuning of the frequency response of the new electronics to satisfy the requirements of the LIGO experiment, testing and characterising the performance of the readouts and the investigation of key components, such as the inductor (see section 2.3.1) and the op-amp in the first transimpedance stage (see section 2.3.5). In section 2.1 the background theory and motivation behind the new photodetector electronics are described. In section 2.2 I describe the design of the new photodetector readouts. In section 2.3 the results from the new readouts are discussed. In section 2.4 I project new LIGO displacement noise curves based on previous noise models and

¹The LIGO arms have a storage time of ~ 1 s but are comprised of more than one cavity

new dark noise measurements with the upgraded readouts. In section 2.5 I outline the improvements that could be made for a future upgrade to the readouts.

Chapter 3 covers my work within the ALPS experiment. ALPS is a light-shining-through-walls (LSW) experiment attempting to detect axions or axion-like particles that serve as a strong dark matter candidate particle. My work entailed the design and implementation of an auto-alignment scheme that maintains the input laser's alignment with the cavity eigenmode and successfully reduced low frequency cavity noise by a factor of 10. In section 3.1 the theory and motivation behind the search for axions is covered. In section 3.2 the design of the experiment and the design of the auto-alignment system is discussed. In section 3.3 the results of the auto-alignment system are discussed and in section 3.4 I outline the areas of improvement for auto-alignment before the upcoming science run at the end of 2022.

Chapter 2

LIGO

In 2015 gravitational waves were first detected [10], almost 100 years after they were famously first predicted by Einstein [11]. Since that first detection there has been an immense worldwide effort to build new detectors and upgrade the existing ones, involving over 100 institutions in 18 different countries and over 1400 authors listed on each collaboration paper. This collaboration now contains detectors around the world with LIGO [9], VIRGO [12], GEO600 [13] and KAGRA [14] actively searching for gravitational waves and several next generation detectors in development such as the Einstein Telescope [15], Cosmic-Explorer [16] and the space-based LISA detector [17]. As the existing detectors get iteratively upgraded and new ones come online, the rate of detections has steadily increased, with a total of 90 successful detections at the end of the third observing run.

That being said, gravitational wave astronomy is only 7 years old and is still very much a new field of work when it is compared to traditional astronomy that, arguably spans back thousands of years to the very first sky mappings. As such, gravitational wave astronomy poses an incredibly exciting prospect. With each new advancement in astronomy we have increased our knowledge and understanding of the universe. Exclusively optical light was utilised by early astronomers then radio, infra-red, UV, X-rays, γ radiation and neutrinos. Each of these advancements have led to new theories and understanding by looking at distant objects throughout the universe. Who is to say what more we will discover now that we are finally starting to listen to the universe as well.

The first half of this thesis concerns the LIGO detectors and the work I carried out to improve them. In section 2.1 a familiar ‘text-book style’ approach is taken, where I provide some basic overview of what gravitational waves are, where they come from and how we detect them. In section 2.2 the motivation behind the design of the upgrade to the main LIGO readout electronics is described. In section 2.3 the measurements and results from this new design are shown. In section 2.4 I

project the new sensitivity of the LIGO detector now that the electronics have been upgraded. In section 2.5 I describe potential improvements to the design that could constitute a further upgrade to the readout electronics.

2.1 LIGO Theory

2.1.1 Riding the Wave

Einstein’s theory of general relativity differs from Newtonian gravity by defining gravity as a curvature in space time rather than a force that is independent of the space-time in which it acts. Free falling objects will follow whatever curvature that region of space-time has and in turn any object with mass will curve space-time around itself. In general terms: “Space tells matter how to move; matter tells space how to curve” [18].

One of the byproducts of Einstein’s approach to gravity was the emergence of wave solutions to the field equations. These solutions run counter to the Newtonian approach to gravity which predicts that a gravitational field is always proportional to $1/r^2$, implying that for a moving mass there is some instantaneous action at a distance no matter how much that mass may be accelerating towards or away from the observer. General relativity addresses this “action-at-a-distance” problem by stating that the field does not change instantaneously, rather an accelerating mass will cause a gradient in the gravitational field to travel through space-time at the speed of light. This is the basic premise of gravitational waves. Accelerating masses cause waves in space-time.

To understand how these waves distort space-time and by extension how they interact with regular matter we must delve into some general relativity. In general relativity, space-time is defined by:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \quad (2.1)$$

where μ and ν can take values 0 – 3 to represent one of t, x, y or z in the Cartesian coordinate system and $g_{\mu\nu}$ encodes all information about the space time. In the case of small perturbations it is described by:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (2.2)$$

where $h_{\mu\nu}$ is a perturbation in flat space time and $\eta_{\mu\nu}$ is the Minkowski metric defining flat space time:

$$\eta_{\mu\nu} = \begin{vmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix}, \quad (2.3)$$

Any distortion in flat space time is therefore described entirely by the metric $h_{\mu\nu}$ and is thus the focus of our discussion moving forward. There are many forms that the metric $h_{\mu\nu}$ can take, but by working in the transverse traceless gauge, Einstein's equation becomes a wave equation [19]:

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right) h_{\mu\nu} = 0, \quad (2.4)$$

where $h_{\mu\nu}$ can take the form $h(2\pi ft - \mathbf{k} \cdot \mathbf{x})$ with $f = \frac{|\mathbf{k}|}{2\pi c}$ representing a wave propagating in the direction $\hat{k} = \frac{\mathbf{k}}{|\mathbf{k}|}$ at the speed of light, c .

For a wave propagating along the \hat{z} axis and considering the perturbation $h_{\mu\nu}$ is transverse and traceless we can describe $h_{\mu\nu}$ by:

$$h_{\mu\nu} = \begin{vmatrix} 0 & 0 & 0 & 0 \\ 0 & a & b & 0 \\ 0 & b & -a & 0 \\ 0 & 0 & 0 & 0 \end{vmatrix}, \quad (2.5)$$

which can then be further decomposed into two basis tensors \hat{h}_+ and \hat{h}_\times that represent orthogonal polarisations of the gravitational wave.

$$h_{\mu\nu} = \hat{h}_+ + \hat{h}_\times = \begin{vmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & 0 & 0 \\ 0 & 0 & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{vmatrix} + \begin{vmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & h_\times & 0 \\ 0 & h_\times & 0 & 0 \\ 0 & 0 & 0 & 0 \end{vmatrix}. \quad (2.6)$$

As we will see in section 2.1.2, a common source of GW is from the orbits of a pair of massive objects around their common centre of mass in a binary system. Maggiore gives the definition of h_+ and h_\times for a circular binary orbit as [20]:

$$h_+(t) = \frac{1}{r} \frac{4G\mu\omega_s^2 R^2}{c^4} \left(\frac{1 + \cos^2 \theta}{2} \right) \cos(2\omega_s t), \quad (2.7)$$

$$h_\times(t) = \frac{1}{r} \frac{4G\mu\omega_s^2 R^2}{c^4} \cos \theta \sin(2\omega_s t), \quad (2.8)$$

where G is Newton's gravitational constant, r is the distance to the binary system,

μ is the reduced mass of the binary system, R is the radius of the binary orbit, c is the speed of light, θ is the angle between the normal of the orbital plane and the line of sight to the observer, and $2\omega_s$ is the gravitational wave frequency.

From equations 2.7 and 2.8 some interesting features of gravitational wave strain become apparent:

1. The strain in space-time scales with $1/r$ as opposed to the intensity of electromagnetic radiation which scales with $1/r^2$. This is important for gravitational wave sources that release high frequency EM radiation such as gamma-ray bursts (GRBs), as the EM radiation oscillates too fast to detect their amplitudes and we must instead detect the power of the radiation.
2. The strain signal is maximised when the orbital plane is orthogonal to the line of sight of the observer, i.e. $\theta = 0$.
3. The two strain signals h_+ and h_\times are orthogonal to each other (see the $\cos(2\omega_s t)$ and $\sin(2\omega_s t)$ terms).

Interestingly, although these two basis tensors are orthogonal they actually rotate into each other with a 45° rotation around their direction of propagation, \hat{z} . Figure 2.1 illustrates this by showing how the two polarisations would act on a ring of free falling test masses.

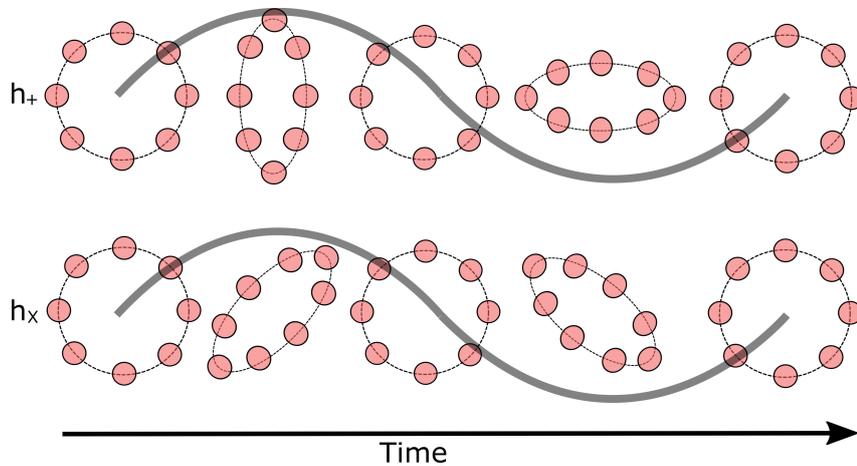


Figure 2.1: Visual representation of the effect of the two polarisations of a gravitational wave, \hat{h}_+ and \hat{h}_\times , from a gravitational wave propagating in the \hat{z} direction into the page. The red circles represent freely falling test masses that are distributed around a ring in the x-y plane. This distribution is distorted by the gravitational wave as it passes through the ring. The waveform of the gravitational wave is superimposed in **grey**. Clearly if the ring from the \hat{h}_+ polarisation is rotated 45° it would resemble the distortion caused by the \hat{h}_\times polarisation. A principle *disadvantage* of using an interferometer to detect gravitational waves is that the device is insensitive to one of these polarisations (see section 2.1.3).

The exact form taken by \hat{h}_\times and \hat{h}_+ does not always follow the simple circular

binary system described in equation 2.7 and 2.8. The models for the strain in space-time depend on the source being considered and become increasingly complex for more complex systems. Fortunately for our discussion however, these models go beyond the scope of this thesis, so we will simply consider the strain in the form of equation 2.7 and 2.8. In the following section we will explore the various sources of gravitational waves throughout the universe.

2.1.2 Astronomical Sources

Following the prediction of gravitational waves from solving Einstein's equation in section 2.1.1, there were many areas of physics considered for potential sources of gravitational waves. When it comes to terrestrial sources (those being sources that can be produced in some laboratory here on earth) they have, for one reason or another been found to be lacking. One of the simplest concepts would be to move a heavy mass through space and then detect the gravitational waves emitted from it. In [19] they calculate that the strain produced by such an experiment would be on the order of $h \sim 9 \times 10^{-39}$, about 16 orders of magnitude below what the LIGO detector is able to resolve. Others considered high energy particle accelerators [21] [22] or even nuclear explosions [23] but all predicted strains that would be undetectable even with today's technology. The best sources for producing gravitational waves with a detectable amplitude are those that are astrophysical.

Brief descriptions of some of the key astrophysical sources for gravitational waves are given below.

Pulsars: As long as the pulsar is not axisymmetric the steady rotation of the star will emit gravitational waves [24] [25]. The amplitude of the gravitational wave is proportional to the square of the rotation frequency and pulsars typically rotate at quite low frequencies. This gives even the fastest rotating pulsar a gravitational radiation amplitude around $h \sim 10^{-39}$ [19]. Experiments searching for this low frequency gravitational wave signal exploit the fact that the pulsar signal is continuous and integrate the signal over long timescales. The resolved signal will increase the longer a measurement runs for and hopefully reach a point of statistical significance. There is one such search that has been looking for over 12 years but has thus far been unsuccessful in making a detection [26]. The inference of gravitational waves from a decaying pulsar binary is discussed in section 2.1.3.

Supernovae: When a star has burned all of its fuel the inner core begins to contract and become supported largely by electron degeneracy pressure. Once this inert inner core reaches the Chandrasekhar mass ($> 1.5 M_{\odot}$) [27], electron degeneracy is no longer sufficient to counteract the gravitational collapse, so the inner core implodes and rapidly reaches neutron star-like densities, only stopped by the neutron degeneracy pressure [28]. Without the support of the now imploded inner core, the

outer core also collapses and rebounds off the inner core, causing a *bounce* in the collapse of the dying star. As long as this collapse is not symmetric there will be a burst of gravitational waves emitted due to this bounce in the core [29]. These events are somewhat rare with only 2-3 Core Collapsing SuperNovae (CCSN) per century occurring within our galaxy. Somewhat optimistic estimates give a peak strain at $h \sim 10^{-21}$ for sources 10 kpc away [29] which would be within LIGO sensitivity range for the third observing run (O3) however there have been no observations of GW associated with supernovae yet, implying that there have either not been any viable sources within that range or the supernovae model is too generous in the predicted strain. The most recent data from the third observing run of LIGO was used to try to identify the GWs produced by two known CCSN at a distance of 11 MPc but were unsuccessful [30] and only managed to set upper limits.

Stochastic Background: This can be the result of numerous overlapping merger signals [31] or alternatively from quantum gravitational processes in the early universe. The former is described in the **compact binaries** description. In some ways the latter is analogous to the 2.7 K cosmic microwave background (CMB), however CMB experiments are only able to observe photons produced when matter and radiation decoupled from each other - around 3×10^5 years after the big bang [32]. Information about earlier epochs can of course be inferred from these snapshots, e.g. the density fluctuations of the present epoch were produced much earlier and are compatible with inflation predictions [33]. GW on the other hand are predicted to have decoupled as early as 10^{-44} s after the big bang and so represent a more direct observation of the state our early universe. Expected stochastic ‘signals’ would actually appear as an additional noise source in the detector and are expected to scale with $f^{-1/3}$ which means only detectors that can reach low (\sim mHz) frequencies would be able to detect them. LISA will be a space based detector that aims to probe the GW parameter space with a sensitivity down to $h \sim 10^{-20}$ at mHz frequencies [34].

Black Hole Formation: The asymmetric collapse of a massive star into a black hole can emit GW in a similar way to the emission of GW from supernovae. This is predicted to constitute a part of the stochastic background for LISA and would occupy predominantly low frequencies at \sim nHz [35].

Compact Binaries: The most well known source of gravitational waves and the only type that has been successfully detected. Compact binaries generally refers to two dense objects in an orbit together that gradually lose energy due to gravitational radiation. This causes the orbit radius to decrease and the two objects to orbit at higher and higher frequency. This initial process is known as the *inspiral* and emits low frequency and low amplitude gravitational waves. Just before the two objects completely merge their orbital frequency rapidly increases and a burst of

high frequency, high energy gravitational waves are released in what is known as the *chirp* (see figure 2.2). The term ‘Compact binaries’ refers to any two objects that are in a binary object together and so includes white dwarf binaries that are expected to make up the stochastic background of GW for the LISA detector. The most commonly detected binary mergers are Binary Black Hole (BBH) systems which produce a larger amplitude GW (due to their higher mass) but are less common than white dwarf systems and hence would not make up a stochastic background, but rather a distinct signal in space-time strain. Most of the signal detected from BBH systems comes specifically from the high frequencies emitted during the merger which occur at around 50 Hz.

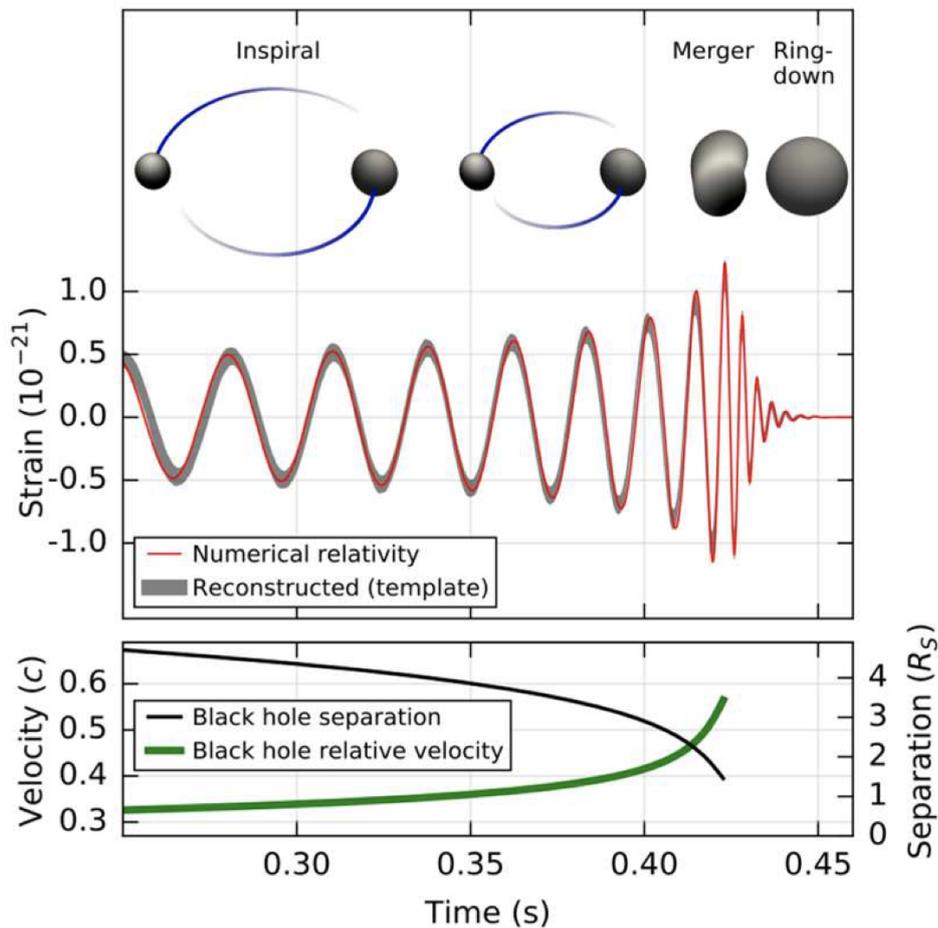


Figure 2.2: **Grey:** The reconstructed gravitational wave signal from Hanford of the first gravitational wave detection GW150914 [10]. **Red:** The best matching waveform predicted with numerical relativity. GW150914 was produced by the merger of two black holes with masses $M_1 = 36M_\odot$ and $M_2 = 29M_\odot$ at a distance of one billion light years from earth.

As of the end of the third observing run a total of 90 compact binary mergers had been observed using the LIGO and VIRGO detectors (see section 2.1.3 for a description of their operation) [36]. The detected inspiral mergers have predominantly

been binary black hole mergers (BBH) with only occasional neutron star - black hole mergers (NSBH) or binary neutron star mergers (BNS). The latter two are significantly harder to detect, primarily due to the fact that the gravitational wave signal scales linearly with the mass of the system so for lighter objects such as BNS a fainter signal is produced. For context, during the O3b LIGO run, of the 35 total detections, 32 were reported as BBH events, 3 were suspected NSBH events and there were 0 BNS events [36].

Some examples of interesting observations are described below to give an overview of the scientific outcomes of gravitational wave detection:

1) The event GW190814 consisted of a merger with an extremely imbalanced mass ratio [37]. One object had mass $m_1 = 23.2 M_\odot$ (clearly a black hole) and the other had $m_2 = 2.59 M_\odot$. $2.59 M_\odot$ is an intriguing compact object for astronomers since it appears to be too large for a typical neutron star ($m \sim 1.3 M_\odot$) and too small for a typical black hole ($m \sim 5 M_\odot$) [38]. The mysterious object exists in the so called *mass gap* and has prompted speculation over the formation processes of high mass neutron stars or low mass BH [39] and made headlines around the world [40] upon its announcement.

2) Binary Neutron Star (BNS) mergers were first announced with GW170817 [41] in 2017. This event also prompted suitable observatories to search the predicted sky-location of the merger and γ -ray bursts were observed [42] alongside the initial GW detection. The joint observations prompted follow up research on the population of BNS systems [43], the GW emitted from the decay of BNS systems [44], physical constraints on gravity [42] and cosmology [45].

BNS systems have lower mass than BBH systems and emit gravitational waves at a higher frequency. Typically the low frequency inspiral is around 100-400 Hz as in GW170817. Although the inspiral emits lower amplitude GWs than the signal emitted during the merger, the inspiral signal can be summed coherently as the objects orbit each other which means that, in BNS systems, the inspiral is typically detected rather than the chirp. The chirp in BNS systems occurs in the kHz region [46] however the noise floor of the LIGO detector is normally too high to observe the GW emissions from the merger and post merger. Lowering the noise floor at these frequencies would potentially allow further understanding of the physics behind neutron stars. As we will see in section 2.1.5 the dark noise (that being the noise of the photodetector readout detector electronics when there is no light on the photodetector) significantly contributes to the overall noise floor of the detector in the frequency region 600 Hz - 3 kHz.

Detectors all around the world now routinely undergo iterative upgrades to increase their sensitivities to space-time distortions. Increasing the sensitivity of the detectors will enable the LIGO Scientific Collaboration (LSC) to detect not only fainter signals

coming from quieter mergers but also increase the range of the detector to events further out in the universe. For example just taking BNS events, it is expected that the full network of upgraded LSC detectors may be able to detect up to 100 BNS events each year [41], compared to the 0 BNS events detected during O3b.

In the following section we will explore how the gravitational waves from these mergers are converted into a detectable signal.

2.1.3 From Source to Detection

In principle a gravitational wave detector can be constructed in many different ways, indeed there were many experiments that sought to convert passing gravitational waves into some detectable signal. Most notably Joseph Weber in 1968 (erroneously) claimed to have detected gravitational waves using the so called ‘Weber bars’ [47]. Weber bars are resonant devices that consist of large cylindrical structures that would change size as a gravitational wave passed through them and this distortion would be measured using a piezoelectric device. Since they are resonant devices, they only have a narrow sensitivity band and could detect changes in length down to 10^{-16} m but only at specifically 1657 Hz (its first longitudinal mode). Despite repeatedly claiming to have made successful detections [48], his results were never reproduced by other experiments. Various theories were put forward to explain his false positives including mistakes and subjectivity in his analysis or in the incorrect prediction of induced size fluctuations as the wave passes through the bar but by the late 1970s the general scientific consensus was that he was incorrect in his announcement.

Others looked to pulsars as a possible measuring stick for gravitational waves. There was actually some previous evidence for gravitational waves with a pulsar origin. In 1982 an indirect measurement of gravitational waves was carried out by observing the evolution of the interval between the arrival of pulses from a binary pulsar system in order to measure the decaying orbit of the two stars. They measured an overall shift of around 10 s over the period of 15 years which matches the predicted energy loss due to gravitational radiation of the system to an accuracy of about 1% [49]. Although this observation is a strong indicator of the existence of gravitational waves, it crucially told physicists nothing about how the gravitational waves interacted with matter. In a sense this inference of gravitational waves in 1982 served as the smoking gun for the existence of gravitational waves but in terms of direct measurements the pulsar approach has been less successful. Observatories have investigated the fluctuations in the arrival time of emissions from pulsars [50] that could be explained by a stochastic background of gravitational waves in the nHz frequency range. Despite these observatories working with 12.5 years worth of data there have been no statistically significant observations as of writing this thesis although it is hoped this will change within the next few years of data [26].

The only approach that has been successful in detecting gravitational waves is the one that utilises the optical technique first developed in the famous Michelson-Morley experiment to investigate the so-called ‘aether’ [4]. The gradual optimisation of an optical technique that dates back to 1887 has resulted in the development of the most precise instruments ever created by mankind, measuring displacements smaller than $1/1000^{\text{th}}$ the diameter of a proton. The following section details one of the most well known instruments in the field of gravitational wave research, the Michelson interferometer.

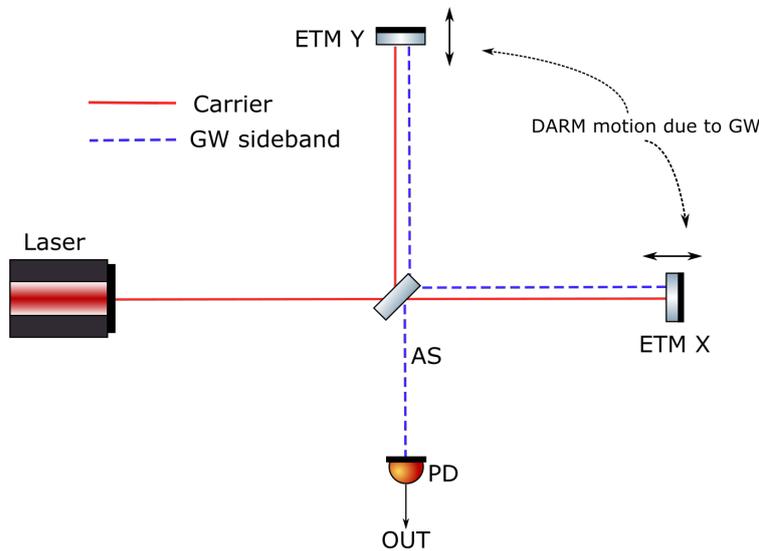


Figure 2.3: A basic Michelson interferometer locked to the dark fringe. Some common nomenclature for the LIGO experiment is introduced. ETM stands for End Test Mass in the X or Y direction. AS stands for Anti-Symmetric port, sometimes also referred to as Dark Port due to the light level at this port being held at a minimum. DARM refers to Differential Arm Motion i.e. when the arms move in different longitudinal directions along their optical axes.

The interaction of light within the Michelson interferometer is well covered in many textbooks; however some basic explanation is covered here for background to future topics.

To understand the eventual signal emerging at the AS port of the interferometer we will be working with electric fields. The electric field of the light is described generally by:

$$E = E_0 e^{i(\omega t - \vec{k} \cdot \vec{z})}, \quad (2.9)$$

where E_0 is the field amplitude, $\vec{k} = \frac{2\pi}{\lambda}$ is the wave number, \vec{z} is the direction of the propagation of the photons, ω is the angular frequency of the photons and t is time. Since the $e^{i(\omega t)}$ is a term common to all fields it is ignored from here on.

The light seen out the AS port, E_{AS} , and the light back reflected from the beam splitter, E_{Refl} , is described by:

$$E_{AS} = t_M E_{in}, \quad (2.10)$$

$$E_{Refl} = r_M E_{in}, \quad (2.11)$$

where E_{in} is the field amplitude input to the interferometer from the laser and t_M and r_M are respectively the transmission and reflection coefficients of the interferometer. The interference of the light within the interferometer arms affects the overall transmission and reflection coefficients and depends on the phase of the light accumulated upon its propagation along the arms:

$$t_M = \frac{1}{2} \left(r_x e^{i2\phi_x} - r_y e^{i2\phi_y} \right), \quad (2.12)$$

$$r_M = \frac{1}{2} \left(r_x e^{i2\phi_x} + r_y e^{i2\phi_y} \right), \quad (2.13)$$

where $r_{x/y}$ are the reflectivities of ETM X/Y respectively and $\phi_{x/y}$ is the phase accumulated by the light as it travels down the respective x/y arm. The accumulated phase is related to the $e^{i(-\vec{k} \cdot \vec{z})}$ term in equation 2.9. The phase accumulated by the light as it propagates along the interferometer arm of length, L , is given by:

$$\phi = \frac{2\pi}{\lambda} \int_0^L g dx, \quad (2.14)$$

where λ is the wavelength of the light, L is the effective length of the interferometer arm and g is the metric defined by equation 2.2 in section 2.1.1. For a time dependent perturbation $h(t)$ caused by a gravitational wave, this can then be expressed by:

$$\phi(t) = \frac{\omega L}{c} h(t), \quad (2.15)$$

where ω is the angular frequency of the light in the interferometer and c is the speed of light. If the interferometer arms are of equal length, there is no light transmitted through the interferometer and the AS port is dark. By conservation of energy all the light is reflected from the interferometer and is sent back to the laser.

From equation 2.15 we therefore see that the phase accumulated by the light within

the interferometer arms encodes information about the strain induced by the gravitational wave. The question then becomes how do we extract the signal from the phase of the light. By considering the common and differential motion of the interferometer arms the situation becomes clearer.

Differential arm motion, coming from differential noise or a GW signal, will result in a differential phase. For simplicity we will simply refer to the phase within the arms as ϕ without explicitly mentioning the time dependence. The differential phase is given by:

$$\phi_- = \phi_x - \phi_y, \quad (2.16)$$

whilst common noise in the interferometer arms, coming from either common arm motion or common input noise such as laser frequency noise, will result in a common phase:

$$\phi_+ = \frac{1}{2}(\phi_x + \phi_y). \quad (2.17)$$

This allows us to solve equations 2.16 and 2.17 to find:

$$\phi_x = \phi_+ + \frac{1}{2}\phi_-, \quad (2.18)$$

$$\phi_y = \phi_+ - \frac{1}{2}\phi_-. \quad (2.19)$$

Similarly we can make substitutions for the reflection coefficients of the x and y mirrors in the interferometer arms:

$$\Delta r = r_x - r_y, \quad (2.20)$$

$$r = \frac{1}{2}(r_x + r_y), \quad (2.21)$$

which can be solved to give:

$$r_x = r + \frac{1}{2}\Delta r, \quad (2.22)$$

$$r_y = r - \frac{1}{2}\Delta r. \quad (2.23)$$

We can now use equations 2.18, 2.19, 2.22 and 2.23 to express equations 2.12 and 2.13 in terms of the sum and difference of the x and y arm phase and mirror reflectivity:

$$t_M = e^{i2\phi_+} \left(ir \sin \phi_- + \frac{\Delta r}{2} \cos \phi_- \right), \quad (2.24)$$

$$r_M = e^{i2\phi_+} \left(r \cos \phi_- + i \frac{\Delta r}{2} \sin \phi_- \right). \quad (2.25)$$

In a perfect interferometer we would have identical mirror arms such that $\Delta r = 0$ and would operate the interferometer on the dark fringe such that no light from the input couples to the AS port. In practice this is impossible to achieve and there will always be some light present at the AS port that is referred to as the contrast defect.

When it comes to detecting the phase of an electric field we encounter a new problem as photodiodes are not fast enough to detect light fields that are oscillating faster than ~ 300 THz frequencies. Instead the *power* from the AS port is detected which is given by $P = E^*E$:

$$P_{AS} = P_{BS}|t_M|^2 = P_{BS} \left(r^2 \sin^2 \phi_- + \left(\frac{\Delta r}{2} \right)^2 \cos^2 \phi_- \right), \quad (2.26)$$

$$P_{RefI} = P_{BS}|r_M|^2 = P_{BS} \left(r^2 \cos^2 \phi_- + \left(\frac{\Delta r}{2} \right)^2 \sin^2 \phi_- \right). \quad (2.27)$$

From equations 2.26 and 2.27 we can see that the common phase term has cancelled out and that the power at the AS port is only dependent on the differential phase of the interferometer arms. In section 2.1.1 we saw that a gravitational wave signal can be decomposed into two basis tensors that rotate into each other by a 45° rotation along their direction of propagation. Due to the interferometer's insensitivity to common arm motion, it will be sensitive to only one of these quadratures. The key benefit of using an interferometer instead comes from the rejection of common noise in the interferometer arms.

To increase the sensitivity of the interferometer to phase changes in the arms, Fabry-Perot cavities are formed within the X and Y arms by inserting cavity mirrors ITM X/Y.

For our purposes we do not need to consider basic optical cavity theory which can be found in [51] [52] and will instead skip to the end result which tells us that for an interferometer with resonant arm cavities our signals at the reflected port and AS port now have terms containing the phase gain off the cavity arms $A_\phi = \frac{r'_c}{r_c}$. r'_c and r_c are the cavity reflectivity on resonance and r'_c is its derivative with respect to intra-cavity phase, and are defined respectively by:

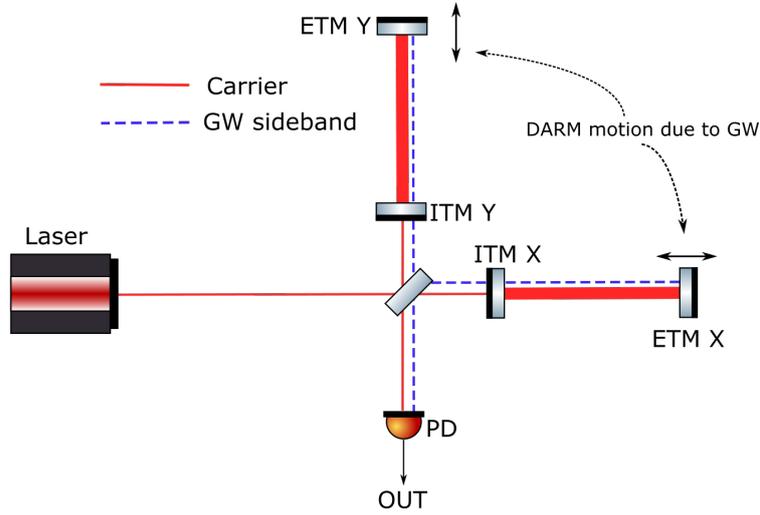


Figure 2.4: A Michelson interferometer with Fabry-Perot arm cavities formed by the insertion of Input Test Masses (ITM) X/Y with the end mirrors ETM X/Y respectively.

$$r_c = \frac{r_1 - r_2}{1 - r_1 r_2}, \quad (2.28)$$

$$r'_c = \frac{(1 - r_1^2) r_2}{(1 - r_1 r_2)^2}. \quad (2.29)$$

The phase gain, A_ϕ , can be understood conceptually by imagining that one of the end mirrors is moved by a microscopic amount which results in some small phase shift in the light's electric field as it travels up and down the interferometer arm according to equation 2.15. If we were to force the light to repeatedly traverse the same distance by forming a Fabry-Perot cavity, it would experience the same microscopic offset with each lap and would cause the phase to accumulate with the number of round trips that the light completes within the arm. Implementing this phase gain, equations 2.24 and 2.25 then become:

$$t_M = e^{i2A_\phi\phi_+} \left(ir \sin A_\phi\phi_- + \frac{\Delta r}{2} \cos A_\phi\phi_- \right), \quad (2.30)$$

$$r_M = e^{i2A_\phi\phi_+} \left(ir \cos A_\phi\phi_- + \frac{\Delta r}{2} \sin A_\phi\phi_- \right), \quad (2.31)$$

where Δr and r now refer to the reflectivities of the two arm cavities rather than

just the two end mirrors.

The light power levels in the arms are increased once again by installing another mirror, this time in the input path to the interferometer.

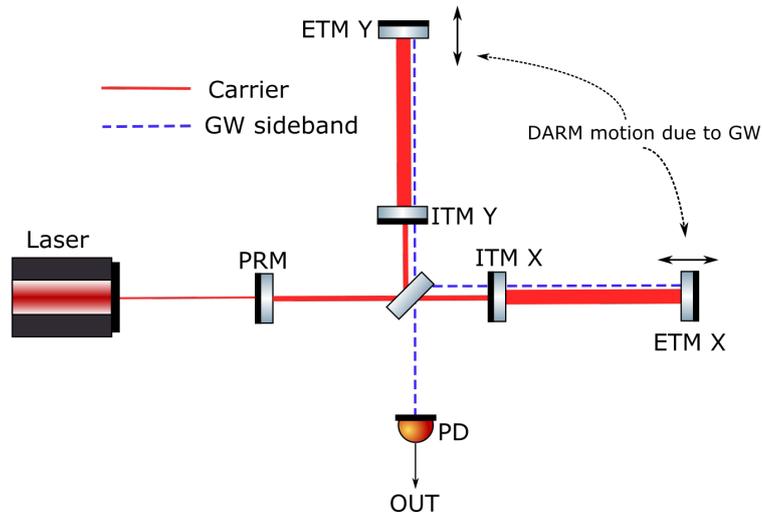


Figure 2.5: Layout for an interferometer with a power recycling cavity formed between the Power Recycling Mirror (PRM) and the end mirrors ETM X/Y.

The insertion of the Power Recycling Mirror (PRM) forms a Power Recycling Cavity (PRC) with ETM X/Y respectively. The reflectivity of the PRM is selected such that it matches the reflectivity of the rest of the interferometer. The PRC is then critically coupled and all of the input power from the laser is coupled into the interferometer arms where it circulates until the light is eventually lost via scattering and other losses. This effectively increases the P_{BS} term in equations 2.26 and 2.27.

Even with all of this work we would still be unable to detect the small changes induced by the gravitational wave in the phase of the electric field. To do this we need a method to convert the phase modulation into amplitude modulation. This is achieved by the various readout schemes that have employed by LIGO over the years.

2.1.4 Readout Schemes

The method of converting the gravitational wave phase modulation into an amplitude modulation signal is not a straightforward problem and has been the subject of previous LIGO upgrades. During the early science runs of the LIGO detector (pre

2008) a heterodyne readout was used, shown in figure 2.6.

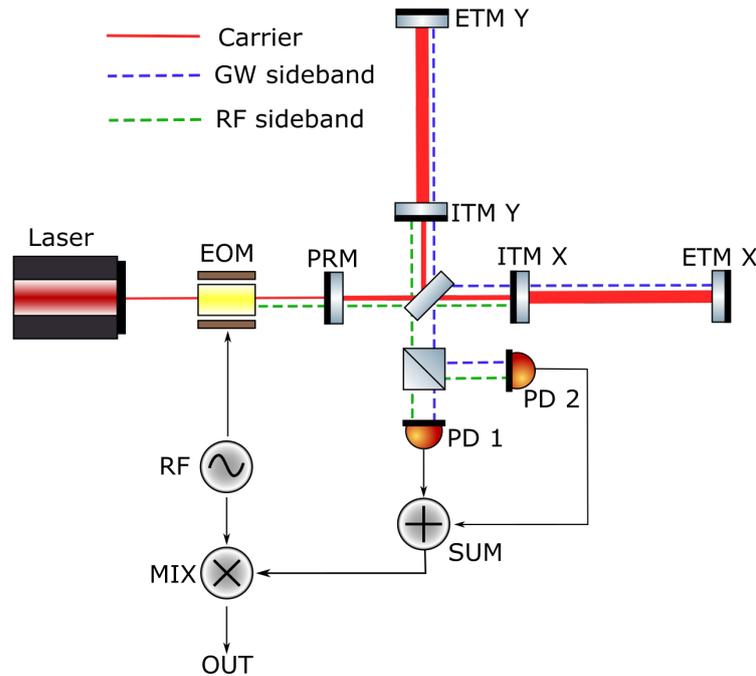


Figure 2.6: Heterodyne detection readout employed during the early ‘initial LIGO’ runs (pre 2008). **Red:** The carrier beam at 1064 nm. **Blue:** the Gravitational wave sidebands induced by a gravitational wave passing through the detector. **Green:** RF sidebands imposed on the 1064 nm carrier. The RF sidebands are resonant within the PRC but not the arm cavities. Two photodiodes (PDs) are not strictly required by the readout but are used at LIGO to reduce the total light intensity per photodiode and provide a diagnostic null stream by taking the difference in the PD signals. The gravitational wave signal is provided by electronically demodulating the sum of the two readouts with the LO and is labelled ‘OUT’ in the diagram. In this setup there is ideally no carrier light output of the AS port.

This heterodyne scheme is well known by any early academic working with precision optics as it is frequently used in the Pound-Drever-Hall (PDH) [53] technique to lock a Fabry-Perot cavity. Optical sidebands are imposed on the carrier light by an Electro-optic Modulator (EOM) at the LO frequency f_{het} . The optical combination of the LO sidebands f_{het} and gravitational wave sidebands f_{GW} produces a beat note in the photocurrent generated by the photodiode that contain $f_{het} + f_{GW}$ terms. The gravitational wave signal containing purely the f_{GW} term is extracted via an additional electronic demodulation in the mixer with the original LO frequency, f_{het} . There are two pairs of sidebands used in conjunction for length sensing and

alignment control of the various cavities and length locks within the interferometer. These sidebands are at 9 MHz and 45 MHz.

The heterodyne readout was used up to the end of the S5 science run which completed in mid 2007 and achieved DARM sensitivities down to $10^{-19} \text{ m}/\sqrt{\text{Hz}}$ at the most sensitive frequencies [54] but was of course unsuccessful in detecting gravitational waves and was only able to set upper limits [55]. As part of the upgrades to the detector it was decided to change the readout from heterodyne to DC readout. This was motivated by several noise related reasons:

- There was noise on the RF sidebands that was being coupled into the final signal.
- There was an imperfect spatial overlap between the signal beam and the RF sidebands. This is because the RF sidebands are not actually resonant in the arm cavities, only the Power Recycling Cavity (PRC).
- There is additional shot noise in heterodyne readout compared to other readouts due to interference between the upper and lower RF sidebands.
- There were plans to implement squeezing ¹ at LIGO but implementation of squeezing with heterodyne readout is more complicated than other readout methods because it requires squeezing in the audio and RF band as opposed to other schemes that only require squeezing in the audio band.

Promising results from the GEO 600 detector [56] [57] and the 40 m LIGO prototype in Pasadena [58] motivated the LIGO group to use a form of homodyne readout called DC readout as an alternative to the heterodyne scheme.

DC readout is one of the oldest forms of reading signals from an interferometer [4] but was not considered viable for GW experiments due to the coupling of laser power noise into the readout. The improved noise suppression of later GW experiments meant that DC readout became viable once more. DC readout alleviated many of the problems encountered with the heterodyne readout. Primarily the RF sidebands are filtered by an Output Mode Cleaner (OMC) in the DC readout (not pictured) so there is no longer an interference effect between upper and lower sidebands which ultimately means there is a lower shot noise level. It also means that there is no coupling of noise from the RF sidebands into the light output from the AS port of the interferometer. There is also no longer a problem encountered with the spatial overlap of the LO and the signal sidebands since, in the DC readout, both beams are resonant within the same cavities so spatial overlap occurs by design. And finally since there are no longer RF sidebands involved in the signal extraction it is easier to implement squeezing since squeezing is only required in the audio band and not

¹injecting squeezed light into the interferometer reduces the uncertainty in one quadrature of light at the expense of an increase in the orthogonal quadrature

audio + RF range.

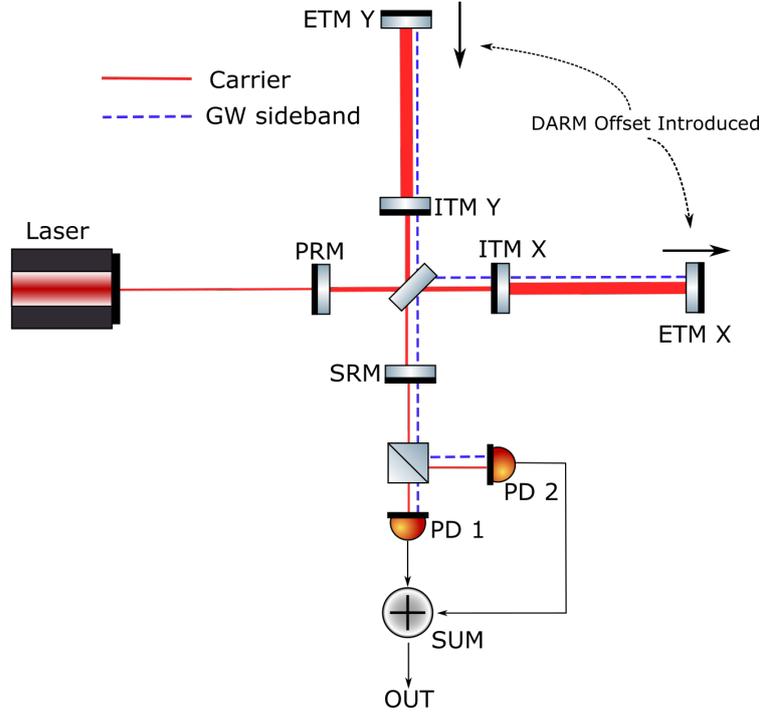


Figure 2.7: DC readout currently installed on the LIGO interferometers. A small differential arm offset is intentionally applied to introduce a *dark fringe offset*. The carrier light exiting the anti symmetric port now acts as the LO. DC readout also introduced a Signal Recycling Mirror (SRM) which is used to tune the interferometer response [59][60]. Not shown is the output mode cleaner OMC used to filter higher order mode light that is seen from the AS port.

From equation 2.30 the electric field at the AS port is defined by:

$$E_{AS} = t_M E_{in} = e^{i2A_\phi\phi_+} \left(ir \sin A_\phi\phi_- + \frac{\Delta r}{2} A_\phi\phi_- \right). \quad (2.32)$$

For perfect interferometer mirrors $r_x = r_y = 1$, which results in, $\Delta r = 0$ and $r = \frac{1}{2}(r_x + r_y) = 1$. A DC field is introduced at the AS port on top of the varying GW sidebands by applying a static differential arm offset. The differential phase may now be written: $A_\phi\phi_- = \phi_{DC} + \phi_{GW}$. The field at the anti-symmetric port then becomes:

$$E_{AS} \approx iE_{in} e^{i2A_\phi\phi_+} (\phi_{DC} + \phi_{GW}) = e^{i2A_\phi\phi_+} (E_{LO} + E_{GW}), \quad (2.33)$$

where in the DC readout scheme the DC field at the AS port serves as the local oscillator and is labelled E_{LO} . The power at the AS port is then:

$$P_{AS} = |E_{LO}|^2 + E_{DC}^* E_{GW} + E_{LO} E_{GW}^* = |E_{LO}|^2 + 2\Re(e(E_{LO} E_{GW}^*)), \quad (2.34)$$

where second order terms in E_{GW} have been neglected. As before we also see that the common phase terms have cancelled out. Given that the GW signal term is contained in $E_{LO} E_{GW}^*$, the measured GW signal is proportional to the square root of the power in the LO. The shot noise arises due to the Poissonian distribution in the arrival times of photons incident on the photodetector. It is straightforward to calculate the generated photocurrent and the induced shot noise (see also section 2.1.5 and 2.1.6) in the photocurrent using equations 2.35 and 2.36.

$$I_{PD} = \frac{e\lambda P}{hc}, \quad (2.35)$$

where I_{PD} is the generated photocurrent, e is the fundamental charge, λ is the wavelength of the carrier, P is the power of the incident light, h is Planck's constant and c the speed of light. Equation 2.35 has also assumed a photodetector with 100% quantum efficiency. The shot noise spectral density in this photocurrent is given by:

$$\tilde{I}_{shot} = \sqrt{2eI_{PD}}, \quad (2.36)$$

where \tilde{I}_{shot} represents the shot noise amplitude spectral density and is expressed in A/\sqrt{Hz} .

In the case of DC readout the signals incident on the two photodiodes are correlated so the sum of the two photodiodes gives the measurement of the gravitational wave signal, whilst the difference gives a measurement of the shot noise which is uncorrelated between the two detectors [61].

An important question then arises in all readout schemes: How large do we set the power in the LO to be? We see from equations 2.34 and 2.36 that both the GW signal level and the shot noise level scale with the square root of the LO power so at first glance it appears that we can choose any arbitrary LO power and achieve the same overall SNR. In reality of course things are never so simple.

The upper limit to the LO power level comes from a technical limit. Increased power on the photodiode can cause the noise on the readout electronics to increase (explored in section 2.1.7) and if the power is large enough it can even burn the photodiode.

We may therefore be tempted to make the LO power very small so that if there is accidentally an increase in the light exiting the AS port there is no chance of it damaging our photodiodes. Here though we run into another limitation which is that other noise sources such as the dark noise and the shot noise from the contrast defect are independent of the LO power. These noise sources are incoherent and therefore sum together in quadrature. If the controllable noise (that being the shot noise from the LO light) gets too close to the technical noise floor (the contrast defect shot noise and dark noise of the electronics) then the overall Signal to Noise Ratio (SNR) decreases due to contamination from the technical noise sources.

A balance must therefore be struck where the LO power is selected to be high enough that the SNR is minimally affected by the technical noise floor, but low enough that there is no additional noise introduced and no risk to the photodiodes. Although there is no official number dictating the headroom one desires between controllable noise and technical noise a good rule of thumb is to have the controllable noise a factor of 10 above the technical noise floor. In the case of the LO power however, the shot noise scales with the square root of the power of the LO so one would need to use a LO power 100 times larger than the contrast defect which would be 1 W of power. This is generally too large a power for the photodiode to handle so instead a LO power 10 times larger is selected with gives a factor of $\sqrt{10}$ headroom between the two noise levels as shown in figure 2.8. Although the factor of 10 rule of thumb is not followed in this case, we will use it as our measuring stick in section 2.1.6 in deciding how successful the upgrade to the LIGO electronics are.

The contrast defect results in about 10 mW of light split between two PDs. This light does not contain any GW signal but *does* contribute to the shot noise. Using equation 2.35 this gives a photocurrent of 4.2 mA per photodiode and using equation 2.36 a corresponding source of shot noise of $37 \text{ pA}/\sqrt{\text{Hz}}$ purely due to the contrast defect.

The dark fringe offset is chosen to allow $10 \times 10 \text{ mW} = 100 \text{ mW}$ of light to leak out the AS port to serve as the LO which corresponds to a DC photocurrent of about 43 mA per photodiode and a shot noise level of $120 \text{ pA}/\sqrt{\text{Hz}}$. In practice the power in the LO varies slightly and it is typical to use 50 mW which corresponds to roughly 20 mA per photodiode. These are numbers we will revisit later as we discuss the design of the electronic readouts within section 2.2.

It is worth pointing out that RF sidebands are still used in the experiment but are not sent to the DC PhotoDiode readouts (DCPDs). The 9 MHz and 45 MHz sidebands are used in conjunction to achieve alignment and initial length lock of the various optical cavities. After initial lock is achieved the length lock is passed off to DC readout whereas alignment control is still maintained using RF readout, using separate photodiode readouts. The light sent to the DCPDs passes through

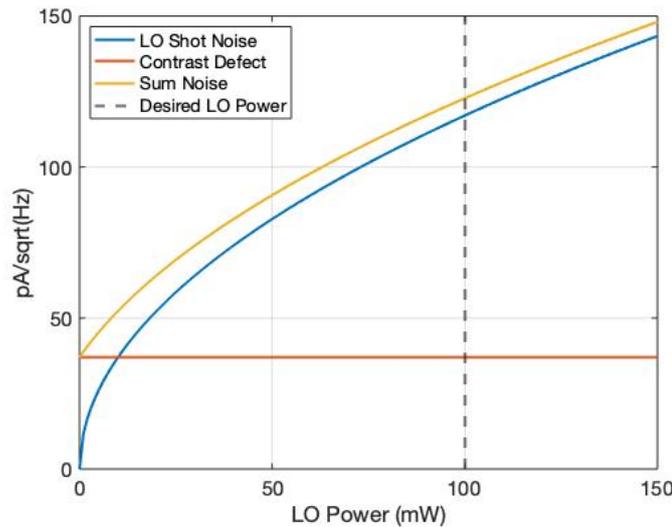


Figure 2.8: The LO shot noise is shown in **blue** and scales with the square root of the power of the LO. The power due to the contrast defect at the dark port is independent of the LO power and its shot noise contribution is shown as the **orange** horizontal line on the plot. The sum in quadrature of the two noise sources is shown in **yellow**. When the LO power is high, the shot noise from the contrast defect contributes less to the overall sum in quadrature. The vertical line labels the LO power that is 10 times larger than the power due to the contrast defect.

an Output Mode Cleaner (OMC) which attenuates the RF sidebands.

Not mentioned so far are the 3 MHz sidebands that will be used in locking the squeezer phase. The signal level for the 3 MHz sidebands has been estimated by [62] to be $54 \mu\text{A}$ in the photocurrent during normal operation and will be read out by the DCPD readouts. This is discussed in section 2.2.5.

DC readout has been employed at LIGO for over a decade now and was successfully used during the first gravitational wave detection in 2015 [10] and in the most recent observation run (O3b) which completed in March 2020 and achieved 35 detections [36]. It is so reliable that it will even be used for the first half of the next observation run (O4a). There are downsides to the DC readout however, which mean that it will soon be replaced for the second half of the A+ upgrade. The key downside is that there is no easy way to change the phase of the LO with respect to the gravitational wave signal since both are co-propagating within the interferometer cavities. Although there is some limited ability to adjust the phase using DC readout it is ultimately unsatisfactory for future upgrades in which it is planned to use quantum non demolition readout schemes [61]. There are also problems associated with having a dark fringe offset at the AS port since back scattering of the light can couple noise back into the interferometer.

Figure 2.9 outlines the setup proposed for the upcoming LIGO A+ upgrade to move

from DC readout to Balanced Homodyne Readout (BHD) [63]. In this design the LO is provided by picking off the back reflection of light from the main PRM at the Pick Off Port (POP).

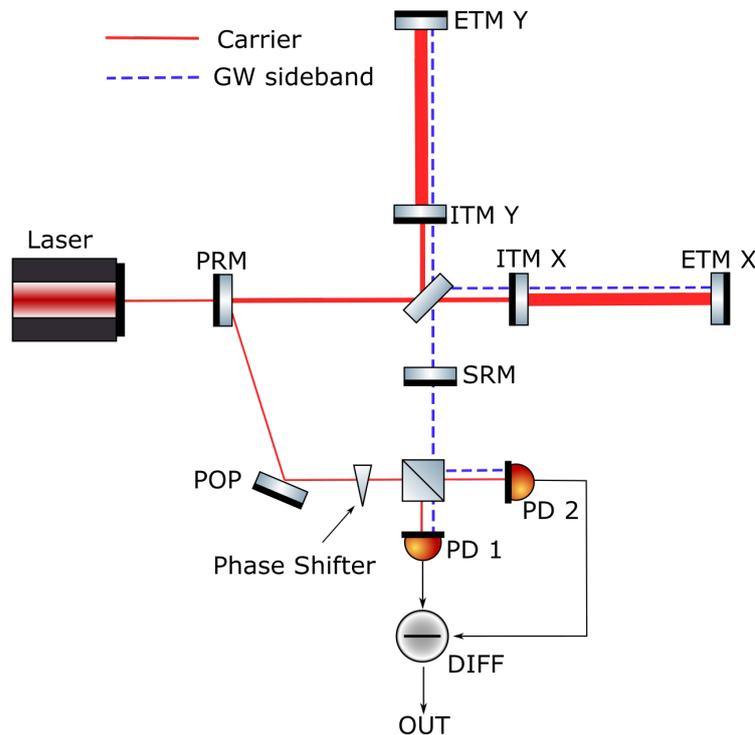


Figure 2.9: Balanced Homodyne Detection (BHD) readout proposed for the Advanced LIGO upgrade. The LO is provided by picking off the main beam from the PRM mirror at the Pick Off Port (POP). This is a simplified diagram that does not show the geometry of the power recycling cavity that enables the pick-off of the beam seen in [63]. The gravitational wave signal comes from the mixing of the LO with the GW sidebands. One benefit of this setup over the DC readout is the ability to tune the phase of the pick off LO, represented here by the labelled phase shifter. Note that in this configuration the gravitational wave signal is produced by the *difference* in the signals from PD 1 and PD 2 whilst the sum would be a measurement of purely the noise on the LO [61].

The BHD scheme provides significant benefits over the DC readout by enabling phase changes of the LO but at the cost of increased complexity. The phase noise on the LO must be lower than the phase noise of the signal, which may necessitate suspended optics and suppression of the jitter of the beam pick off. There is also no longer a guaranteed spatial overlap between the LO and GW sidebands since the

LO is now picked off from the Pick off Port (POP) and does not co-propagate with the GW sidebands out of the AS port. Precise alignment of the LO pick off beam is critical, but with the improvements in optical alignment techniques over the last decade it is expected that this will be accomplished by the LIGO collaboration for the A+ upgrade.

From [64] and [61] the gravitational wave signal for BHD readout now comes from taking the *difference* in the power detected by each photodiode with the sum containing a measurement purely of the shot noise. The intuitive explanation for this is that the LO no longer co-propagates with the GW signal but is instead combined with the GW sidebands on the beam splitter before the photodiodes. This causes the reflected field to acquire an additional π phase shift upon reflection from the beam splitter compared to the transmitted field. Fritschel et al. [61] show that the sum and difference of the signals on PD 1 and PD 2 are given by:

$$P_{PD1} + P_{PD2} = P_{POP} (1 + 2\epsilon) \quad (2.37a)$$

$$P_{PD1} - P_{PD2} = 2\mathcal{R}e \left(e^{i\phi} E_{POP} E_{GW}^* \right) \quad (2.37b)$$

where P_{POP} is the power in the LO from the pick-off port, ϵ represents the amplitude noise present in the electric field of the pick-off port beam and $e^{i\phi}$ now represents the phase between the LO field and GW field. In DC readout the two fields co-propagated inside the interferometer so this term was simply one.

Another point to note is that although we no longer have an intentional DARM offset in the arms, there would still be DC light incident on the photodiodes. This light comes from:

- Imbalance in the power in the interferometer arms due to technical limitations in matching arm properties e.g. slightly different reflectivity and losses in the optics or different scattering levels.
- The pickoff carrier serving as the LO, P_{POP} .

The DC power from these sources is estimated to be around 5 mW based on the design document [65].

Before we move on to discuss the noise sources within the LIGO experiment it is worth remembering that the transition to BHD Readout will not occur until the latter half of the O4 observing run (as of writing O4 is scheduled to begin in March 2023). This means that any upgrade to the LIGO readout should be compatible with not only DC readout but also BHD readout such that additional changes are not required when the transition to BHD eventually occurs.

2.1.5 Noise in the LIGO Experiment

Of course any discussion on the signal extraction from a gravitational wave detector would not be complete without also discussing the various noise sources that sets the limit on the detectors' sensitivity. These limits are constantly being revised and reduced with each upgrade to the detector, with some previously unconsidered noise sources becoming limiting and some noise sources once considered problematic becoming a minor inconvenience. To demonstrate this gradual improvement and development one is directed to Rai Weiss original design study from 1972 [66] where he outlined the noise sources he predicted to be limiting in first detecting gravitational waves via interferometry. A revised and updated version for the era of enhanced LIGO can be found in R. Adhikari's thesis [67] from 2006 and then the most up to date document outlining the noise from the O3 observing run is given in [68].

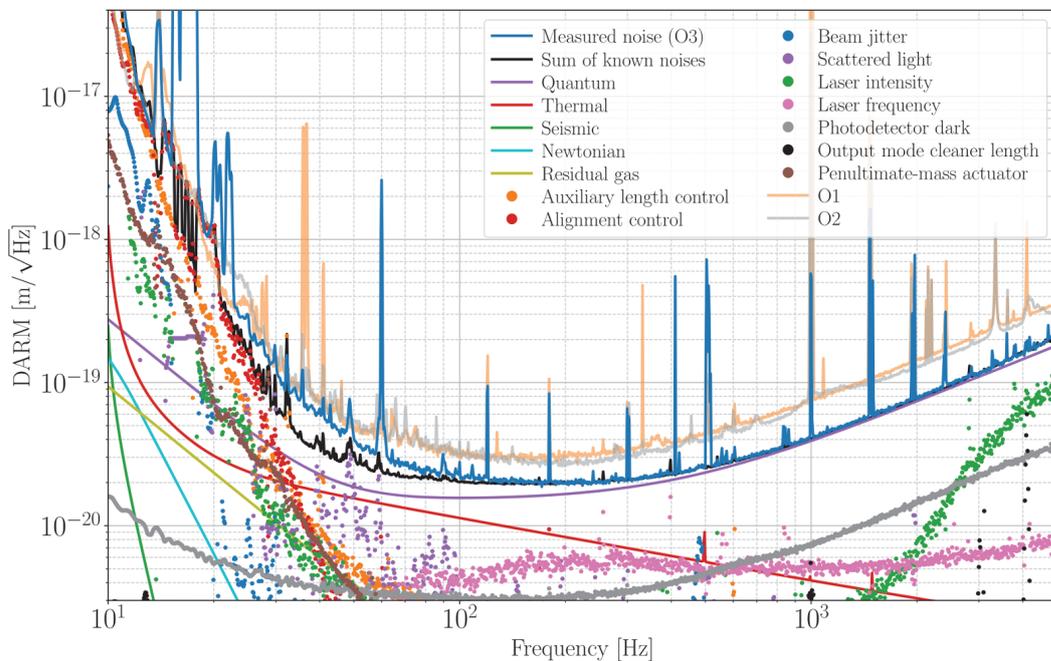


Figure 2.10: DARM sensitivity from Hanford observatory at the end of the third observing run [68] that operated with 2.0 dB of squeezing. **Blue**: the sum of all noises in the third observing run which reached sensitivities as high as $2 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}}$. Also shown in the plot is the sensitivity of the detector during observing runs, O1 (**grey solid**) and O2 (**yellow**) which operated without using squeezed light. The key contribution to the total noise is the quantum noise (**light purple**) and is specifically due to the shot noise of the photons for frequencies $> 50 \text{ Hz}$. The dark noise (**dotted grey**) is roughly a factor of 6 lower than the squeezed shot noise and is the next limiting noise source in the frequency range 600 Hz – 3 kHz.

A brief outline of the limiting noise sources is given below with the colour referring to the colour scheme in figure 2.10. The frequency band provided is only an estimate and due to the nature of noise budgets they will somewhat overlap. From low frequency to high we have:

- < 25 Hz (Red dots) **Alignment Controls:** Precise alignment of beam position is achieved using various wavefront sensors and alignment optics; however this angular alignment couples into DARM motion and must be aggressively filtered to prevent it from dominating the readout at higher frequencies.
- $25 - 50$ Hz (Purple) **Quantum Noise - Radiation Pressure:** the noise in the mirror position caused by momentum transfer from photons hitting the mirrors. Since it arises due to photons striking the mirrors in the arms it increases as the circulating power in the arms increases.
- $40 - 100$ Hz (Red) **Brownian Noise:** Brownian motion of the atoms in the mirror substrate, mirror coatings and mirror suspensions causes a displacement noise in DARM. Thermal noise increases with optical loss, and therefore ultra low loss materials are a requirement for LIGO and its future upgrades.
- > 50 Hz (Purple) **Quantum Noise - Shot Noise:** as mentioned in section 2.1.4, the shot noise in the photodetector readout electronics due to the Poissonian distribution in the arrival time of the photons on the detector. It scales with the square root of the power incident on the photodiodes. Although the shot noise in figure 2.10 seems to have some frequency dependence the shot noise as seen by the photodiode readout has a white spectrum and the shape of the shot noise on the DARM sensitivity curve comes from calibrating the shot noise with the optical gain of the cavity. One of the key targets for future upgrades is to reduce the shot noise via the introduction of increased squeezing levels (a target of 6 dB for the A+ upgrade).

Although not necessarily limiting overall sensitivity, it is worth giving an honorable mention to some of the other sources of detector noise:

- (Green - solid) **Seismic Noise:** Mirrors are suspended on a quadruple pendulum system to decouple ground motion from mirror motion. The dominant sources of seismic activity is at low frequency (< 10 Hz) however large motions can scale non-linearly into DARM and affect higher frequency bands.
- (Pink) **Laser Frequency Noise:** As the name implies this is due to instantaneous fluctuations in the frequency of the laser. Since an interferometer rejects noise that is common to the interferometer arms most of the frequency noise is rejected; however asymmetries in the interferometer (for example the DARM offset mentioned in section 2.1.4) can appear as a differential phase fluctuation in the mirror arms. The laser frequency noise is not limiting but has a fairly flat contribution to the noise budget from frequencies > 50 Hz.
- (Green - dots) **Laser intensity noise:** Similar to frequency noise, the common mode rejection of the interferometer does not completely decouple laser intensity noise from the output of the interferometer. Fluctuations in the in-

tensity of light at the dark port can mask a gravitational wave signal. - This noise is not limiting and only becomes significant at low (< 40 Hz) and high (> 1 kHz) frequencies

The final noise source to be considered - and the focus of the first half of this thesis - is the photodetector **dark noise** (grey dots in figure 2.10). At first glance it may seem that we need not concern our self with the dark noise, indeed it is so low that at points it barely features in the plot! However, in the field of precision interferometry we should never underestimate the impact of a noise source, even when it resides roughly a factor of 6 below the shot noise level.

2.1.6 The Dark Noise Dilemma

The key improvement in both LIGO detectors going from O2 to O3 was using higher laser power and the implementation of squeezed states of light in order to reduce the shot noise at the readout. During O3 the observed squeezing levels were 2.0 dB and 2.7 dB for Hanford and Livingston respectively, however there are plans to increase the squeezing levels to 6 dB during O4 as part of the A+ upgrade and potentially to 10 dB for O5 and beyond. As the quantum noise levels get reduced however, the contribution that the photodetector dark noise has on the overall noise floor becomes increasingly severe.

For this discussion we will refer to the 1 kHz region from figure 2.10 since the dark noise is the second most limiting noise source between 600 Hz to 3 kHz. For simplicity we will also neglect all other noise sources contributing to the overall sensitivity of LIGO. Currently, dark noise levels are about a factor of 6 below the *squeezed* shot noise level which corresponds to roughly a factor 8 below the corresponding *un-squeezed* shot noise level (assuming 2 dB of squeezing and at 1 kHz frequency). To some experimental physicists it may seem that this is an acceptable level for the dark noise, indeed for many experiments it would be. However as mentioned in section 2.1.4 it is considered a good rule of thumb to try to keep controllable noise sources at least a factor of 10 below the shot noise level. Clearly the dark noise at LIGO does not operate with this factor of 10 margin and the margin gets even worse as the shot noise levels continue to get reduced due to squeezing as shown in figure 2.11.

The effect that this has on the overall DARM sensitivity is shown in figure 2.12. Clearly increasing squeezing provides significant benefit to the DARM sensitivity, however we begin to see diminishing returns as the dark noise becomes increasingly limiting at 1 kHz.

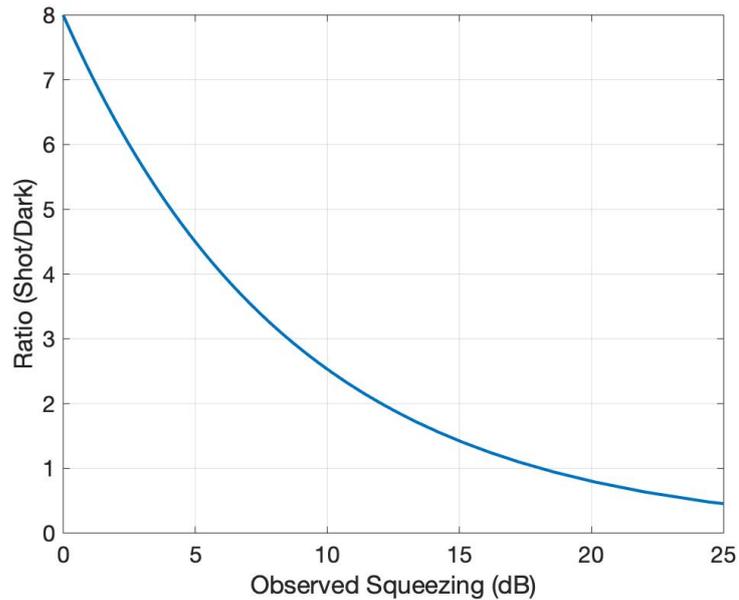


Figure 2.11: The ratio of shot noise to dark noise at 1 kHz against the observed squeezing level. An initial dark noise level was taken to be a factor 8 below the un-squeezed shot noise level.

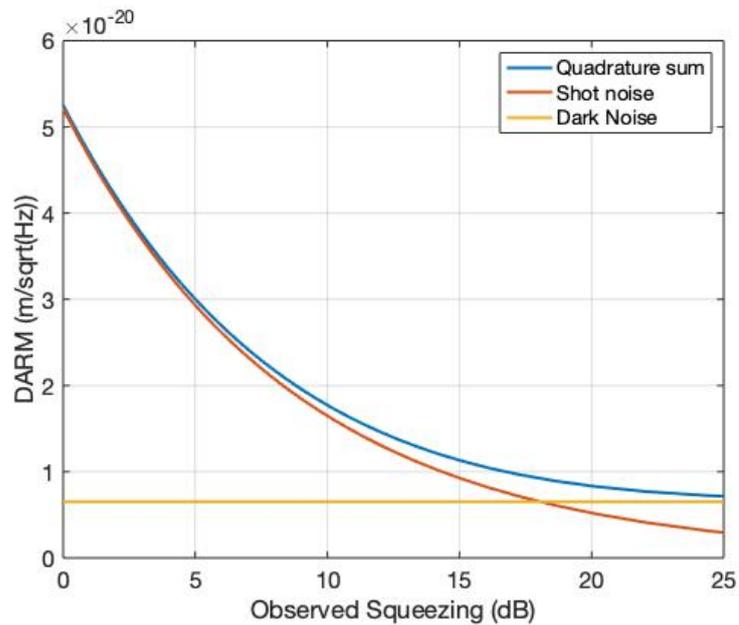


Figure 2.12: Projection of DARM sensitivity at 1 kHz. **Yellow**: Dark noise contribution of the old DCPD readout. **Orange**: shot Noise level as squeezing is increased. **Blue**: sum in quadrature of the two noise sources. This plot ignores contributions to the noise floor from all other sources.

Figures 2.11 and 2.12 surmise the problems that would arise if the readout electronics were kept the same whilst further levels of squeezing are implemented at LIGO.

Figure 2.13 shows the benefit to the LIGO detector if the readout's dark noise were to be reduced in parallel to increasing squeezing levels.

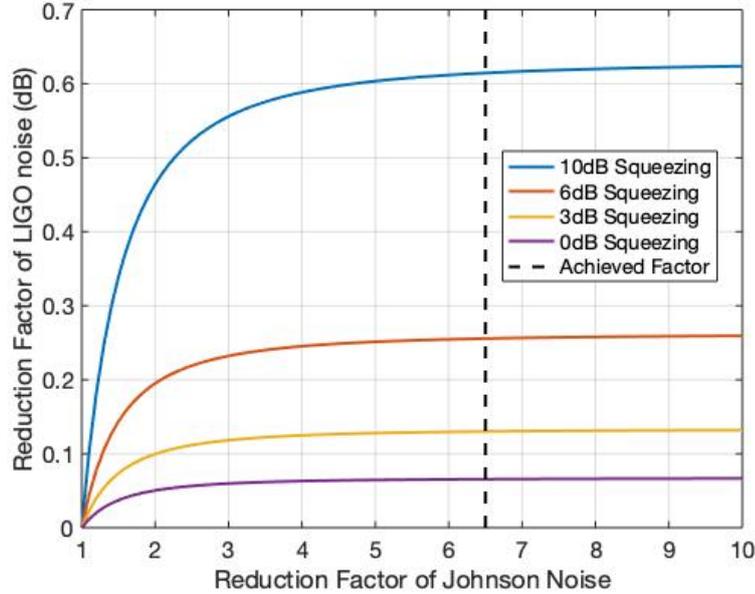


Figure 2.13: The improvement to the overall LIGO noise as the dark noise in the readout is reduced. In this plot only the relative noise levels seen at 1 kHz in figure 2.10 were considered. It was also assumed that the dark noise and shot noise sum in quadrature and all other noise sources were ignored. It should be noted that the reduction on the y-axis is on top of the reduction due to increased squeezing levels.

In figure 2.13 it is assumed that the readout electronics are initially a factor 8 below the *un-squeezed* shot noise level. The dark noise and shot noise sum in quadrature to give the total noise level. This value is then normalised by the new noise in quadrature as the dark noise is reduced by a factor F . This gives the reduction factor of the LIGO noise, A :

$$A = \frac{\sqrt{I_s^2 + \left(\frac{1}{N}I_{s0}\right)^2}}{\sqrt{I_s^2 + \left(\frac{1}{FN}I_{s0}\right)^2}}, \quad (2.38)$$

Where I_{s0} is the shot noise level in the photocurrent in the absence of squeezing, N is the initial factor the dark noise is below the un-squeezed shot noise (assumed to be 8 in the plot) and F is the additional factor reduction in the dark noise.

The impact that this sensitivity improvement has during an observing run is significant. As shown in equation 2.7 and 2.8 the strain amplitude scales with $1/r$ so the volume of the universe that the LIGO detectors are sensitive to scales with the cube of the strain sensitivity [69]. The cube of the curves shown in figure 2.13 can thus be used to express the increase of LIGO's observable universe, shown in figure 2.14.

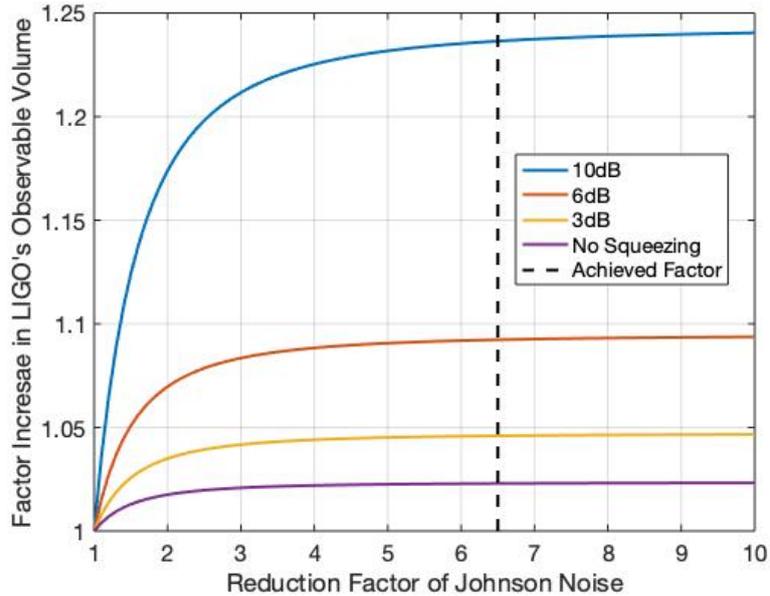


Figure 2.14: The improvement to LIGO’s observable volume of the universe as the dark noise is reduced. As in figure 2.13 the improvement on the y-axis is on top of the improvement due to implementing improved squeezing levels.

This idea can of course then be taken in many different directions. If the observable volume of the universe increases by roughly 10 % as is shown in figure 2.14 (at 1 kHz, 6 dB squeezing and a factor 6.5 reduction in the dark noise), then for a uniform distribution of candidate sources one could assume 10 % more detections at 1 kHz during a run!

Alternatively one could look at the total running cost of the LIGO detector as the price for getting a certain number of detections throughout an observing run. The more successful detections made by the LIGO detector during a run the more efficient it becomes to run the detector. Putting a number on the LIGO cost of operation is slightly complicated (should initial construction and upgrades be factored in or only operational costs? Should the salaries of the many academics around the world be considered or only workers on the LIGO sites?) however [70] cites that LIGO has received \$ 1.1 billion from the National Science Foundation (NSF) over the course of 40 years which equates to a total of \$ 75,000/day. We will be slightly more generous to try to account for other funding sources and assume a round \$ 100,000/day. This number can be used to produce an estimate on the cost per merger shown in figure 2.15.

Although it is slightly unconventional (and slightly amusing) to assign a monetary value to an astronomical event, the point remains that it costs a significant sum simply to operate the detectors during a science run and any improvement that can be made to the sensitivity of the detectors is therefore valuable. Whether one agrees

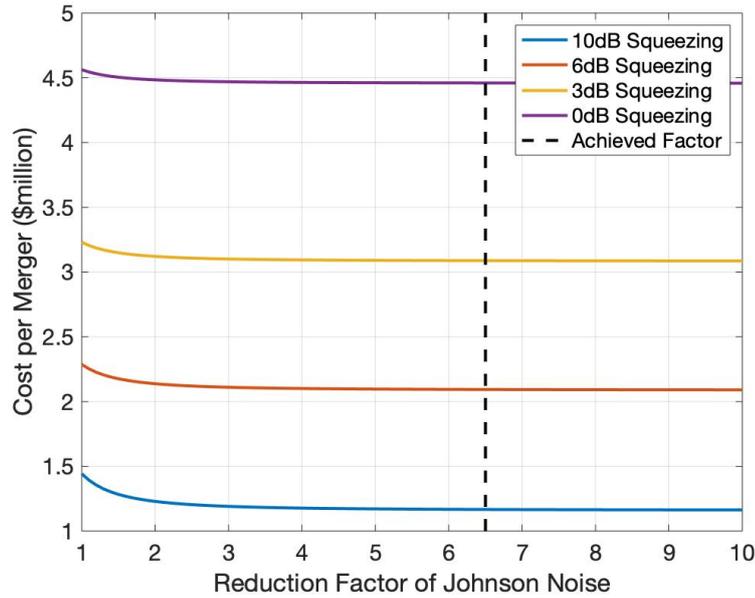


Figure 2.15: The cost per merger during a single observing run. It was assumed that on average 8 mergers per year would be made if squeezing levels were not introduced (based on the number of detections during O2) and that it costs \$100,000/day to operate both LIGO sites throughout a run. This plot has made a rather generous assumption that all detected mergers occur at 1 kHz.

with it or not, the price of receiving funding for an experiment is to produce scientific results and the scientific impact of simply making a gravitational wave detection has severely diminished since the first detection - compare the ripples (pardon the pun) the first BBH detection made in 2015 [10] with the (relatively) muted reaction to each individual BBH detection in O3. Instead the scientific motivation comes from one of two areas: detecting single mergers with unique or surprising characteristics [38] or by performing statistical analysis on a larger sample size [71], something more akin to traditional astronomy. The future of LIGO research will be in these two areas and both rely on increasing the detector sensitivity in order to detect a larger total number of events.

The next section explains the limiting factors to the dark noise of typical photodetector readouts and provides motivation for the new readout design.

2.1.7 Photodetector Model and Other Solutions

Ideal photodiodes generate an individual electron for each photon incident on their surface. Most readout systems then convert this photocurrent to a voltage by passing the current through an ideal resistor as shown in figure 2.16.

The resistor R_{trans} is sometimes referred to as the ‘load resistor’. In figure 2.16 this is true because the current is solely being driven through this one component and

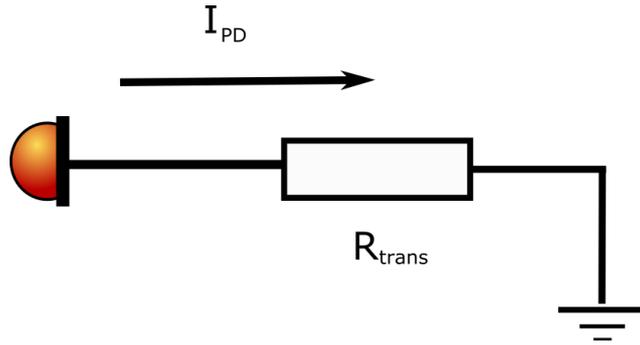


Figure 2.16: A schematic of a simple photodiode readout.

it therefore represents the ‘load’ being driven. This is not always the case however, so we will typically refer to the resistor R_{trans} as the *transimpedance* resistor due to the fact that it converts the photocurrent into a detectable voltage.

The two noise sources that we will consider in this section are the shot noise and the dark noise. In equation 2.36 it was shown that shot noise arises from the Poissonian distribution in the arrival times of the photons impinging on the photodiode which leads to a corresponding shot noise in the photocurrent. This shot noise is proportional to the square root of the photocurrent as shown in equation 2.36.

The dark noise of a photodetector readout is simply the noise of the photodetector and readout when there is no light incident upon it i.e. when it is *dark*. The noise of the readout is then the quadrature sum of the noise contributions from all electrical components, assuming that they are uncorrelated.. For the case of the readout in figure 2.16 we will ignore the $1/f$ flicker noise (experimentally this was seen to be insignificant) such that the only contribution to the dark noise is the Johnson noise of the resistor R_{trans} :

$$\tilde{I}_T = \sqrt{\frac{4K_bT}{R_{trans}}}, \quad (2.39)$$

where \tilde{I}_T is the dark noise amplitude spectral density measured within a bandwidth $\sqrt{\Delta f}$ and is in units of A/\sqrt{Hz} , K_b is the Boltzmann constant and T is the temperature of the component.

By combining equations 2.36 and 2.39 into a ratio of shot noise to dark noise, A_{SD} we can express the two noise sources by the voltage, U , across the load resistor.

$$A_{SD} = \frac{\tilde{I}_s}{\tilde{I}_T} = \sqrt{\frac{2eI_{PD}R_{trans}}{4K_bT}} = \sqrt{\frac{eU}{2K_bT}}. \quad (2.40)$$

As mentioned in section 2.1.4 it is desirable to keep the dark noise significantly lower than the shot noise, ideally at least a factor of 10. This is equivalent to achieving a

ratio, $A_{SD} > 10$.

As mentioned in section 2.1.4, for both DC readout and BHD readout there are constant DC offsets to the light incident on the photodiodes that result in a constant DC photocurrent. The magnitude of the DC photocurrent during a run with DC readout varies from 20 mA to 50 mA but we will use 20 mA for this discussion. The DC current must pass through the transimpedance resistor to convert it into a DC voltage following $U = I_{PD}R_{trans}$. The previous LIGO circuits used throughout O3 (shown in section 2.1.8) operated with a $100\ \Omega$ transimpedance resistor which resulted in a constant 2 V DC voltage across it. Using equation 2.40 this gives a ratio of the shot noise to electronic noise of 6.3 (this is also a simplified model that neglects other noise sources in the readout). To improve the ratio further one may consider increasing the impedance of the transimpedance resistor or decreasing the temperature of the electronics, however Grote et al. [69] found that these were not as straightforward as first appeared and there were practical limits to the improvements.

First let us explore the possibility of reducing the temperature of the readout in order to reduce the dark noise. In order to reach a ratio of 10 for the case of $U = 2\text{ V}$ one would need to reach temperatures below 100 K. To go further and reach the ratio of 47 that has been achieved by the new Cardiff design (see section 2.3.2) one would need to reach temperatures as low as 5 K (assuming a 2 V voltage across the transimpedance resistor), which would require the addition of cryogenics to cool the electronics readout. In the case of LIGO, the electronics are housed within a vacuum system for noise reasons and bringing in cooling systems would be extremely noisy and a major change to the current system. It is also worth considering the additional cost to run these cryogenic systems and the additional outgassing that would occur if cryogenics were introduced into the vacuum chamber.

Now let us consider increasing the impedance of the transimpedance resistor, R_{trans} . The ratio from equation 2.40 scales with the square root of the impedance value, therefore a factor 10 improvement to the ratio would require a load resistor 100 times larger. If the transimpedance resistor was increased from $100\ \Omega$ to $10\ \text{k}\Omega$ then a constant DC voltage of 200 V would be formed across the resistor! In this example the transimpedance resistor alone would need to have a power tolerance of 4 W (typical SMD resistors have a power tolerance of 1/4 W) not to mention the power tolerances of other components in the readout.

So far we have only considered the simplest photodetector readout shown in figure 2.16, but as we will see in section 2.2 such a simple readout is not practical and more complicated electronics are required that include components such as operational amplifiers (op amps). The problem behind DC power saturation then becomes more complex as all electronic components will reach a point where the voltage across

them is too large and results in a non-linear response. This introduces a limit to the improvement one can get by simply increasing the load resistance. If we say that the maximum tolerable DC voltage is 15 V then the transimpedance resistor can only be increased from $100\ \Omega$ up to $750\ \Omega$ which would improve the ratio by a factor of 2.7. For comparison the improvement over the old design using the Cardiff readout is a factor of 6.5 (see section 2.3.8). This also leaves no margin for error and in practice one does not want to operate electronics at the limit of their tolerance.

Any further improvement to the ratio would therefore require some high pass filter to reduce the fraction of the DC photocurrent that is passed through the transimpedance stage of the readout. The filter itself should also not introduce any additional noise into the electronics. The new electronics design is outlined in section 2.2 and is the successor to the previous generation electronics. A brief description of the old electronics is given in section 2.1.8 below.

2.1.8 Old Readout Design

Before outlining the new approach to reading out the photodiode signal let us first outline the old version of the DCPD readout design used throughout O3:

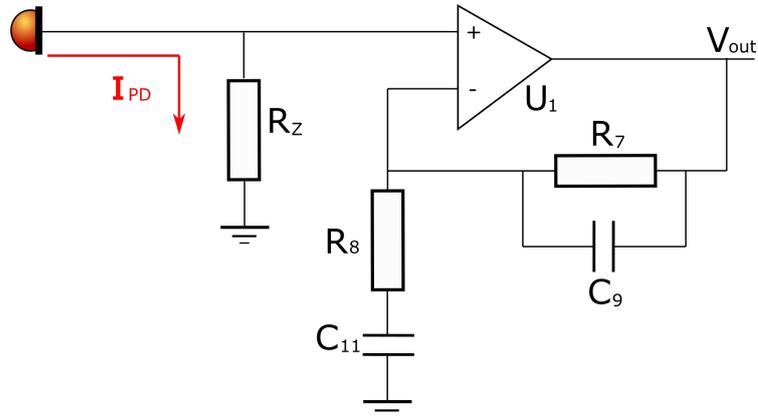


Figure 2.17: First amplification stage of the old DCPD readout. In the full schematic [72] there is a second identical amplification stage.

In this version of the DCPD readout the full photocurrent flows through the transimpedance resistor R_Z . The first amplification stage (shown in figure 2.17) has a zero at 8 Hz and poles at 80 Hz (formed by R_8 and C_{11}) and 16 kHz (formed by R_7 and C_9). The gain between 100 Hz and 10 kHz is set by the ratio of R_7 to R_8 to be a factor of 10.

The transfer function of the total readout including the second amplification stage and the differential output was modelled using the Program for Linear Simulation and Optimization of analog electronic circuits (LISO) [73] and is shown in figure 2.18:

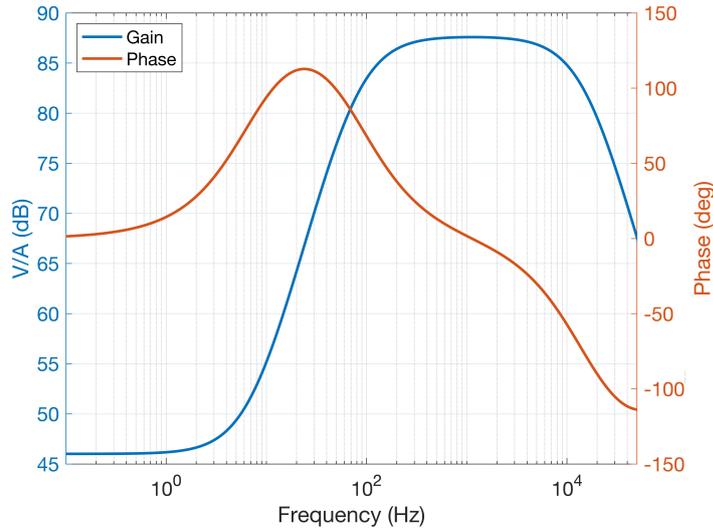


Figure 2.18: **Blue:** LISO model of the transfer function gain from the photodiode input to the differential output of the DCPD [74]. **Orange:** phase of the circuit response. The phase response is only used for accurate transfer function fits, an important requirement in order to reconstruct a gravitational wave signal.

The transimpedance resistor R_Z is $100\ \Omega$ and there is a gain of 10 in each amplification stage which brings the overall gain between 100 Hz and 10 kHz to around 80 dB in V/A. An additional gain of 6 dB comes from a differential receiver bringing the total gain to a little over 86 dB.

The voltage noise at the output is shown in figure 2.19.

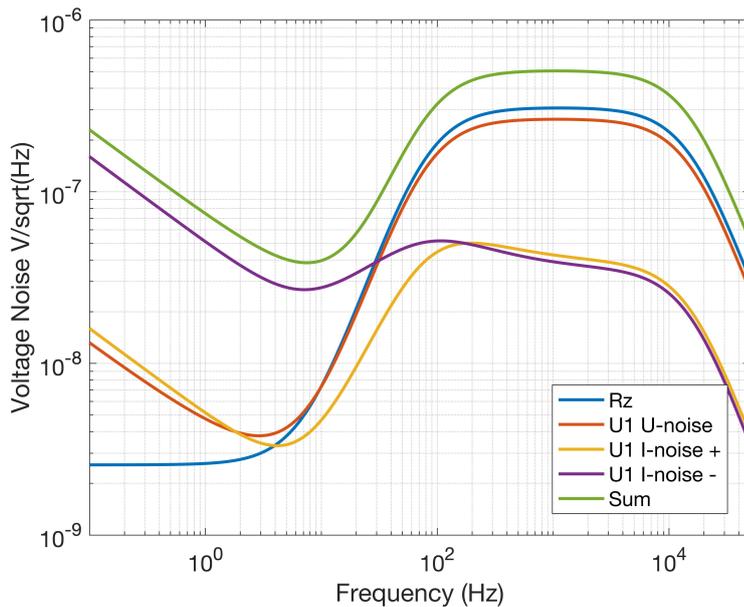


Figure 2.19: Voltage noise floor at the output of the old DCPD. U_1 refers to the op amp in the first amplification stage.

The noise floor is mainly due to the transimpedance resistor R_Z but the op amp in the first amplification stage, U_1 also contributes significantly.

The circuit noise is more intuitive when it is expressed as an *input referred current noise* (IRN) as it is easier to compare to noise in the photocurrent. To obtain the IRN of the circuit we simply divide the voltage noise in $V/\sqrt{\text{Hz}}$ by the circuit frequency response in V/A as shown in figure 2.20.

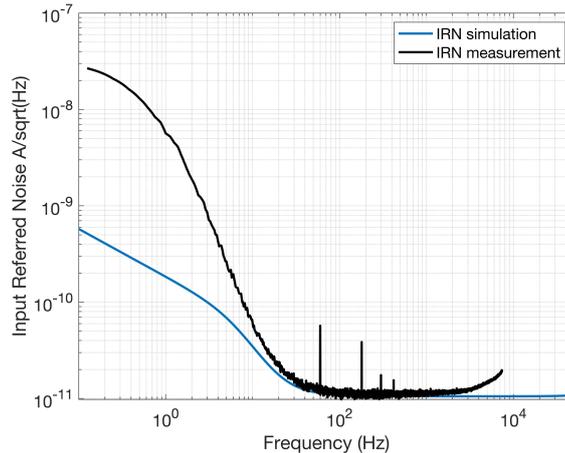


Figure 2.20: **Blue:** LISO simulation of the input referred noise. **Black:** Measurement of the input referred noise from the DCPD used throughout O3. The resonances in the measurement come from the mains frequency used in the US which is 60 Hz. Disagreement of the model and the measurement is seen at low frequencies. This discrepancy has not been investigated previously and was considered unimportant as LIGO DARM sensitivity is limited by angular control noise at low frequency. As is explored in section 2.3.8, I suspect that the difference is due to low frequency noise coming from anti-whitening filters or the CDS system.

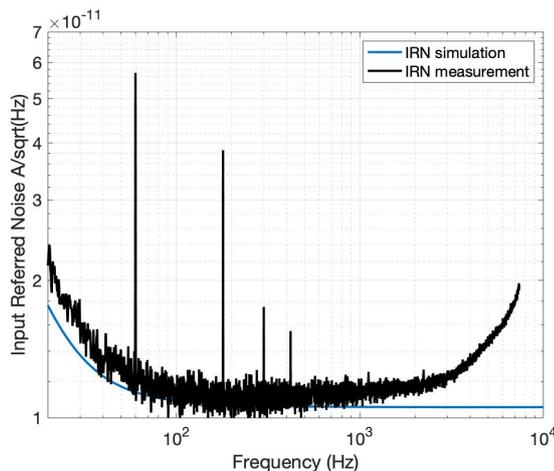


Figure 2.21: Same as figure 2.20 but zoomed to the frequency region 20 Hz to 10 kHz, the sensitive frequency band of LIGO. The deviation of the measurement from the model is not understood but is suspected to arise from the whitening stage.

Throughout the rest of this chapter we will refer back to the input referred noise in figure 2.20 and use it as the benchmark in evaluating the efficacy of the new electronics.

2.2 New Readout Design

As explained in section 2.1.7 the simplest way to reduce the dark noise of the readout is to increase the transimpedance resistance, however one runs into practical limitations due to the DC photocurrent being passed through it. A solution would be to make the transimpedance stage have a frequency dependence such that the DC photocurrent is filtered from the main readout path. The simplest of these configurations is shown in figure 2.22 where a high pass filter is formed when R_{trans} is in parallel to L and a pole is formed when $R_{trans} = \omega L$.

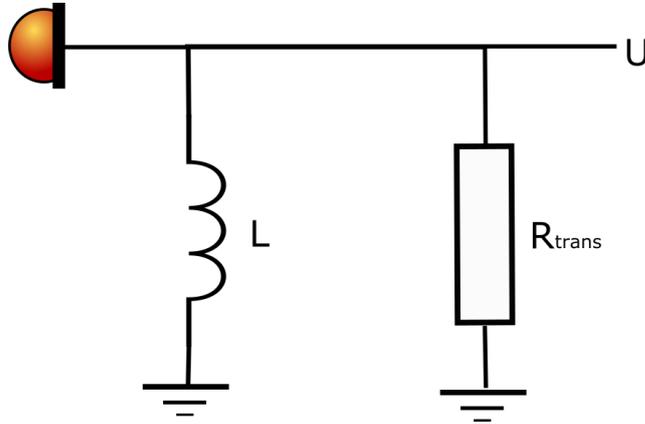


Figure 2.22: The simplest implementation of an inductor into the readout. The inductor in parallel to R_{trans} results in the voltage, U , becoming frequency dependent.

The pole formed by L in parallel to R_{trans} filters the low frequency signal and allows a larger value of R_{trans} without saturating with large DC voltages, however there are downsides to the simplicity of this setup. Firstly, a large inductance value is required in order to form a pole at sufficiently low frequency. For a 10 k Ω load resistor and a pole frequency at 20 Hz one would need an inductance of 80 H. Such an inductance value is non-trivial. Gyrator circuits can be used to replicate an inductance this large but then they will typically introduce new noise sources into the readout which can negate any improvement to the dark noise. Solid core inductors with sufficiently high magnetic permeability are a possibility; however the inductor must be carefully designed such that the core material is not saturated by the DC current passing through it (see section 2.2.3).

A further downside of the setup in figure 2.22 is that it has a very basic (first order) frequency response and there are typically requirements in the transfer function dependent on the noise features seen by the readout.

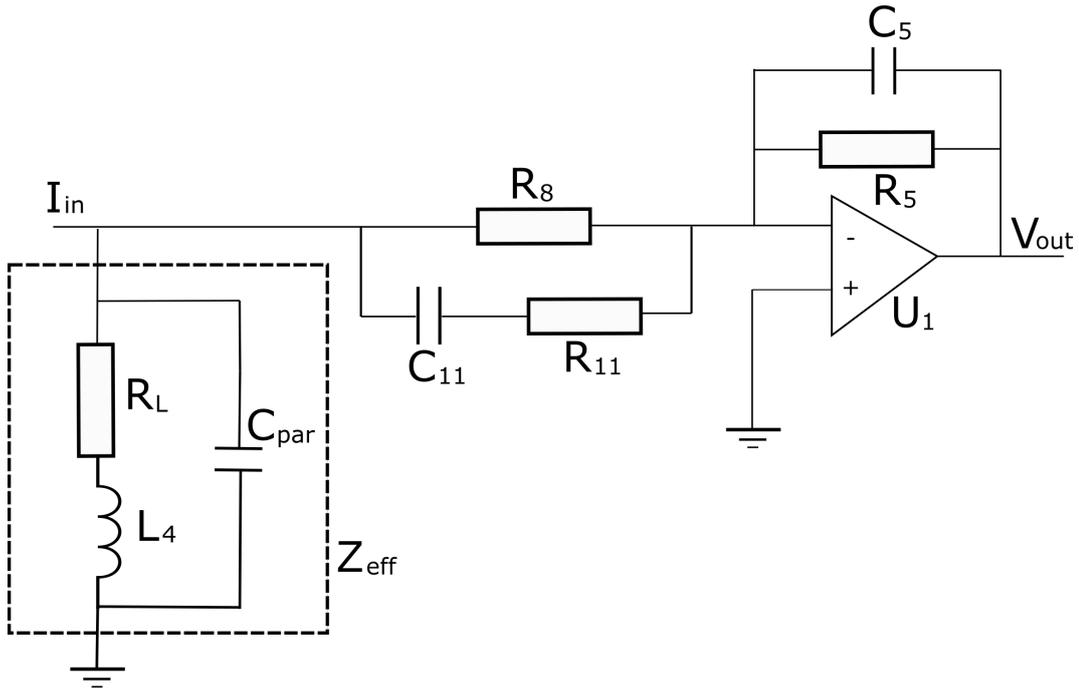


Figure 2.23: Simplified schematic of the audio path photodetector readout now implemented at the LIGO sites. Component names are the same used in the full schematic shown in figure 5.8. This design is based on the readout implemented at the GEO600 detector in 2016 [69] that successfully improved the ratio from a factor of 7 to more than a factor of 30; however it was not optimised for the frequency region of the LIGO detector. The dashed box represents a realistic model for the inductor, Z_{eff} , which is labelled L_4 in the full electronics schematic. This model is explored in section 2.3.1. R_L represents the ohmic resistance of the copper wire in the inductor L_4 . C_{par} represents the parasitic capacitance of the inductor wires.

The proposed alternative is the circuit shown in figure 2.23 which introduces an op amp with one of the pins held at virtual ground and the transimpedance resistor, R_5 , in parallel to the op amp stage. This effectively isolates the transimpedance resistor from the frequency response section of the readout and allows fine tuning of various components to alter the transfer function. The trade-off to this flexibility is that there is a new noise source introduced to the readout by the op amp. This op amp noise contribution is explored in detail in section 2.2.2 and 2.3.5.

The details of the frequency and noise response of the readout in figure 2.23 are outlined in section 2.2.1 and section 2.2.2 below.

2.2.1 Transfer Function of Audio Output

The positive op amp input pin is held at ground which then causes the op amp to drive a voltage that holds the negative input pin at virtual ground. Holding the negative pin at virtual ground isolates the transimpedance resistor from the other components. One can then choose suitable values for components, R_5 , R_8 , R_{11} , C_{11} ,

C_6 and L_4 to tune the frequency and noise response of the circuit.

The first design choice is to decide the fraction of DC photocurrent that is sent through to the transimpedance stage. It is desirable to have some fraction of the DC current being sensed by the readout for diagnostic reasons, but too much and you encounter the same saturation problems outlined in 2.1.7. The proportion of the DC signal sent through the readout is set by the ratio of the impedance of R_8 and the DC resistance of the copper wire in the inductor Z_{eff} , which is represented in figure 2.23 by $R_L = 31 \Omega$. R_8 is then set to $3.1 \text{ k}\Omega$ such that only 1% of the DC photocurrent passes through the transimpedance resistor R_5 . This in turn allows the impedance of the transimpedance resistor R_5 to be increased to 100 times its previous impedance value and thereby improve the ratio from equation 2.40 by a factor of 10.

Since the transimpedance resistor now has an additional factor 100 in the gain (in converting current to voltage) the factor 100 gain provided by the op amps in figure 2.17 is no longer needed. The maximum signal that can be read out by this first transimpedance stage is the same as the previous readout but the predicted noise is 10 times lower, i.e. the dynamic range in the first transimpedance stage has been improved by a factor of 10. Note that in practice a slightly smaller improvement is expected due to other noise sources being introduced.

Although a low current \rightarrow voltage conversion is desired at DC frequencies, a higher conversion is required in the operating frequency range of the LIGO detector between $\sim 20 \text{ Hz}$ to $\sim 10 \text{ kHz}$. To ramp up the current \rightarrow voltage conversion between DC and 20 Hz a double zero and a complex pole is formed by the inductor L_4 and capacitor C_{11} . Since the complex pole is not formed by the transimpedance resistor like in figure 2.22, it is possible to use an inductor with a lower inductance value. In the case of a $20 \mu\text{F}$ capacitance the corresponding inductor value is around 2.5 H. This is still a large inductance value and the component had to be custom made for the readout. Custom made components can be a risk compared to mass produced items because their properties are not well known from previous experiments. This inductor is also crucial in the frequency and noise response of the readout at frequencies below 100 Hz therefore a substantial amount of time was spent characterising the inductor (see section 2.2.4, 2.3.1 and 2.3.7). The complex pole can not be shifted to arbitrarily low frequencies because low-frequency noise from the interferometer would be amplified and saturate the DCPD outputs. This low-frequency noise can be seen in figure 2.12 labelled as ‘alignment control’ and it dominates the DARM sensitivity below $\sim 30 \text{ Hz}$.

R_{11} serves to damp the complex pole formed by L_4 and C_{11} . At frequencies $\gtrsim 80 \text{ Hz}$ the capacitor’s impedance is $Z_C < R_{11}$ so the readout’s gain flattens out.

Capacitor C_5 acts to form a low-pass filter with the resistor R_5 with a pole at

10 kHz to roll off the transfer function and filter any high frequency features of the photocurrent so that they do not saturate subsequent electronics stages.

Although it is not shown in figure 2.23 for simplicity, there is an additional amplification stage after the first transimpedance stage that has another zero/pole pair at low frequency and another low pass with a pole at 10 kHz to give an overall gain of 5 between 20 Hz to 10 kHz. The reason for this amplification stage is to amplify the dark noise to a level of $100 \text{ nV}/\sqrt{\text{Hz}}$ (see section 2.2.2 for the noise model). This is a design choice based on the expected pickup that occurs when sending signals over long lines from inside the vacuum chamber to the electronics racks on the outside of the chamber. By keeping the noise levels above $100 \text{ nV}/\sqrt{\text{Hz}}$ the effect of this pickup contamination is minimised; however this comes at the cost of reducing the dynamic range by a factor of 5. Given the improvement in dynamic range in the first stage of a factor of 10 this gives the readout an overall improvement in dynamic range of a factor of 2. The second amplification stage can be seen in figure 5.8.

A LISO model of the final transfer function (in V/A) from the photodiode input pin to the single ended output of the audio readout is shown in figure 2.24. In appendix 5.9 I outline an analytical approach to the detector response that achieves the same result.

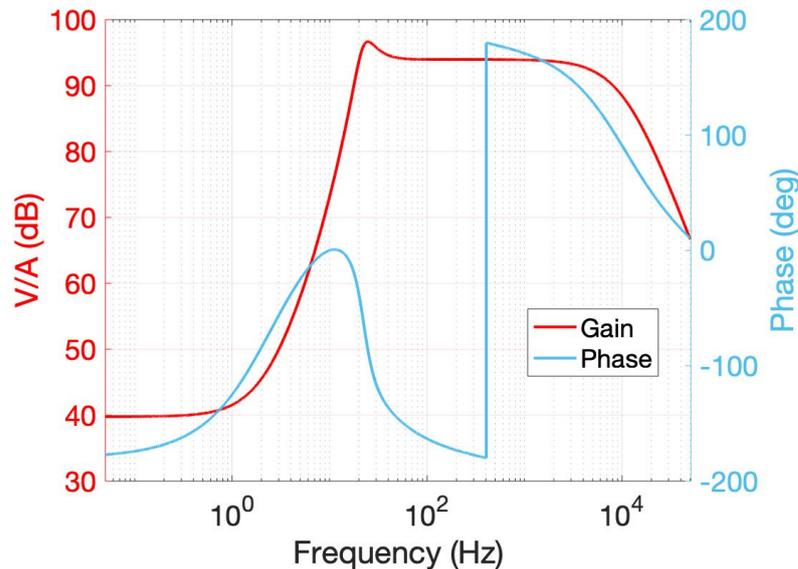


Figure 2.24: LISO [73] model of the transfer function from the photodiode input to a single-ended audio output of the DCPD. At DC the fraction of the photocurrent through the audio path is 1%. A complex zero/pole pair ramps up the transfer function between DC and 20 Hz. The transfer function is roughly flat between 80 Hz to 8 kHz. Two single low-pass filters with a pole frequency at 10 kHz rolls off the transfer function to prevent any high frequency signal content from saturating the subsequent amplification stages. Note that, for LIGO the output will be received differentially so the effective gain will be 6 dB higher, giving an overall gain of 100 dB.

Other outputs in the readout include an RF signal output (see section 2.3.4) and an output to monitor the bias voltage noise during a run (see section 2.3.3).

Figure 5.8 in appendix 5.7 shows the full schematic and PCB layout. Also shown in figure 5.8 is a relay used to input a signal to the circuit across a $100\text{ k}\Omega$ resistor at the PD input to the circuit. This resistor has a significantly higher impedance than the rest of the circuit and approximates a current source. This relay therefore allows remote calibration of the audio path using signal injection over the LIGO CDS system.

2.2.2 Noise Model of Audio Output

The dominant noise sources in the audio output are shown in figure 2.25

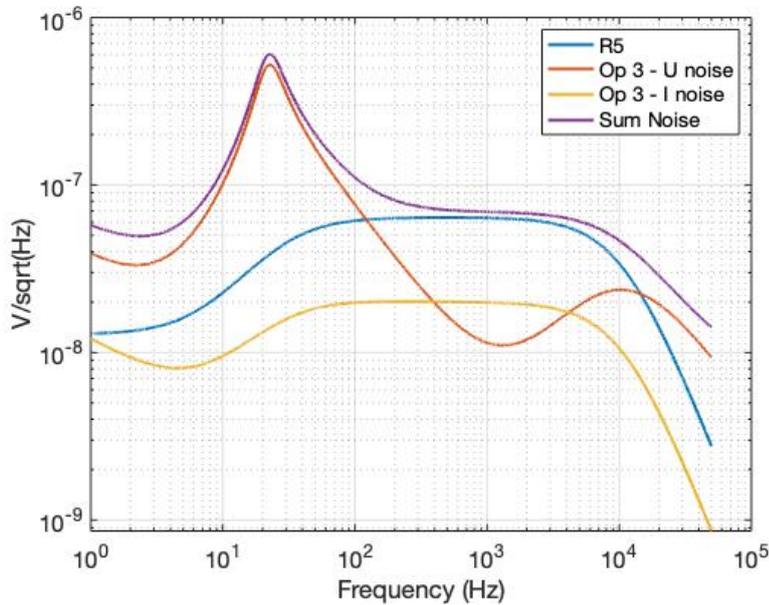


Figure 2.25: LISO [73] model of the audio output noise. Only the key components contributing to the noise are shown, however all noise sources are accounted for in the sum. **Blue**: the Johnson noise of the $10\text{ k}\Omega$ transimpedance resistor. **Orange**: the voltage noise of the transimpedance op amp. **Yellow**: the current noise of the transimpedance op amp. **Purple**: the sum of all noise sources. Note that Op 3 refers to the op amp in the first transimpedance stage (see figure 5.8), which in this LISO model was an OPA2210. Op amp noise is discussed further in 2.3.5 and op amp specifications are shown in table 2.2.

Figure 2.25 shows a LISO model of the noise at a single output of the audio path on the DCPD. Below 100 Hz the noise is dominated by the voltage noise of the op amp in the first transimpedance stage, labelled in the schematic as U_3 . Over 100 Hz the noise is dominated by the current converting resistor, R_5 .

Figure 2.25 may be misleading in some ways as it could appear at first glance that one could decrease the impedance of R_5 to reduce the noise output in the audio

path. It is worth remembering that what we are interested in is the *input referred* noise and not the voltage noise specifically.

This is shown by dividing the voltage noise at the output in figure 2.25 by the transfer function of the audio path in figure 2.24 such that the circuit noise is now using the input photocurrent as a reference level. To further illustrate this point the IRN of a readout with a 10 k Ω and a 1 k Ω resistor for R_5 is shown in figure 2.26.

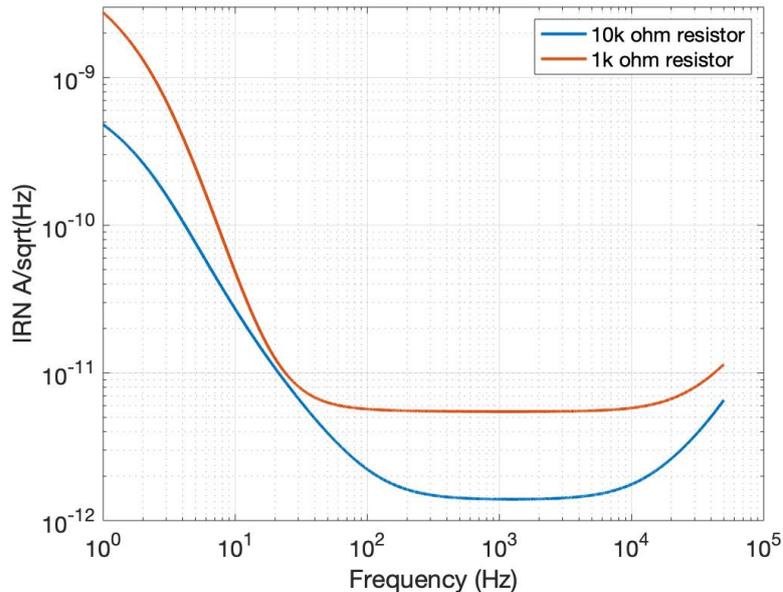


Figure 2.26: **Blue:** LISO simulation of the IRN of a readout with a 10 k Ω resistor. **Orange:** LISO simulation of the IRN of a readout with a 1 k Ω resistor. Despite the 1 k Ω resistor having a lower Johnson voltage noise it results in worse sensitivity in the readout when referred to the input.

A key role of the inductor L_4 is not only to form the zero/pole to increase the current \rightarrow voltage conversion above DC but also to reduce the IRN of the readout at frequencies < 200 Hz. If we have a low inductance value for L_4 there would be a lower impedance path to ground at low frequencies which would amplify the voltage noise of the op amp. This is actually measured later in section 2.3.8 - figure 2.75, where the IRN of one of the readout channels at LIGO Livingston is worse at low frequencies as a result of a mechanical shift in the inductor's air gap during transport.

2.2.3 Inductor Theory

A key component in the electronics design outlined in section 2.2.1 is an inductor with a mu-metal core. Some basic understanding of inductors can be found in [75]; however this section focuses on the role that the mu-metal core has on the inductance value.

The core material in an inductor serves to increase the amplitude of the magnetic field and hence the inductance. In the case of mu-metal the material is composed of a crystalline-like structure of magnetic domains. All atoms within a magnetic domain have their magnetic moments aligned in the same direction. When the core material is in an unmagnetised state the magnetic domains are randomly orientated with respect to each other and will cancel each other out. When an external field is applied the domains will re-orientate themselves to align with the field. As the domains suddenly change size and re-orientate themselves within the material there will be a sudden change in the total magnetisation of the material which results in electrical pulses from the inductor, referred to as Barkhausen noise [76].

If the applied magnetic field (induced by currents through the wire coil or by external fields) continues to increase, eventually all magnetic domains will be aligned in the same orientation as the applied field. At this point the inductor core is completely saturated and the coil will behave as if it is an air-core coil. Materials with a high magnetic permeability can be used as magnetic shields for the inductor to prevent saturation from ambient magnetic fields. The magnetic shield surrounds the core and redirects any external magnetic field around the magnetic path length of the shield rather than passing through the inductor core. The magnetic shield in this experiment is built of the same mu-metal as the core material and is shown in figure 5.11.

The magnetic shield will of course do nothing to attenuate magnetic shields induced by current passing through the inductor. In this case an air gap can be introduced in the core in order to decrease the overall permeability of the material and prevent core saturation. Due to the higher reluctance of air [77] the magnetic flux spreads out as it crosses the gap in a process known as ‘flux fringing’. The permeability of a core with an air gap is:

$$\mu_e = \frac{\mu_i}{1 + \mu_i \frac{G}{l}}, \quad (2.41)$$

where μ_i is the magnetic permeability of the core material without the air gap, G is the gap width, l is the magnetic path length of the core and μ_e is the effective magnetic permeability with the air gap introduced.

The effective permeability can then be used in the standard inductance equation to calculate the overall inductance:

$$L = \frac{\mu_e \mu_0 N^2 A}{l}, \quad (2.42)$$

where L is the inductance, N is the number of turns on the coil and A is the cross-sectional area.

2.2.4 Inductor Design

The inductors were procured by the UK based company Sowter. The inductor follows the $E-I$ core design [77] where E-shaped sheets make up the bulk of the core and I-shaped sheets make up the lid; see figures 2.27, 2.28 and 2.29. These sheets are packed together to form the inductor core. The wire is wrapped around the middle section of the ‘E’ sheets and sealed with heat shrink, then the I-shaped lid can be attached and the whole structure held in place using a metal frame.

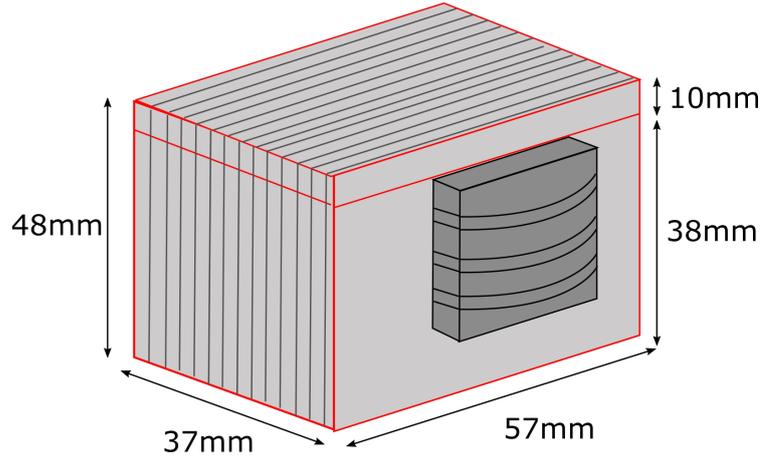


Figure 2.27: Dimensions of the inductor once assembled. Not shown is the metal casing and the bolts holding the sheets together. The dark grey box indicates the wire windings around the centre of the E-sheets.

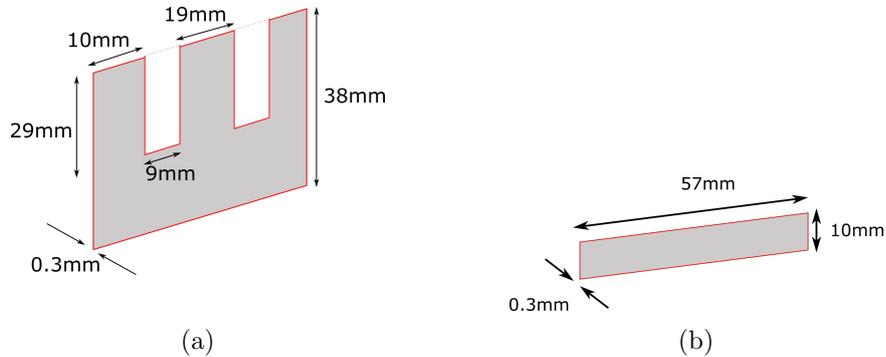


Figure 2.28: **a)** Sketch of the E sheets. The wire is wrapped around the centre prong of the E-sheets. **b)** Dimensions of the inductor I-sheets.

The magnetic permeability of mu-metal depends on the conditions used for each batch, however manufacturer specifications state a magnetic permeability of $\mu_i = 11,000$ for this batch of inductors.

The magnetic path length of the core is the total distance the magnetic flux must flow through which on average is $2 \times 57 \text{ mm} + 2 \times 48 \text{ mm} = 204 \text{ mm}$. The cross-sectional area of the core is taken as the area encased by the wire windings which is approximately $19 \times 15 \text{ mm}^2$. There are 1000 wire windings around the core and an

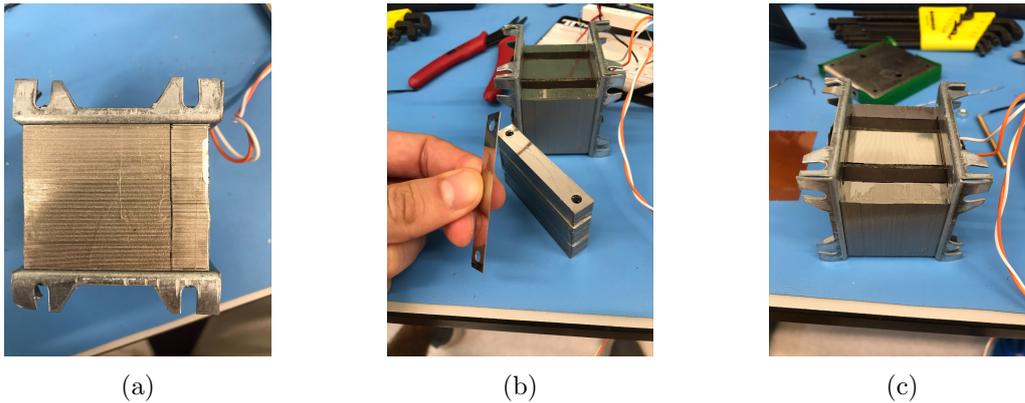


Figure 2.29: **a)** side on view of inductor showing the small air gap. **b)** I-sheets making up the ‘lid’ of the inductor core. **c)** View of the E-sheets making up the rest of the core. Plastic shims can be placed in this space to increase the air gap in the core.

air gap of roughly 0.1 mm. Using these numbers in equation 2.41 we find that the initial magnetic permeability of 11,000 H/m drops to 1,720 H/m due to the inclusion of the air gap. This effective magnetic permeability can be used in equation 2.42 which gives an inductance value of 3.0 H.

The inductance value is strongly dependent on the air gap. For the scenario where the air gap increases from 0.1 mm to 0.3 mm the inductance is predicted to drop from 3 H to 1.13 H, a drop of factor 2.7. For small variations, the inductance is approximately inversely proportional to the air gap.

The air gap of the inductors can be partially adjusted by placing shims between the two layers of sheets shown in figure 2.29, however it is hard to make large changes to the gap, since they are secured via a bolt passing through all the sheets. A larger air gap would require a redesign of the system used to secure the inductor together. It is also hard to make precise changes to the air gap as the gaps of individual sheets vary slightly with respect to each other, making quantitative estimates on the air gap fairly difficult.

A preliminary investigation into air gap variation was carried out and found that replacing the 0.1 mm shim with a ~ 0.3 mm shim caused the air gap to increase and the inductance to drop by $\sim 30\%$. This change is smaller than the 67% decrease predicted above, but the discrepancy is likely due to the aforementioned uncertainty in measuring the air gap width. Applying force to the inductor core whilst the bolt was still secured also decreased the air gap and caused an increase in the inductance of $\sim 10\%$ showing that it is possible to change the air gap without removing the bolt and changing the plastic shims.

Future upgrades of this readout could investigate implementing a system of easily adjusting the air gap to allow precise selection of the inductance value. A priority

should also be to secure the air gap in a better way to avoid any unwanted changes to the inductance during transportation. Presently the metal sheets are held in place using bolts at the edges which may allow the sheets to bend in the middle. A new support could be designed that supports the sheets all the way along their length to prevent any potential bending. The details of how the air gap affects the circuit response is covered in section 2.3.8.

2.2.5 RF Readout Design

The RF output is used to monitor the 3 MHz sidebands that are used to control the phase of the squeezer. From [62], a $54 \mu\text{A}$ signal at 3 MHz is expected in the photocurrent. This photocurrent must be converted into a detectable voltage that is significantly higher than the electronic noise of the RF readout. Trying to read out these signals in the audio readout would be a difficult task due to gain-bandwidth limitations of the op amps in the audio path and saturation of the op amp outputs. Instead the RF signals are picked off from the main photocurrent using a current divider with a cross-over frequency at 200 kHz.

The cross-over frequency to the RF path must be set high enough that there is no significant loss of signal in the audio band and low enough that there is not saturation of the audio amplifier stages by high frequency signal content in the photocurrent or significant loss of the sideband signal in the RF readout. This cross-over is dependent on all the components shown in figure 2.30 (except the transimpedance resistor) but is mostly determined by the inductor L_2 in parallel to the capacitor C_1 . A cross-over frequency at 200 kHz between the audio and RF outputs was chosen by selecting $L_2 = 82 \mu\text{H}$ and $C_1 = 4.7 \text{ nF}$. The impedance of the audio and RF paths shown below in figure 2.31 takes into account all other components.

Above 200 kHz the majority of the photocurrent is sent to the RF readout. The capacitance of the photodiode will reduce the photocurrent that is sent to the RF readout over approximately 1 MHz, however there is still approximately 80% of the photocurrent passing through the RF amplifier stage at 3 MHz.

To convert the current into a voltage, a single 50Ω transimpedance resistor is used (it actually comes from the 50Ω input to the RF amplifier), which results in $0.8 \times 54 \mu\text{A} \times 50 \Omega = 2.2 \text{ mV}$ signal at 3 MHz. An RF amplification stage (MAR-6SM) is used that provides $\sim 22 \text{ dB}$ of amplification that brings the output signal to 28 mV.

The noise of this setup is dominated by the 50Ω resistor input to the amplifier and the slightly smaller noise figure of the MAR-6SM amplifier. The noise figure of the amplifier is defined by the ratio of the input and output signal to noise ratios:

$$F = \frac{SNR_i}{SNR_o}, \quad (2.43)$$

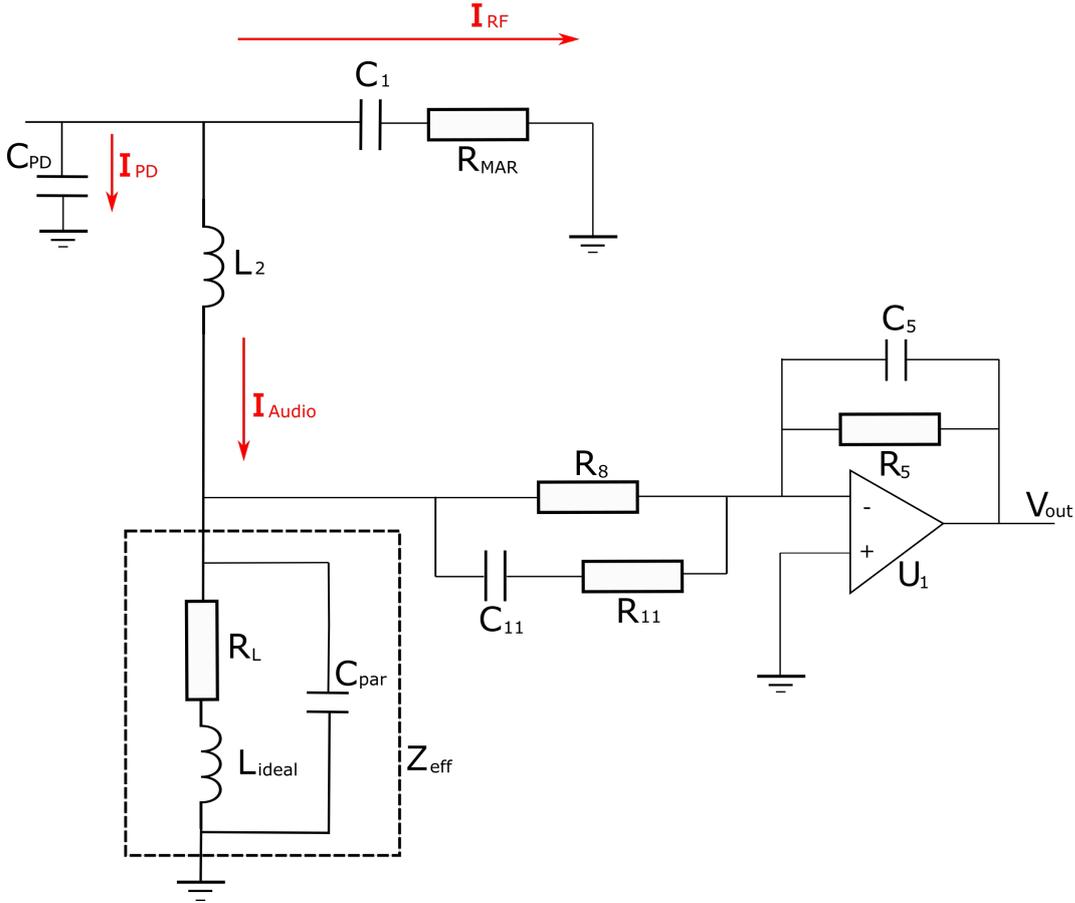


Figure 2.30: A section of the schematic from figure 5.8 that only includes key components in establishing the cross-over frequency between the audio and RF path. C_{PD} represents the photodiode capacitance which was measured by LIGO commissioners to be 200 pF. The current through each path is labelled I_{PD} , I_{RF} and I_{Audio}

where SNR_i and SNR_o are the signal to noise ratio of the input and output to the amplifier respectively. The noise figure for the RF amplifier is given as 2.3 dB. This means that the noise contribution of the MAR-6SM will be 1.3 times less than the $50\ \Omega$ resistor which has a dark noise at the input of $0.9\ \text{nV}/\sqrt{\text{Hz}}$. The expected dark noise contribution of the MAR-6SM at the input is therefore $0.6\ \text{nV}/\sqrt{\text{Hz}}$ input referred i.e. pre-amplification. The sum in quadrature of the two noise sources gives $1.1\ \text{nV}/\sqrt{\text{Hz}}$ which is amplified by 22 dB to $14\ \text{nV}/\sqrt{\text{Hz}}$ at the RF output.

The shot noise of a 20 mA photocurrent through the $50\ \Omega$ input resistor and then 22 dB amplification would be $50\ \text{nV}/\sqrt{\text{Hz}}$ at the output of the amplifier, about a factor 4.5 above the electronic noise. The sum in quadrature of the shot noise and electronic noise would be around $52\ \text{nV}/\sqrt{\text{Hz}}$ at the RF output. For a bandwidth of 1 Hz this is a factor 5.3×10^5 below the sidebands' signal at 3 MHz (28 mV).

The maximum output of the MAR-6SM amplifier is stated as 18.1 dBm which corresponds to 65 mW. The MAR-6SM is driving $50\ \Omega$, therefore we can use $P = V^2/R$

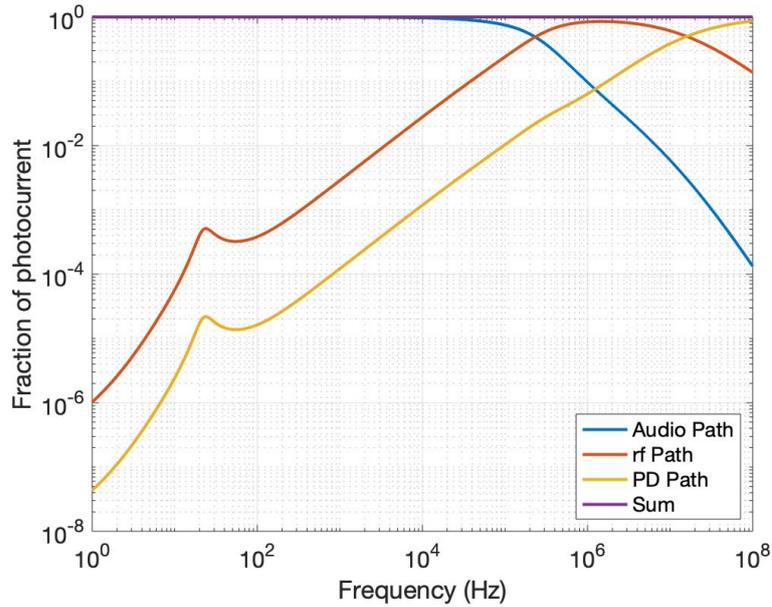


Figure 2.31: Model of the fraction of photocurrent sent to each path of the readout. A 200 pF photodiode capacity was assumed based on measurements taken by LIGO.

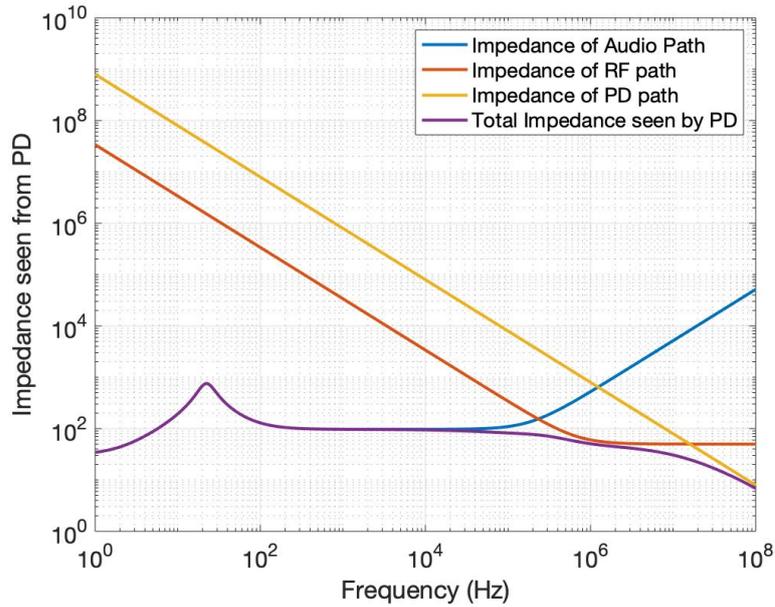


Figure 2.32: Model of the impedance seen from the photodiode. Since the negative input pin on the op amp in the first transimpedance stage is held at virtual ground there is a path to ground that does not include the transimpedance resistor or any other subsequent components. A 200 pF photodiode capacity was assumed.

to see that voltages over 1.8 V will saturate the output of the RF amplifier. Given that the 3 MHz sideband's signal is expected to be 0.028 V, it is predicted that there is sufficient headroom to prevent saturation from occurring.

2.2.6 Bias Readout Design

Photodiode theory and construction is a well documented area with many textbooks on solid state physics covering the topic in detail [78]. This section does not go into specifics on the devices but rather provides a basic explanation of the role of the bias voltage in PN photodiodes such that the importance of a bias voltage monitor is emphasised.

In almost all photodiode applications the diodes are operated with a reverse bias in order to obtain a linear relationship between light intensity and generated photocurrent, as well as a high responsivity and quantum efficiency. The Shockley equation describes the I-V relationship for a PN diode with an additional photocurrent I_P induced by photons falling on the depletion region [79]:

$$I_T = I_{sat} \left(e^{\frac{qV}{k_B T}} - 1 \right) - I_P, \quad (2.44)$$

where I_T is the total current, I_{sat} is the reverse saturation current, q is the fundamental charge, V is the applied bias voltage, k_B is the Boltzmann constant, T is the temperature and I_P is the induced photocurrent which is dependent on many factors including the incident light power and the photodiode responsivity. Generally, the induced photocurrent is independent of the bias voltage and simply shifts the overall I-V curve of the photodiode up or down as shown in figure 2.33.

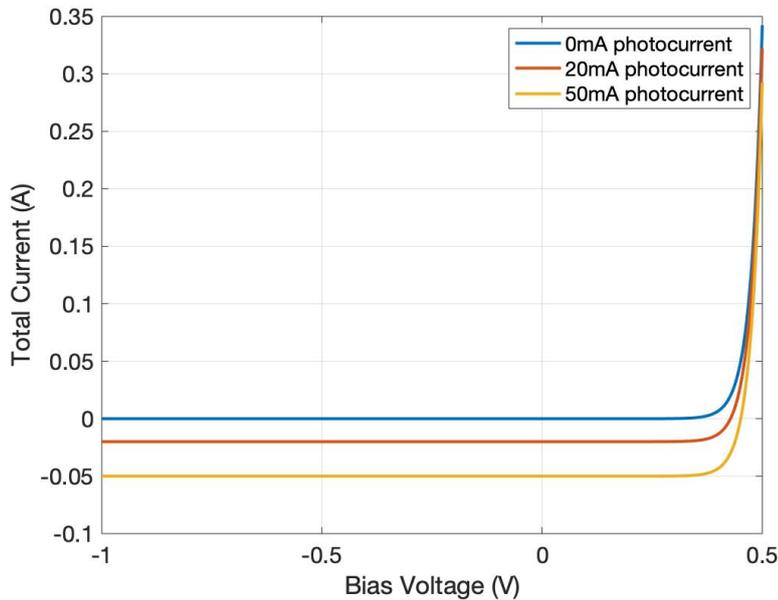


Figure 2.33: A plot of the I-V curve using equation 2.44 for a photodiode with 0mA, 20mA and 50mA of photocurrent. The reverse saturation current I_{sat} was chosen to be 1 nA in this plot.

Typically, the noise in the photocurrent (stemming from noise in the incident light

on the detector) will dominate the noise in the total current whilst the bias voltage noise will not couple significantly. Moreover, voltage regulators are frequently used in electronics to generate stable sources of voltage for devices such as op amps and bias voltage power supplies which further reduces the impact of the bias voltage noise on the total current, I_T .

For this reason, the bias voltage noise has rarely been the subject of investigation at LIGO. Work carried out by R. Abbott reduced the bias noise on the previous generation bias voltage regulators by roughly a factor 10 at frequencies around 100 Hz [80]. There was also a preliminary investigation to measure the coupling of bias noise into DARM where noise was intentionally injected into the bias voltage but there was no bias voltage monitor so no strong conclusions could be drawn. Besides that the noise on the bias voltage has not been the subject of significant investigation and there has never been a way to actively suppress or monitor the bias voltage noise during a run.

With this in mind, it was considered to implement bias noise suppression in the new DCPDs as feed-back or feed-forward on the in-vacuum circuit, however, it was decided that it was impractical to add passive filtering to the circuit. Instead, monitoring outputs were added so that bias noise coupling into DARM could be measured. If the DARM coupling is found to be significant then the bias noise can be compensated for with additional digital filtering using the LIGO CDS system.

The schematic for the bias voltage monitoring stage in the DCPD is shown in figure 2.34.

For DC signals the capacitors can be considered open in the circuit. The voltage at the monitoring output is then only determined by the fraction: $\frac{R_{24}}{R_{24}+R_{20}}$. Since the bias voltage signal content will be dominated by the signal at DC, this ratio is set to be equal to 10 such that there is a factor 10 attenuation of the signal content at DC. Since there has not been in depth investigation into the bias noise floor before, there are no clear requirements on what features were desired in the bias voltage whitening stage. It was therefore decided to have maximum gain in the sensitive frequency region of LIGO (region 30 Hz to 10 kHz) and implement a double zero/pole pair (zeros at 3.4 Hz and poles at 31 Hz) and a third pole with pole-frequency at 10 kHz. These are set by components around the amplification stage U5; zero 1 set by $R_{20} = Z_{C20}$ and pole 1 set by $R_{24} = Z_{C20}$. Zero 2 is set by $R_{16} = Z_{C21}$ and pole 2 is set by $R_{15} = Z_{C21}$. This double zero/pole combination results in the output bias monitoring voltage having a gain of 10 over ~ 30 Hz. Finally, a pole is formed at 10 kHz when $R_{16} = Z_{C20}$ which rolls off the transfer function to prevent saturation from any high frequency pick up that may occur.

A model of the circuit response is shown in figure 2.35 and the result from the bias monitoring measurement on the finished circuit is shown in section 2.3.3.

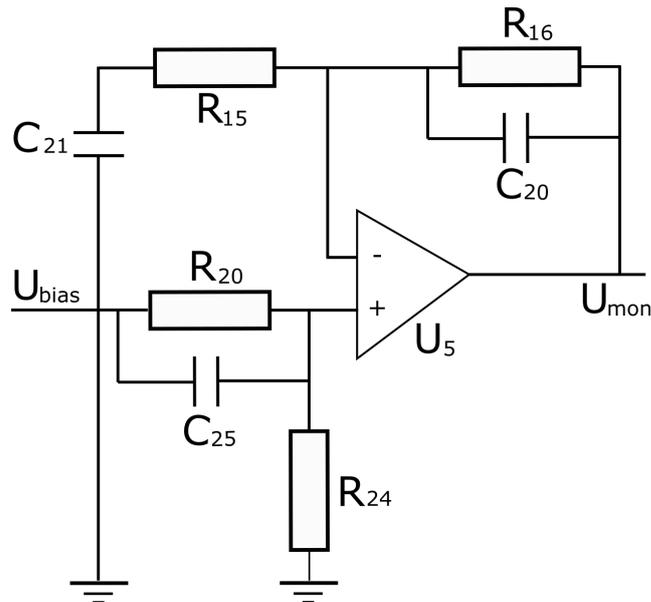


Figure 2.34: Bias monitoring scheme on the DCPD. U_{bias} is the bias voltage applied across the photodiode and U_{mon} is the voltage sent to the differential sending stages of the DCPD. Differential outputs are not included in this diagram but can be seen in the full schematic found in appendix 5.7.

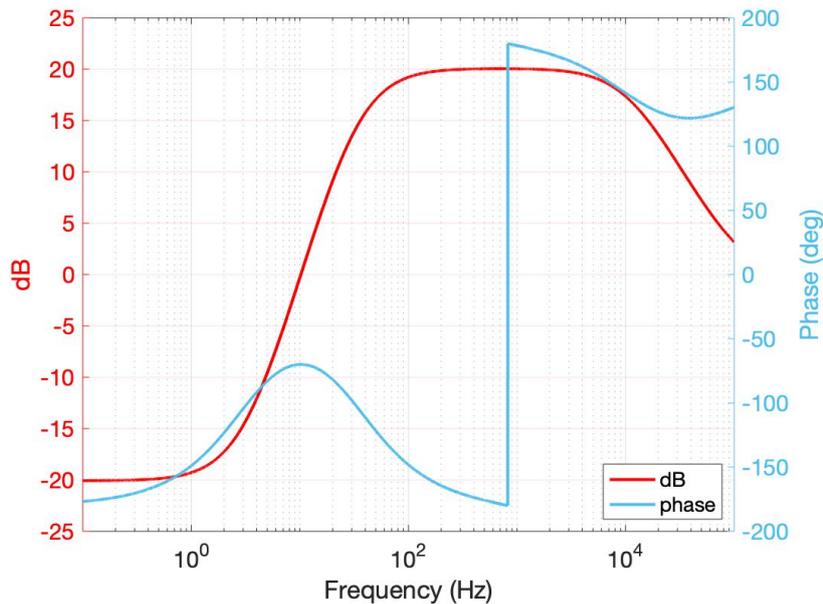


Figure 2.35: Model of the transfer function from the photodiode bias voltage to the bias monitoring port. There is 20 dB attenuation at DC and 20 dB amplification between 30 Hz to 10 kHz. Currently the phase of the transfer function is only used for generating accurate fits. If the bias monitor becomes part of an active feedback system then the phase margin of the loop may need to be reconsidered.

2.2.7 Vacuum Case and Assembly

The vacuum-compatible case is used to house a PCB with two parallel channels and the two inductors used in the audio path of the readout. The circuit needs a gas

inside the enclosure - in this case helium - in order to dissipate heat from the op amps in the circuit (conduction with the PCB is insufficient and would result in the components overheating) but since the enclosure goes in vacuum, there must be a way to achieve a vacuum seal around the enclosure. Previously the enclosure for the photodetector readouts had to be welded shut [81] before being placed into vacuum, which is a fairly restrictive design. Any changes to the electronics would require the case to be cut open and discarded and a new one to be produced. The new case uses a Helicoflex sealed gasket to maintain a vacuum seal, that can be removed if necessary, which allows the case to be reused.

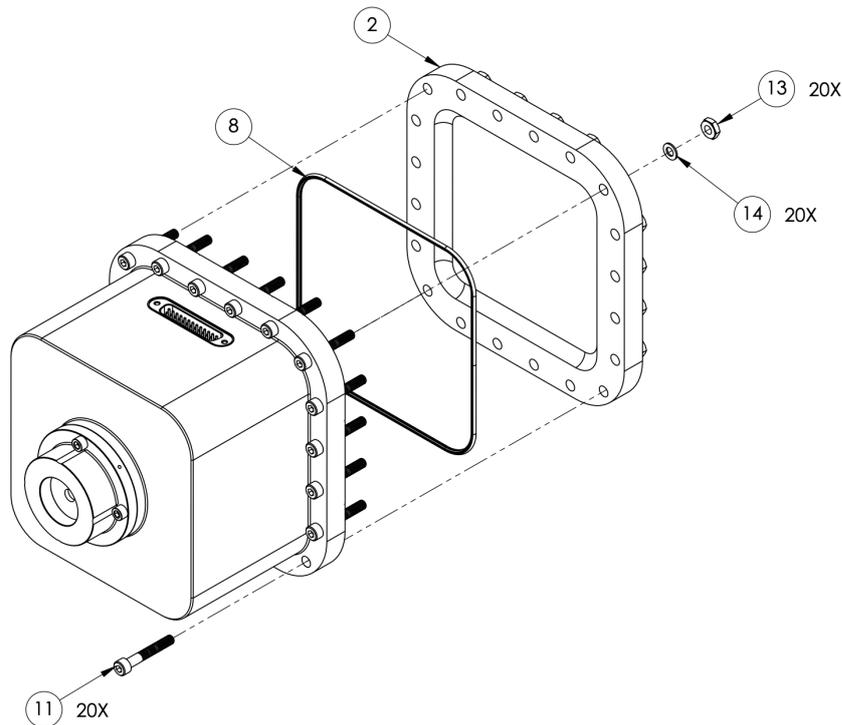


Figure 2.36: Assembly diagram of the DCPD vacuum enclosure [82]. Screws and nuts, labelled **11**) and **13**), secure the main casing to the lid, labelled **2**). The Helicoflex seal ring, labelled **8**), is secured around the rim of the lid and achieves the vacuum seal.

The size of the new case is compatible with the available space that was used by the old design in the LIGO HAM-6 chamber, so the new readout could easily be exchanged with the old version.

Two inductors are housed within the same mu-metal shield inside the case and mounted in opposite directions such that their sum magnetic pickup will cancel out. This mounting procedure is shown in appendix 5.8.

The inductor wires are then connected to the PCB (care must be taken not to reverse the polarity of the wires) as in figure 2.37 and then the PCB is secured into the stainless steel case (this test prototype is made of aluminium) with the inductor

shield secured on top.

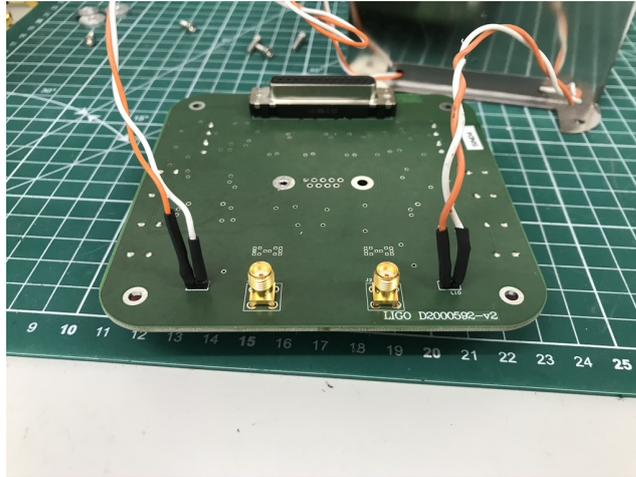


Figure 2.37: Inductor wires are soldered to the PCB as shown. Gold-plated SMA outputs provide the RF signals. The holes in the middle of the circuit are a D-sub 9 socket that provides the photodiode input and bias voltage pins. The socket at the back accepts a D-sub 25 plug that provides all other input and output pins including the necessary power supplies for the op amps. The pin layout is shown in appendix 5.7 and more photos of the PCB are found in appendix 5.4.

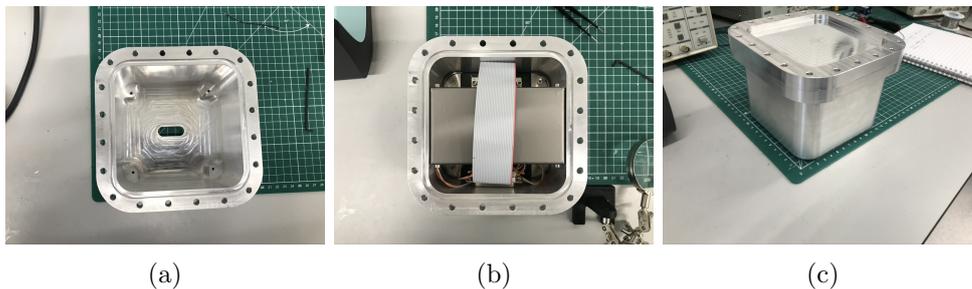


Figure 2.38: **a)** Empty vacuum case. **b)** View of the PCB and inductors mounted into the vacuum case. A ribbon cable passes over the top of the shielded inductors. **c)** DCPD vacuum case fully assembled. This is a test case used to check the dimensions of the assembly and is made of aluminium, the actual case is made of stainless steel.

A potential improvement for a future design for O5 design would be to move the readout electronics outside the vacuum chamber. This would provide benefits to heat dissipation in the readout (see section 2.2.8) and more flexibility in accessing and modifying the electronics. The downside is that the photocurrent signal from the photodiodes would have to pass through long lines (see figure 2.70 to see the PD lines) to exit the vacuum chamber and there could be significant magnetic coupling into the photocurrent. Sufficient magnetic shielding of these lines would therefore be a requirement for this potential upgrade.

If the electronics were migrated outside the vacuum chamber there is a further - more ambitious - upgrade that would then be more feasible, which is to introduce

cryogenic cooling to the readout. The benefit that cryogenics could have to the electronics is described in section 2.1.7.

2.2.8 Heat considerations

The total electrical power dissipation of the PCB (sum of both parallel channels) is $0.06 \text{ A} \times 15 \text{ V} + 0.12 \text{ A} \times 15 \text{ V} = 2.7 \text{ W}$. The asymmetry in power comes from the asymmetric power supply of the MAR-6SM amplifier. Power dissipation from the photocurrent through a single inductor is $(0.05 \text{ A})^2 \cdot 31 \Omega = 78 \text{ mW}$, which does not contribute significantly on top of the 2.7 W , even for as high as 50 mA of photocurrent.

A colleague, Aldo Ejlli [83], simulated the case in vacuum in a room temperature environment, with a thermal load of 3 W . The resulting equilibrium temperature for steel with emissivity of 0.33 is about 36°C . It should be noted however that this is a pessimistic estimate that assumes a slightly higher power dissipation and (importantly) does not take heat conduction from the case to the HAM table into account so the actual equilibrium temperature will be slightly lower.

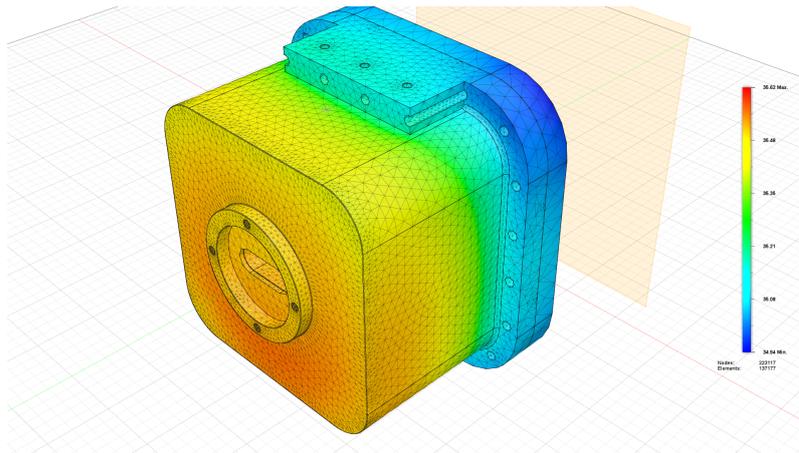


Figure 2.39: A thermal model of case with 3 W of heat load in vacuum with ambient temperature vessel walls. The color legend goes from 34.94°C (**blue**) to 36.62°C (**red**). This model was constructed by a colleague in Cardiff [83].

The effect that the temperature rise has on the readout audio response is explored in section 2.3.8.

2.3 Measurements and Results

2.3.1 Inductor Characterisation

The mu-metal core inductor, L_4 , plays a crucial role in the DCPD both in its noise response (see section 2.2.2) and in the transfer function response (see section 2.2.1). It is also a custom made component and uses a somewhat uncommon compound

(mu-metal) the properties of which vary depending on its production. Then there are factors associated with the assembly of the inductor itself that all contribute to the final properties of the component and its role within the circuit. This section concerns the characterisation of the inductor, and its final performance as part of the DCPD used in the LIGO experiment.

DC offsets

As outlined in section 2.2.3, saturation of the core material with a large magnetic field will result in a drop in the effective inductance which would cause a change in the ultimate transfer function of the DCPD. The majority of the DC photocurrent ($\sim 99\%$) passes through the inductor L_4 therefore an investigation into the DC saturation of the inductor was carried out.

A simple setup was used to generate a DC offset on a sinusoidal signal and inject it into the inductor.

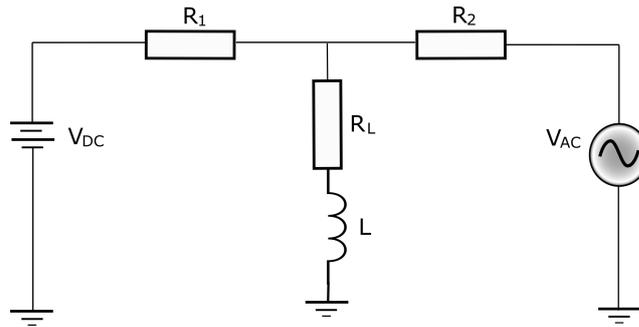


Figure 2.40: DC current was generated by applying a DC voltage across R_1 . An 80 Hz current was generated by passing V_{AC} through R_2 . R_L represents the ohmic resistance of the copper wire in the inductor and L represents the ideal inductor. For simplicity the parasitic capacitance of the inductor has been omitted from this model. This is an acceptable simplification because the parasitic capacitance is about 300 pF (see section 2.3.1) so can be approximated as an open circuit at DC and 80 Hz.

The phase, ϕ , of the AC signal across the inductor was measured for varying DC offsets and used to calculate the inductance according to:

$$L = \frac{\left(\sqrt{R_1^2 \omega^2 (R_1^2 R_2^2 - 4R_L(R_1 + R_2) \tan^2(\phi)(R_L(R_1 + R_2) + R_1 R_2))} + R_1^2 R_2 \omega \right)}{2R_1 \omega^2 (R_1 + R_2) \tan(\phi)}, \quad (2.45)$$

where ω is the angular frequency of the AC source. Equation 2.45 was derived by considering the phases of both the AC current through the circuit and the impedance of the component under test (in this case the inductor).

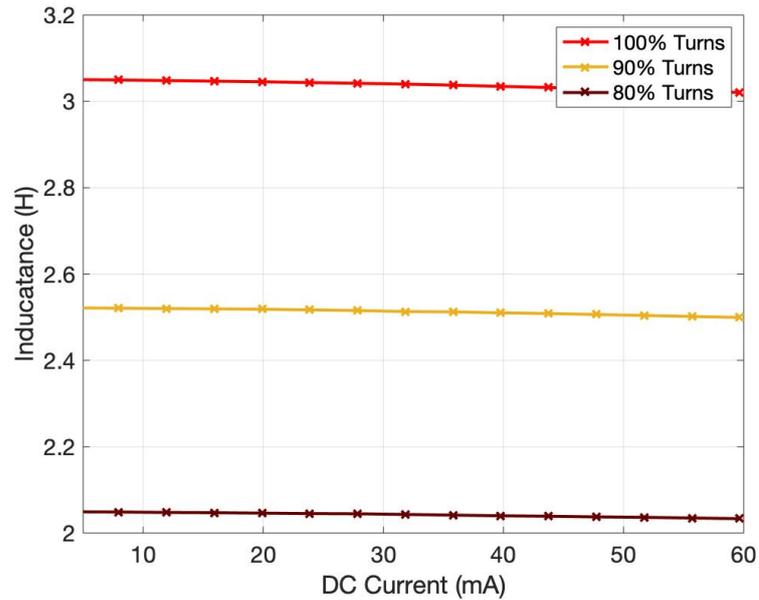


Figure 2.41: Effect of DC current on the inductance at 80 Hz at three different wire taps. These taps were only included in early inductor prototypes in case a different inductance was desired for testing. Subsequent versions did not have individual wire taps.

Figure 2.41 below shows that the inductance gradually decreases as DC current increases due to partial magnetic saturation of the mu-metal core (see section 2.2.3). The variation in inductance going from 5 mA to 60 mA is on the order of 1 – 2%. The effect that this would have on the final transfer function is small enough to not be a concern as shown in appendix 5.6.

Complete saturation of the core material results in the inductor behaving as an air-core inductor, as described in [77], but this was not observed during the test. It could potentially be a topic for further research in future versions of the readouts for O5.

Temperature Dependence

The vacuum case shown in section 2.2.7 is housed within the HAM-6 vacuum chamber at LIGO and hence only dissipates heat radiatively and via conduction through the base mounting which results in moderate heating of the components. (See Section 2.2.8 for a heat model.) To test the temperature behaviour of the inductor, it was heated with current sent through resistors mounted at the circumference of the inductor, but not connected in the same circuit as the inductor. As before, an 80 Hz signal with an increasing DC offset was used to simultaneously characterise heating and possible saturation from DC current. Results are shown in Figure 2.42.

A further effect that heating has on the inductor is that it causes a change to the

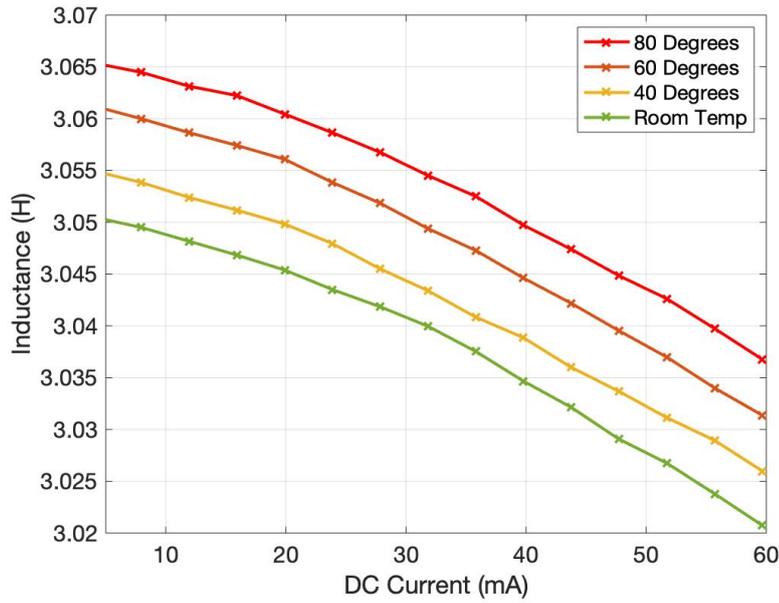


Figure 2.42: Inductance vs DC current with changing temperature. The AC signal was 80 Hz. Due to the time taken to stably heat the inductor, only the 3 H tap was investigated.

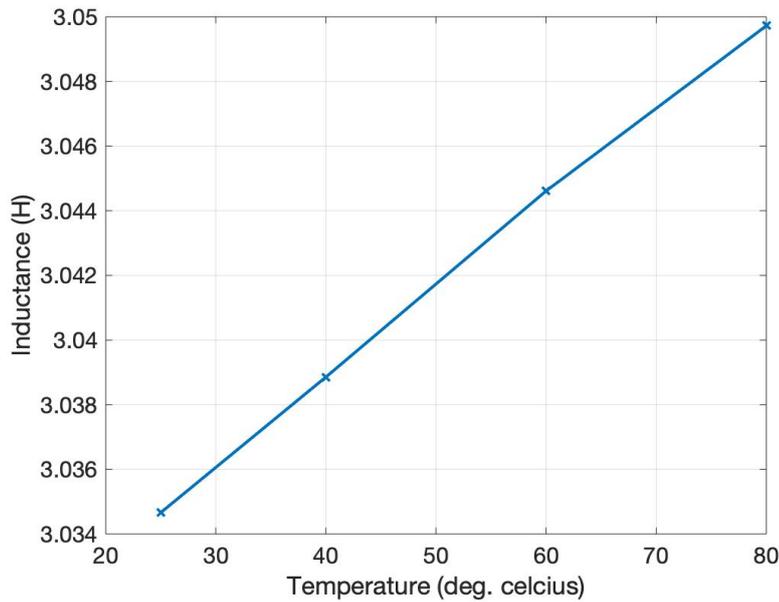


Figure 2.43: Plot of inductance vs temperature. The 4 points were taken from the 20 mA points on figure 2.42 and show an approximately linear relationship. The change in inductance going from room temperature to 80°C whilst maintaining a 20 mA DC photocurrent is on the order 1%, giving the gradient to be about 0.1% per 10°C.

ohmic resistance of the copper wire in the coil. This effect was not considered until the units had been installed at Hanford and Livingston and so there was no direct

measurement of this change in copper resistivity. Section 2.3.8 discusses the effect that a change in temperature had on the transfer function when the readouts were installed.

Parasitic Capacitance

A parasitic capacitance forms between the tightly packed wires of the inductor which results in the inductor more realistically being resembled by Z_{eff} in figure 2.44.

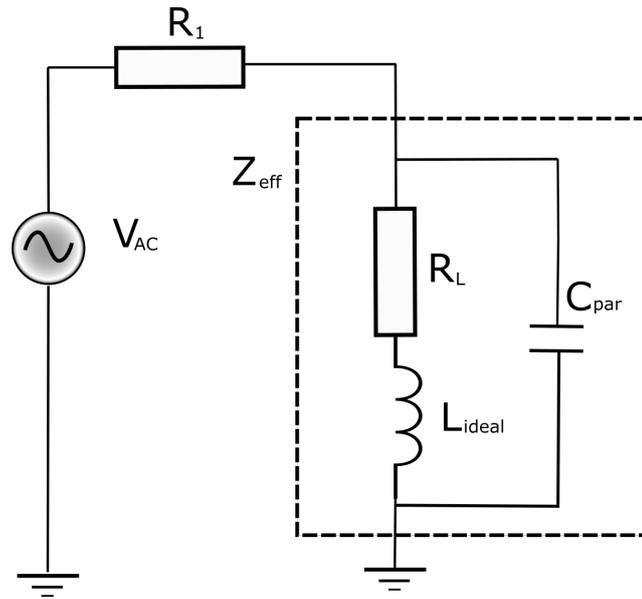


Figure 2.44: Setup used to measure the parasitic capacitance. Z_{eff} represents the large 2.5 H inductor. The copper wire resistance, R_L , is in series with an ideal inductor, L_{ideal} , whilst the parasitic capacitance C_{par} is modelled in parallel.

A swept sine measurement was performed on the inductor in series with a 1 k Ω resistor and the phase across the inductor was measured. The model of the inductor shown in figure 2.44 was compared to the measurement in figure 2.45 and used to estimate that the parasitic capacitance as $C_{par} = 300 \pm 10$ pF.

The parasitic capacitance will form a self resonance within the inductor when the parasitic capacitor impedance equals the series impedance of the resistor and inductor: $Z_{C_{par}} = R_L + Z_{L_{ideal}}$. For a parasitic capacitance of 300 pF this gives the self-resonance to be around 6 kHz.

Although this frequency is well within the sensitive frequency region of the LIGO detector, when it comes to the response of the DCPD the self resonance has almost no effect. The impedance of the ideal inductor at 10 kHz is already high compared to the rest of the circuit (and only higher at resonance), hence the resonance does not appear as a feature in the DCPD transfer function shown in figure 2.24.

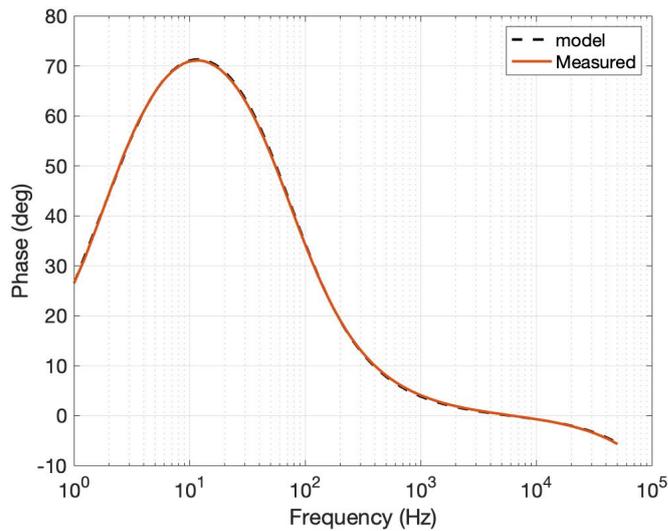


Figure 2.45: A swept sine measurement was taken across the inductor in series with a $1\text{ k}\Omega$ resistor. The ohmic resistance R_L was measured using an ohmmeter and the inductance value L_{ideal} was measured at low frequencies where the impedance of parasitic capacitance was assumed to be infinite.

Magnetic Coupling

Due to the high magnetic permeability of mu-metal ($\sim 1.1 \times 10^4$), the inductors are susceptible to magnetic pick-up, therefore a mu-metal shield was designed to house the inductors. These shields redirect the magnetic flux through the magnetic path of the shield and prevent it from entering the inductor core material.



Figure 2.46: **a)** the original inductor prototype. The brown and red wires are optional taps at 2 H and 3 H respectively which are only present on this evaluation model. **b)** prototype mu-metal case that was used for testing. The new design shown in figure 5.11 was designed for two inductors, but has the same wall thickness of 1 mm.

To test the magnitude of the magnetic coupling a known magnetic field was passed

through the inductor with a Helmholtz coil. The inductor was aligned such that the magnetic field lines were parallel to the normal direction of the wire-wound coil, to achieve maximal magnetic coupling for the measurement. It was found that the magnetic shield provides a suppression of the magnetic coupling of a factor of 26, as shown in figure 2.47.

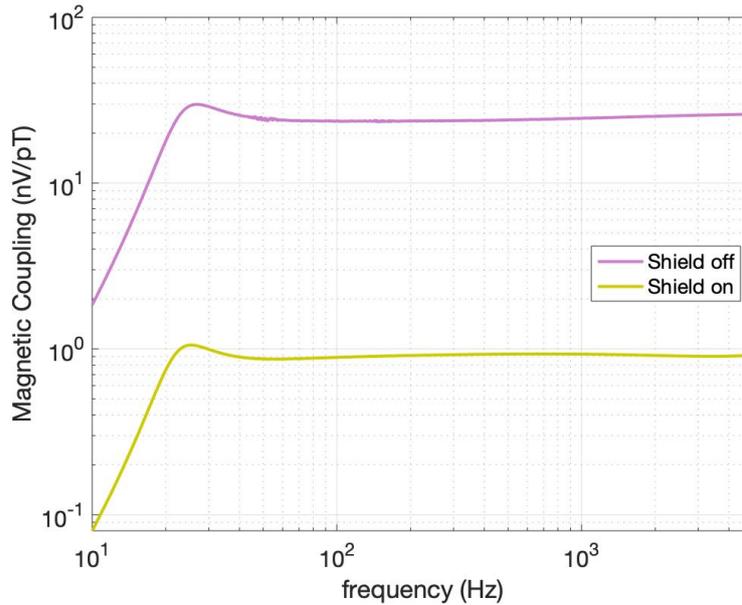


Figure 2.47: The voltage noise induced by a magnetic field with and without the mu-metal shield on. Roughly a factor 26 attenuation is seen at 10 Hz to 5 kHz

Figure 2.48 shows an FFT of the magnetic noise floor at the Livingston HAM 6 chamber, close to where the photodetector electronics will be located but outside the vacuum chamber.

The magnetic coupling measurement from figure 2.47 was used with the ambient magnetic field data from figure 2.48 to estimate the magnetic noise coupled into the inductor at the HAM 6 chamber; see figure 2.49. The magnetic field at the LIGO HAM 6 chamber is fairly isotropic [84] (when averaged with time), therefore it is not anticipated that there is a preference in orientation when installing the inductors.

As mentioned in section 2.2.7 the effective magnetic pickup in the implementation at LIGO can be further reduced by mounting the two inductors in one housing in opposite directions, such that their magnetic pickup nearly cancels if their sum output is processed (at least for spatially uniform magnetic field components across the two inductors). As mentioned in section 2.1.4, for DC readout it is always the sum of the two channels that carries the scientific signal. Although it is not shown in figure 2.9, BHD will actually have two photodiodes in each path after the LO has been combined with the GW signal. This means that the inductor orientation will not need to change when BHD is implemented.

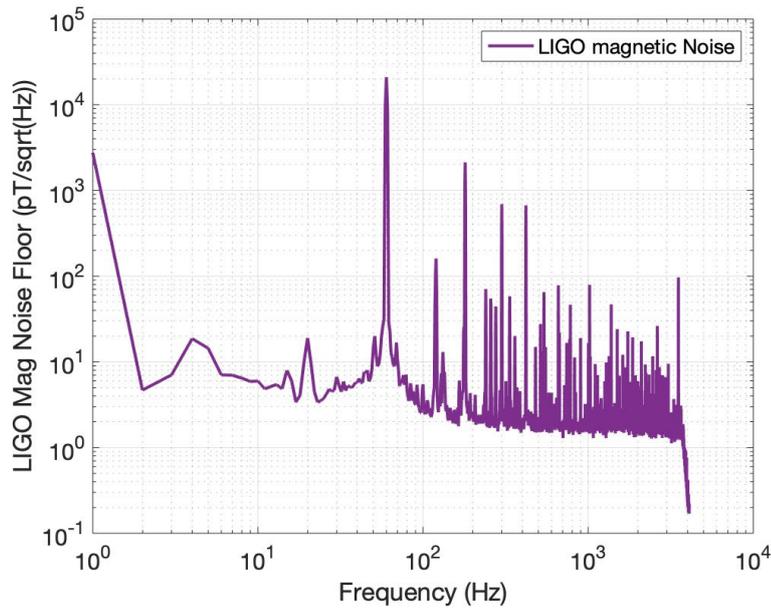


Figure 2.48: LIGO magnetic noise floor from the outside of the HAM 6 chamber at Livingston. The magnetic fields inside the chamber are even lower when the chamber is sealed and have a $1/f$ roll off with a pole frequency at ~ 10 Hz.

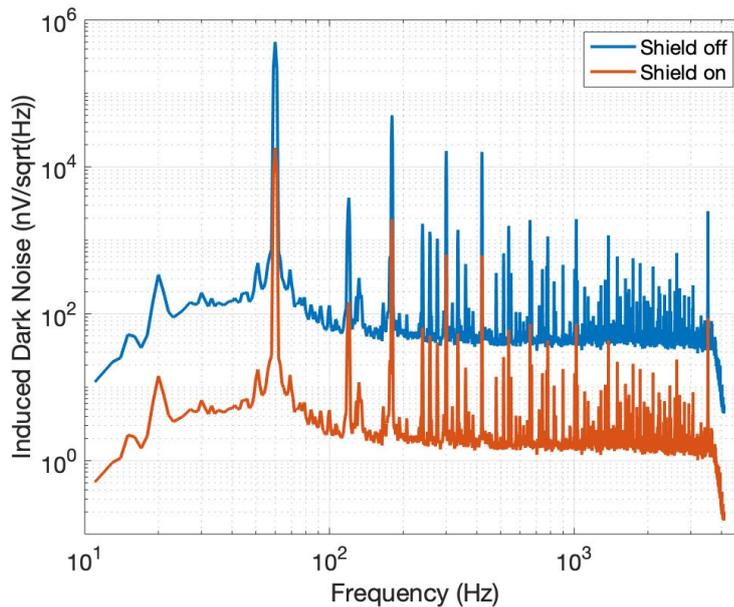


Figure 2.49: The voltage noise estimated to couple into the inductor due to the ambient magnetic fields at the outside of the LIGO HAM 6 chamber. The actual coupling values are even lower due to the reduced magnetic fields inside the HAM 6 chamber when it is sealed.

It is also worth noting that the largest source of pickup in figure 2.49 comes from the mains electricity. In the UK this is at 50 Hz and in the US it is at 60 Hz. A comparison of the estimated magnetic pickup noise to the dark noise of the readout

is shown below in section 2.3.2.

2.3.2 Audio Readout

There are two boards that have been installed at the LIGO sites, labelled SN01 (Livingston) and SN02 (Hanford). Each have two independent channels to process the signals from the two photodiodes in the main readout path (channels A and B), with unique transfer functions on each channel. This uniqueness was not by design, rather it is a by-product of the uncertainty in the inductor properties (see section 2.3.1) and component tolerances. For the sake of succinctness, the following section details measurements carried out on only channel A of board SN01 and ignores the measurements on other boards and channels. There is some discrepancy between the measurements taken in Cardiff before they were installed and upon their installation at the LIGO corner stations which is explored in section 2.3.8.

The pin layout of the circuit and voltage levels used for the measurements are shown in appendix 5.3.

Transfer Function

The following audio path transfer functions were measured by using a test relay on the DCPD (see figure 5.8). A voltage is sent to the test input which has a $100\text{ k}\Omega$ resistor before the normal photodiode input pin. The $100\text{ k}\Omega$ resistor is much larger than the impedance of the rest of the circuit across all frequencies (see figure 2.31), so the test input is a good model of a current input to the readout on the photodiode pin. Each DCPD readout contains two parallel channels with differential outputs of the audio path. There was limited lab equipment in Cardiff during testing so the output and drive voltage were all single ended.

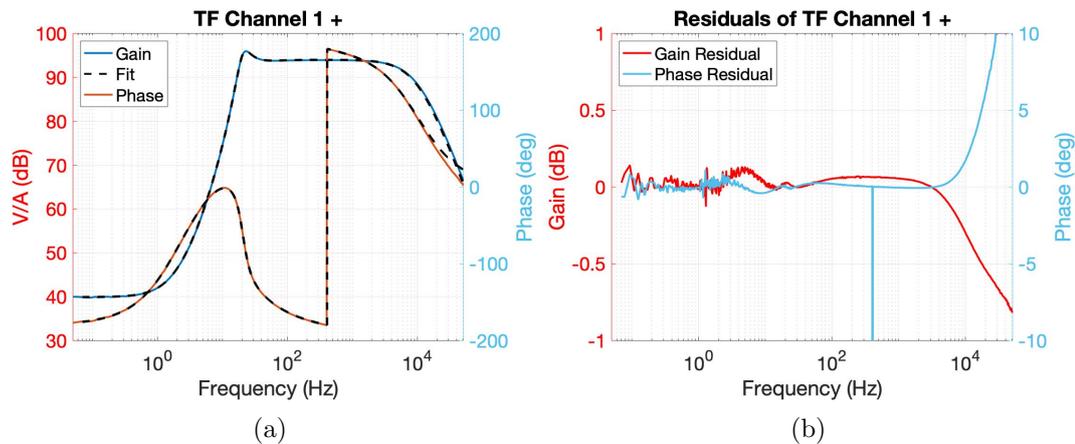


Figure 2.50: **a)** transfer function from the test input to audio output 1+. **b)** Residuals of the fit. Note that in the transfer function the $100\text{ k}\Omega$ resistor has been accounted for to give the units of the transfer function in V/A. Units of dB are used i.e. $20\log_{10}(V/A)$.

Parameter	Frequency (Hz)	Q-value
Zero 0	1.97	0.495
Pole 0	20.57	2.08
Zero 1	7.076	-
Pole 1	33.74	-
Pole 2	10.18 k	-
Pole 3	10.18 k	-

Table 2.1: Parameters obtained from a LISO [73] fit on the measured transfer function in figure 2.50

Noise Floor

The noise of the audio readout was measured using the differential outputs with no drive source connected and amplifying the measurement by a factor of 50, bringing the overall noise floor amplification up to a factor of 100 so that it is larger than the ADC noise of the analyser. On its smallest range the input noise of the analyser is $20 \text{ nV}/\sqrt{\text{Hz}}$. The circuit is of course sensitive to various forms of pickup (see section 2.3.1), therefore efforts were made to ensure the lab environment was as quiet as possible upon measurement. The noise levels in figure 2.51 refer to a single ended circuit output, i.e. the factor 50 pre-amp gain has been compensated for, as well as another factor 2 for the differential output. As mentioned in section 2.2.2 it is more useful to display the noise floor referred to the input of the DCPD rather than as a voltage at the output. This is done by dividing the voltage noise at the output by the transfer function and is shown in figure 2.52.

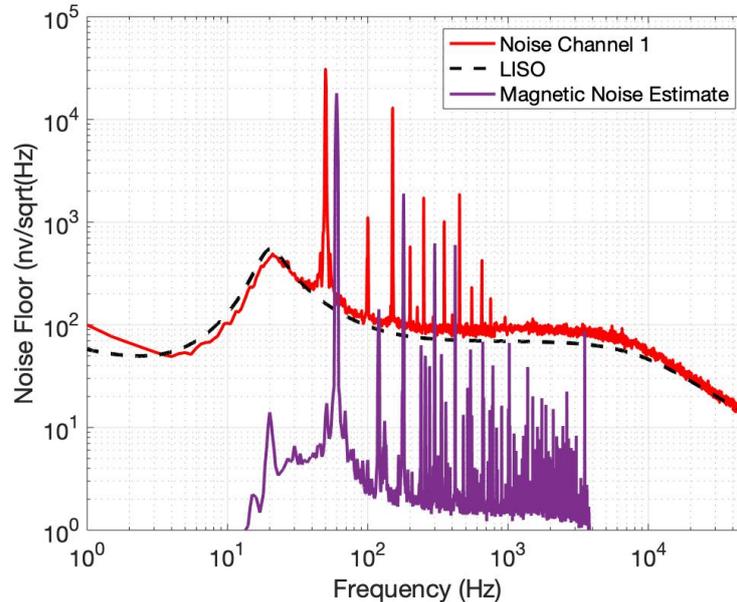


Figure 2.51: **Red**: noise floor measured on channel A audio output of SN01. **Black**: LISO model of the noise floor. **Purple**: projection of the magnetic pickup in the 2.5 H inductor from the ambient fields outside the HAM 6 chamber at LIGO. Note the peaks at the 60 Hz mains frequency used in the U.S. rather than the 50 Hz used in the UK. The only significant contribution of the magnetic noise to the readout noise floor is expected at the 60 Hz harmonics. If these signals are coupling into the inductor, this pickup should be significantly reduced when the sum of the photodiodes signals is taken (see section 2.2.7)

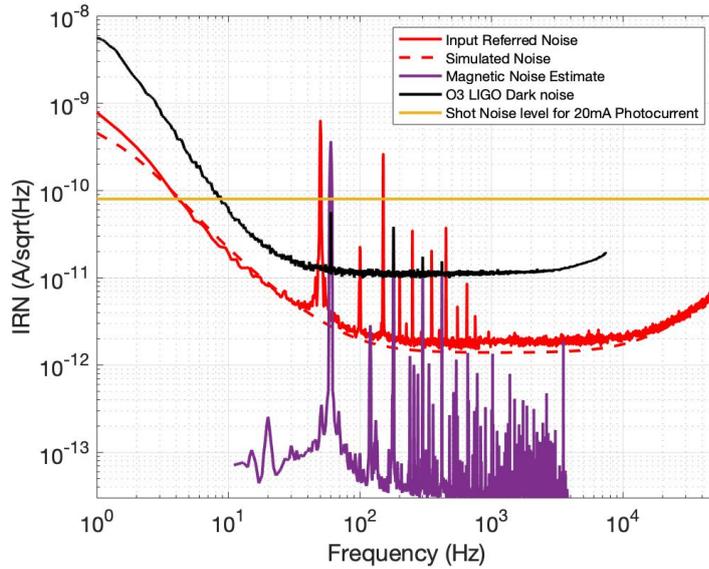


Figure 2.52: **Red**: the measured Input Referred Noise (IRN) for channel 1. **Red-dashed**: the corresponding noise simulation from LISO. **Black**: the noise floor of the previous readout used in LIGO during O3. **Yellow**: the noise contribution of an un-squeezed shot noise from a 20 mA photocurrent. **Purple**: projection of the magnetic pickup in the 2.5H inductor from the ambient fields outside the HAM 6 chamber at LIGO (see section 2.3.1).

Figure 2.52 shows that the new input referred noise is a factor 47 below the un-squeezed shot noise for a 20 mA photocurrent at 1 kHz. We can refer back to section 2.1.4 and recall that a good rule of thumb is to keep the technical noise sources such as dark noise at least a factor of 10 below the shot noise level. With these new readouts there is enough headroom for a 20 mA un-squeezed shot noise level to be reduced by a factor of 4.7 and still maintain a factor 10 between the dark noise and shot noise level. This would correspond to squeezing levels up to 13 dB. The upgrade plan for the fourth observing run (March 2023- March 2024 [85]) is to implement 6 dB of squeezing and the upgrade plan for the fifth observing run (tentatively scheduled for 2026-2029 but likely to change as the date gets closer [85]) is to implement 10 dB of squeezing. With these plans in mind, it appears that the dark noise with the new detector readouts should not be significantly limiting to the LIGO DARM sensitivity for at least the next 7 years. The new dark noise is also a factor 6.5 below the dark noise of the circuit used during the previous observing run at 1 kHz but also has an improved dynamic range over the old readout as mentioned in section 2.2.

2.3.3 Bias Readout

Transfer Function

The transfer functions for the bias path were taken by injecting a swept sine at the bias test port and measuring the signal at the bias output. As in section 2.3.2 the measurements were performed single ended. The result is shown in figure 2.35.

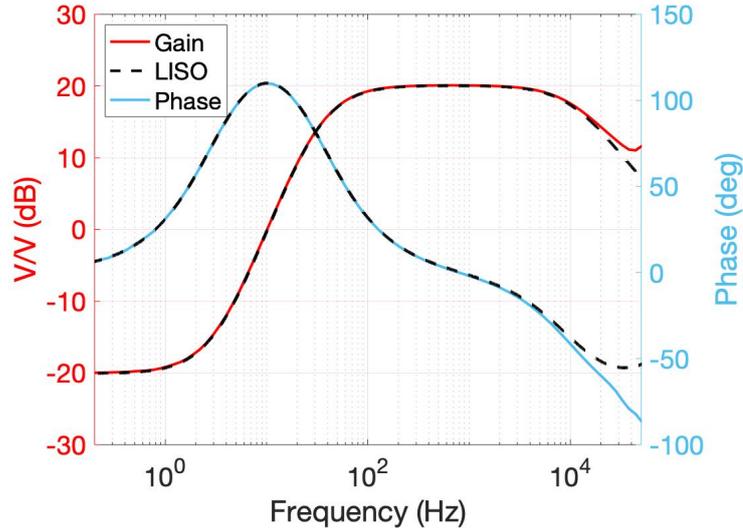


Figure 2.53: transfer function from bias test input to the positive output pin on channel A of SN01 and LISO simulations of the expected transfer function. It is not known why the transfer function diverges from the model beyond ~ 40 kHz but it was decided this was at a high enough frequency to not be a concern for the bias monitoring at this time.

Noise Floor

Noise floor measurements were performed by taking the differential output of the bias path without any drive source connected. Each output was amplified by a factor of 50 to bring the overall gain of the differential measurement to a factor of 100. This has been compensated for in the graphs below such that they represent the noise from a single ended output.

Below 100 Hz the noise is dominated by the current noise from the LT1028 op amp labelled ‘U5’ then beyond 100 Hz it is dominated by the 1.1 k Ω resistor R_{15} .

As mentioned in section 2.2.6 the bias noise has not been the subject of significant investigation in the past. The expected noise from the regulator used in the bias voltage supply (LT3045) is quoted as being between $2 \text{ nV}/\sqrt{\text{Hz}}$ and $70 \text{ nV}/\sqrt{\text{Hz}}$ depending on the setup. A colleague used the same regulators in a similar setup to the one used by LIGO and measured the noise to be $16 \text{ nV}/\sqrt{\text{Hz}}$ at 1 kHz. Assuming this measurement will be consistent with the LIGO measurement, the expected

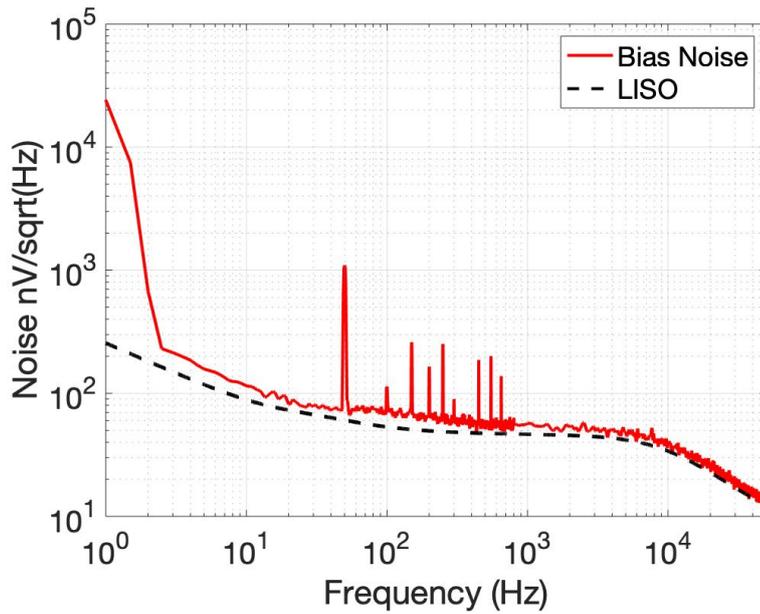


Figure 2.54: **Red**: voltage noise floor at the output of the bias monitoring output. **Black dashed**: simulation of the predicted noise floor from LISO.

$16 \text{ nV}/\sqrt{\text{Hz}}$ would be amplified to $160 \text{ nV}/\sqrt{\text{Hz}}$ by the bias monitor amplification stage, roughly two times higher than the noise floor of the monitor itself. It will also be possible to detect any unexpected signals in the bias noise (e.g. EM coupling) that exceed the noise floor of the bias monitor. It will also now be possible to perform coupling measurements of the bias noise into DARM by injecting and monitoring the noise on the bias voltage.

2.3.4 RF Readout

Transfer Function

The bandwidth of the differential amplifier (AD620) in the relay path is not high enough to obtain a transfer function at RF frequencies. Instead the transfer function was obtained by using a swept sine measurement injected directly to the photodiode input pin and measuring the signal at the RF output.

The transfer function shows a pole at approximately 250 kHz, formed primarily by the 4.7 nF capacitor C_1 and the 50Ω input resistor to the RF amplifier (see figure 2.30 for details), and a flat response until roughly 30 MHz.

From section 2.2.5 the expected 3 MHz signal in the photocurrent is $54 \mu\text{A}$ which after passing through a 50Ω transimpedance resistor and being amplified by 22 dB will give an expected signal level of 28 mV. This number also take into account the effect of the photodiode capacitance mentioned in section 2.2.5. The measurement in figure 2.55 did not have a photodiode (or any capacitance to simulate one) in the

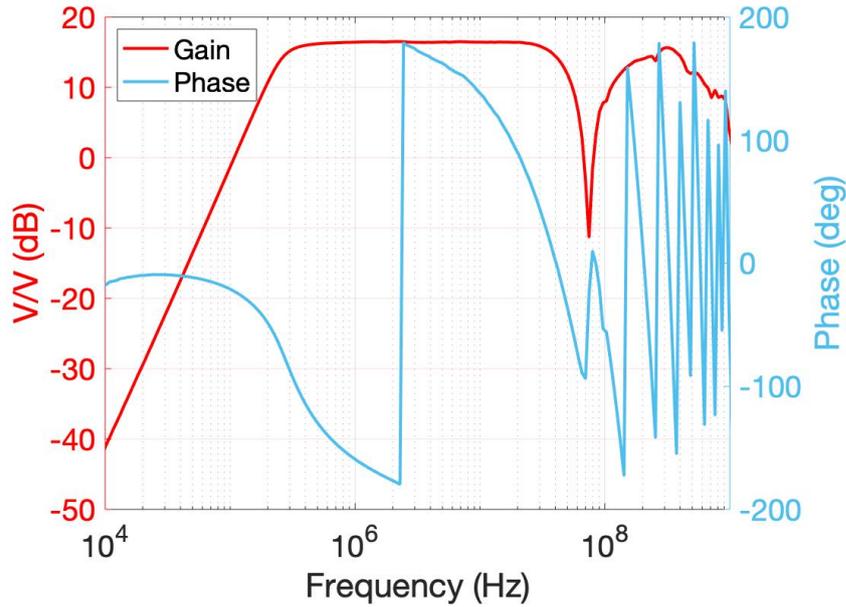


Figure 2.55: Transfer function from photodiode input pin to the RF output. The features at frequencies over ~ 30 MHz are thought to arise from the injection setup. The gain of the RF amplifier is 22 dB but the output is $50\ \Omega$ so the measurement shows 16 dB gain from 250 kHz to 30 MHz. The feature at 70 MHz is also believed to be a feature of the injection setup.

setup, so this drop at higher frequencies is not seen.

Noise Floor

As mentioned in section 2.2.5 the dominant contribution to the RF output is the $50\ \Omega$ transimpedance resistor and the smaller noise figure of the amplifier which brings the overall electronic noise to about $14\ \text{nV}/\sqrt{\text{Hz}}$.

A measurement at the RF output confirmed this and found $14\ \text{nV}/\sqrt{\text{Hz}}$ of noise at 3 MHz which gives an expected SNR of 530,000 at 3 MHz for a 1 Hz bandwidth. It was decided that further investigation was unnecessary.

2.3.5 Op Amp Selection

The key contributions to the noise of the DCPD are plotted in figure 2.25, which models an OPA2210 op amp in the first transimpedance stage. This section covers why this particular op amp was selected and its importance in the readout.

The initial prototype of the DCPD used an LT1028 amplifier in the first transimpedance stage (U3), however testing showed that the measured noise floor was roughly a factor 2 higher than predicted between 30 Hz – 30 kHz as shown by the IRN plot in figure 2.56.

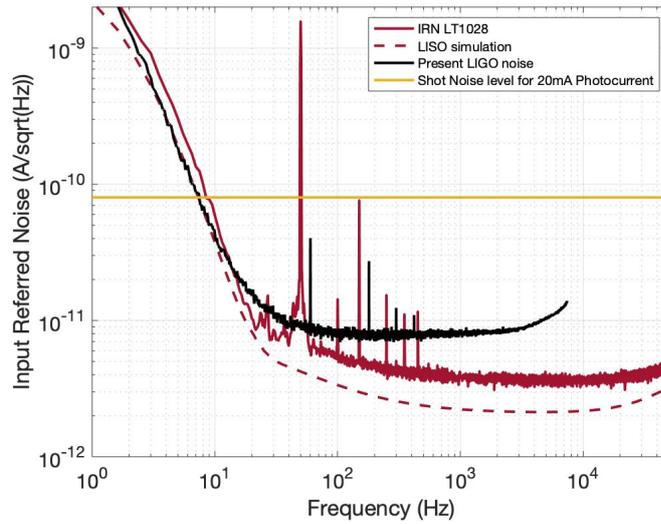


Figure 2.56: **Maroon solid:** Noise floor with LT1028 as U3. Data was taken from four FFTs; one in the range 0–0.8 kHz with a 0.5 Hz bin width, one in the 0.8–7 kHz range with a 4 Hz bin width, one in the 7–30 kHz range with a 16 Hz bin width and one in the 30–55 kHz range with a 16 Hz bin width. **Maroon dashed:** LISO model of the noise floor for LT1028. **Yellow:** Shot noise level due to 20 mA photocurrent. **Black:** LIGO dark noise from O3 [86].

The noise floor contributions of the DCPD audio output with an LT1028 op amp (*not* the OPA2210 shown in figure 2.25) is shown in figure 2.57.

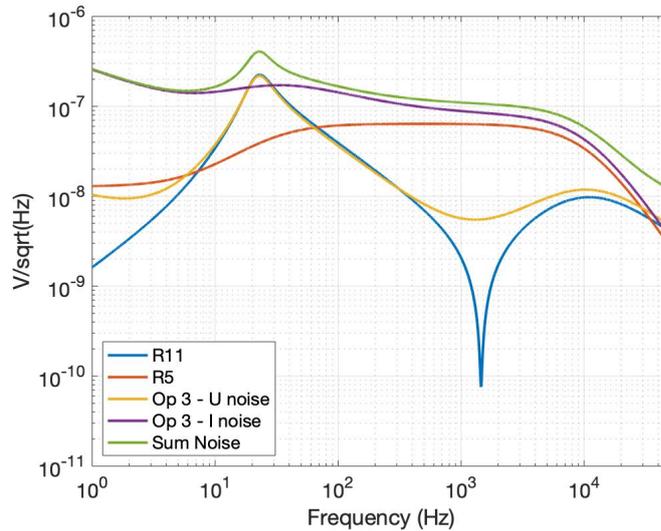


Figure 2.57: Key components contributing to the total noise in the DCPD for LT1028 installed in the first transimpedance stage. All components were used when calculating the sum of all noise sources. R11 is the 100 Ω damping resistor. R5 is the 10 k Ω transimpedance resistor. Op 3 is the op amp in the transimpedance stage, in this case it is an LT1028. The noise floor between 30 Hz to 30 kHz is limited by the current noise of the LT1028.

Figure 2.57 shows that in the case of the LT1028 the noise floor is limited by the current noise at the op amp inputs. It was therefore assumed that the op amp current noise for the LT1028 was higher than the data sheet specification which was resulting in the measured noise floor increase in figure 2.56.

The noise floor measurement was then repeated for 4 different op amps occupying the transimpedance stage (U3) which were selected based on the voltage noise, current noise and speed (see table 2.2). These op amps are LT1028, TLE2027, OPA2209 and OPA2210. It should be noted that testing of LT1028, TLE2027 and OPA2209 were carried out before changes to the audio/RF cross over were made whereas testing of OPA2210 was carried out after these changes, hence there are some small differences in the transfer function. All measurements are shown here, as a comparison of each opamp's performance.

Table 2.2 shows some critical performance data of the op amps under test, taken from their data sheets.

	LT1028	TLE2027	OPA2209	OPA2210
U noise at 10Hz ($\text{nV}/\sqrt{\text{Hz}}$)	-	-	3.3	2.5
U noise at 1kHz ($\text{nV}/\sqrt{\text{Hz}}$)	1.1	2.5	2.2	2.2
I noise ($\text{pA}/\sqrt{\text{Hz}}$)	1	0.4	0.5	0.4
GBW (MHz)	72	15	18	18
Slew Rate ($\text{V}/\mu\text{s}$)	11	2.8	6.4	6.4

Table 2.2: LT1028, TLE2027, OPA2209 and OPA2210 key values from product data sheets.

The circuit input referred noise with all of the tested op amps is shown in figure 2.58.

Both OPA2209 and its newer counterpart, OPA2210, were tested as U3. In most ways they are near identical, except that the OPA2210 is stated to exhibit slightly lower noise at low frequencies (see table 2.2). The IRN for OPA2209 exhibited lower noise than the LT1028 and TLE2027 at all frequencies. It is thought that the increase in the IRN of OPA2210 at 10 kHz is due to the changes of the circuit board that were carried out before testing of this opamp. The feature in the noise floor at 40 kHz is only a feature of this version of the circuit and is not seen in future versions.

The drawback of the OPA2210 opamp is the lower slew rate compared to the LT1028. This is explored in section 2.3.6.

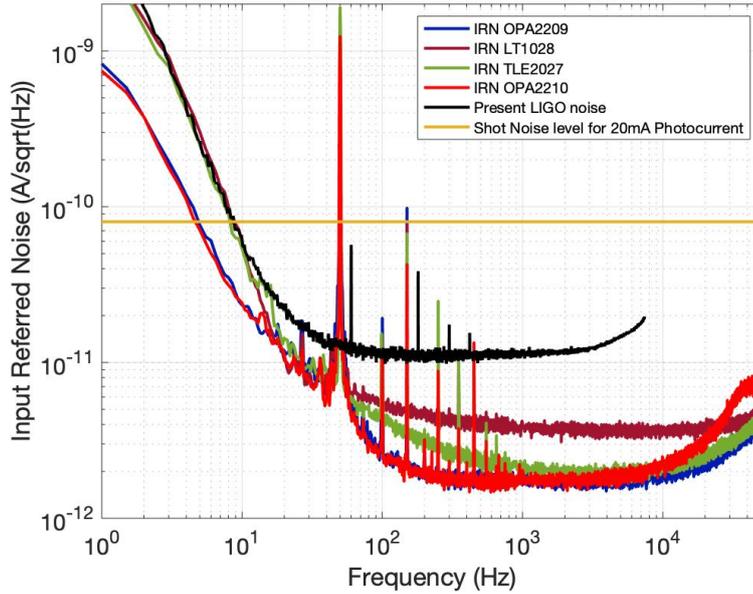


Figure 2.58: **Maroon:** Input Referred Noise (IRN) of the measured noise floor for LT1028 as U3. **Green:** IRN of the measured noise floor for TLE2027 as U3. **Blue:** IRN of the measured noise floor for OPA2209 as U3. **Red:** IRN of the measured noise floor for OPA2210 as U3. **Black:** current LIGO dark noise (O3) [86]. **Yellow:** shot noise of a 20 mA photocurrent. The higher IRN above 10 kHz for OPA2210 comes from the change in the transfer function of the new circuit rather than an increase in the op amp noise.

2.3.6 AC Signal Saturation

The op amps in the audio path are not as fast as the MAR-6SM amplifier in the RF path and there is a concern that they could become saturated by any high frequency noise that is still present in the audio path, (see figure 2.32 for the fraction of the photocurrent sent to the audio path). Figure 2.59 shows a measurement of input referred photocurrent at LLO during O3 out to 100 kHz.

Considering the signal structure at 75 kHz which is $\sim 4 \times 10^{-8} \text{ A}/\sqrt{\text{Hz}}$ and assuming a BW of 1 kHz this corresponds to $10^{-6} \text{ A}_{\text{rms}}$ at 75 kHz. Taking this as reference a signal was injected to the PCB up to $7 \times 10^{-5} \text{ A}_{\text{rms}}$ at 100 kHz, almost two orders of magnitude higher than the signals seen in the LIGO readout. The increase to the noise in the audio band is shown in figure 2.60.

Based on this result it can be assumed that the op amps are fast enough to keep up with signals between 10 – 100 \times typical LIGO levels before a noise floor increase is observed. This is slightly less than expected from the datasheet value. The OPA2210 have slew rates of $6.4 \text{ V}/\mu\text{s}$ and the LT1028 has a slew rate of $11 \text{ V}/\mu\text{s}$. A $100 \mu\text{A}$ current passing through the $10 \text{ k}\Omega$ load resistor generates 1 V at the output of the OPA2210 and 4 V at the output of the LT1028 in the second amplification stage at

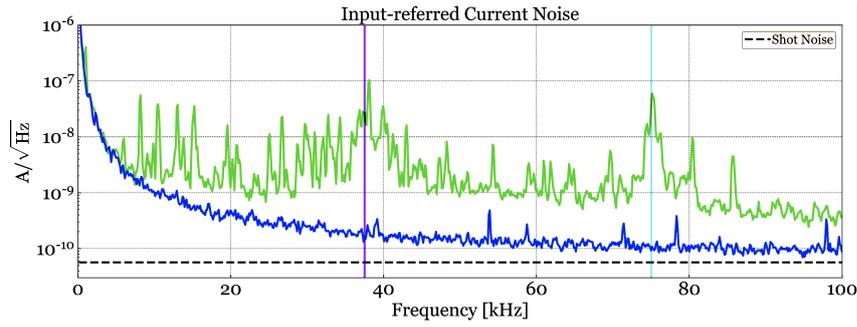


Figure 2.59: **Green:** input referred current noise from the Livingston detector during O3 [87]. **Black dashed:** the predicted shot noise level for the corresponding DC photocurrent level. **Blue:** measured laser noise from a pick-off before the interferometer. The noise in this measurement no longer contains features from the interferometer and comes close to the predicted shot noise level shown by the yellow-dashed line.

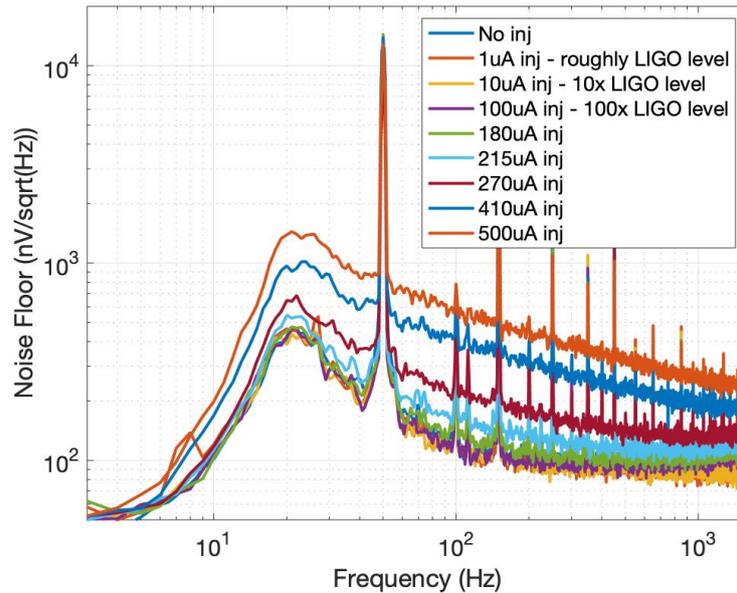


Figure 2.60: Noise floor of audio readout for increasing amplitudes of 100 kHz signal. The $1\ \mu\text{A}$ and $10\ \mu\text{A}$ injection measurements are indistinguishable from the 'no injection' setup in the plot.

100 kHz, which is within the specified operating range of both op amps. It is possible that the op amps are roughly 10 times slower than the manufacturers specifications but this seems unlikely. More probable is that there is additional noise injected to the circuit board from the signal generator. Although interesting, this point is simply an academic one as the OPA2210 has been measured to be fast enough to keep up with the signal content levels 10 times higher than seen during a typical LIGO run. Based on signal saturation results and the noise results from section 2.3.5 the OPA2210 was selected as the op amp in the transimpedance stage.

The RF readout was also monitored during this test and no noise floor increase was observed; however a small distortion of the signal at these amplitudes resulted in harmonics in the RF region. Increasing the signal from $1 \times 10^{-5} A_{\text{rms}}$ to $5 \times 10^{-5} A_{\text{rms}}$ (a factor of 5) caused an increase in the 200 kHz harmonic from $23 \mu V_{\text{rms}}$ to $230 \mu V_{\text{rms}}$, a factor of 10. This compares to the fundamental signal at 100 kHz of $250 mV_{\text{rms}}$. However these signals are over $10\times$ the expected signal content during a LIGO run.

2.3.7 DC Signal Saturation

Since the LIGO detector operates just off the dark fringe, an investigation was carried out to investigate whether DC photocurrent gives rise to excess noise. Due to the ultra-low noise of the DCPD, the noise from a bench-top power supply will easily obscure the noise floor that we are trying to measure. A filter setup was designed to inject DC signals to the readout without injecting noise over the input referred current noise floor of the readout.

The first stage is an active filter setup shown in figure 2.61:

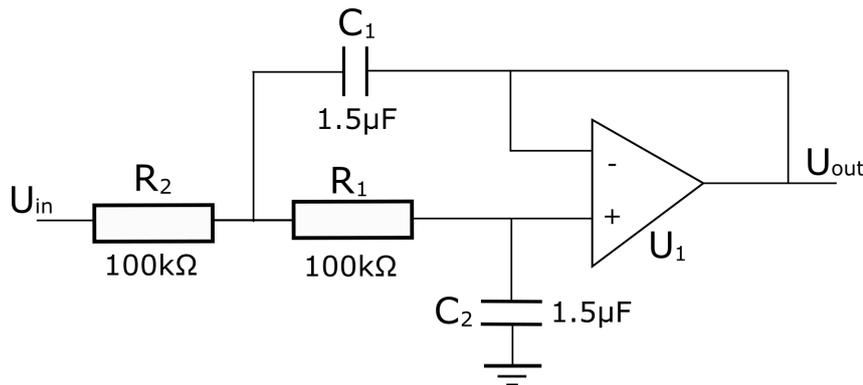


Figure 2.61: Schematic for the active filter used to inject DC current to the readout.

This setup forms a double low pass filter with poles at 0.7 Hz, effectively filtering the noise from the DC supply, however it becomes dominated by the noise from the TLE2027 op amp (U_1) ($4 \text{ nV}/\sqrt{\text{Hz}}$) at around 100 Hz. This noise is still large enough to result in an increase in the noise floor of the readout; therefore an additional, passive filtering stage was used.

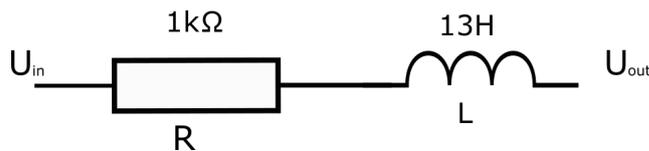


Figure 2.62: Series resistor and inductor forming a passive low pass filter

By choosing $R = 1 \text{ k}\Omega$ resistor and $L = 13 \text{ H}$ this filters the TLE2227 noise contribution below the noise floor of the readout whilst allowing injection of DC signals

up to roughly 10 mA. The 13 H inductance value is of course a fairly large value so one of the Sowter inductors was used in this passive filter. As we will see in section 2.3.7 this may have resulted in more noise being injected into the readout than was realised.

The results from the DC signal injection are shown in figure 2.63 and show an approximately linear increase of the noise floor between 100 Hz and 10 kHz as DC current is increased. Figure 2.64 shows the average voltage noise in the bandwidth 1-5 kHz for increasing DC current injections and shows a gradient of roughly 9 nV/mA.

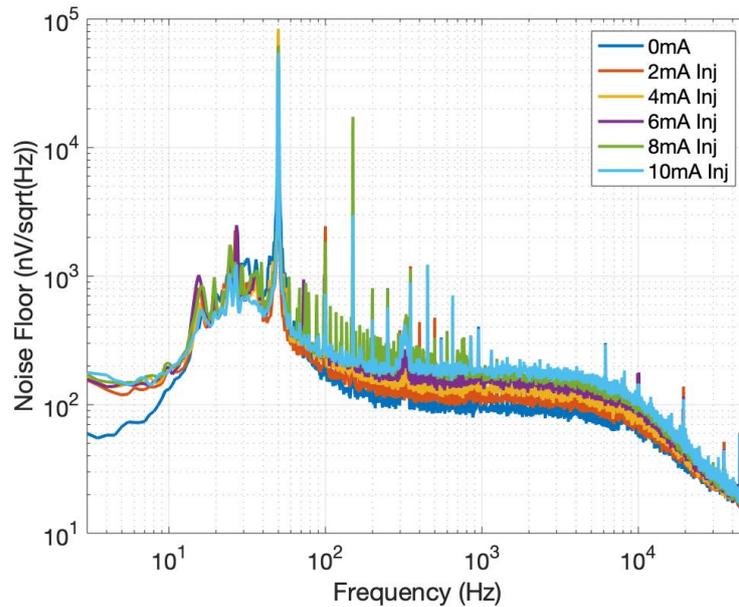


Figure 2.63: Noise floor of the readout for increasing DC current injected at the photodiode input pin.

Due to concerns about the noise produced by the active filter, the setup was also repeated using standard 9 V chemical batteries which exhibit low enough voltage noise that they do not need the active filter in figure 2.61. The same noise floor increase was observed.

Inductor Noise Testing

The majority of the DC current passes through the large 3 H inductor which as mentioned in section 2.2.4 is a custom made component with unknown noise characteristics. The first assumption was therefore that this additional noise could be coming from the inductor, potentially the result of a magnetic effect induced by the DC current within the inductor. If the noise is coming from the inductors it is unfortunately very close to the noise floor of the SRS pre-amp at $4 \text{ nV}/\sqrt{\text{Hz}}$ and there is a lot of pickup noise across the inductors which makes measurement of any induced current noise non-trivial.

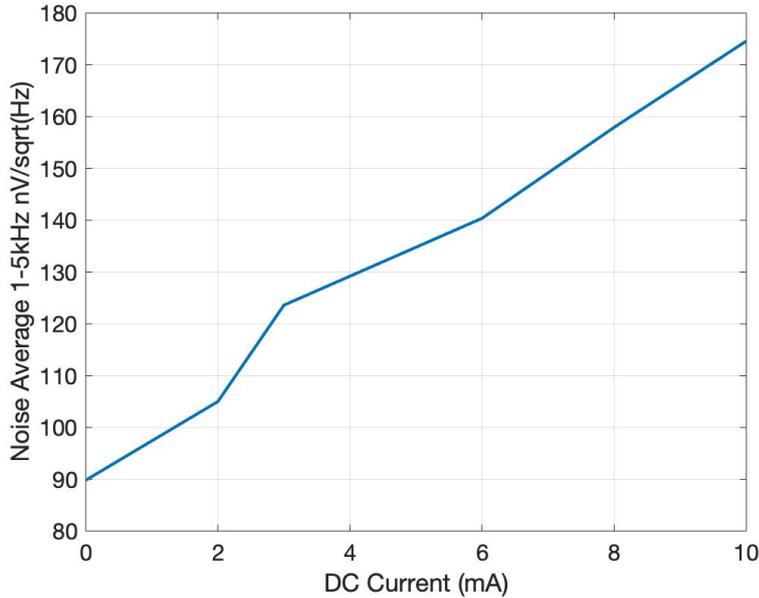


Figure 2.64: The average noise within the 1-5 kHz bandwidth shows an approximately linear increase of the noise floor as DC current is increased.

A test was carried out to try to measure the noise of the inductors independent from the DCPD readout by injecting a 9 mA current through two $2\text{ k}\Omega$ resistors and one inductor in series.

Two inductors were tested, one with a significantly larger inductance to see if there was any relationship between the inductance value and the observed noise. The voltage across resistor R_2 for each measurement is shown in figure 2.66.

The current noise from the inductor can then be calculated by first taking the difference in quadrature of the baseline measurement (‘shorted wires’) to the test measurement and then dividing by the impedance of the test resistor R_2 . The result is shown in figure 2.67.

The noise at 1kHz is $\sim 1.7\text{ pA}/\sqrt{\text{Hz}}$ for both inductors, implying the noise is independent of inductance value.

Converting the voltage noise from figure 2.63 into an input referred plot, one can see that the IRN increases from roughly $2\text{ pA}/\sqrt{\text{Hz}}$ to $3.8\text{ pA}/\sqrt{\text{Hz}}$ at 1 kHz for a 10 mA DC injection. Assuming that the DC current induced noise is incoherent with the dark noise this would indicate the additional noise has a spectral density of roughly $3.2\text{ pA}/\sqrt{\text{Hz}}$ at 1 kHz and 10 mA DC injection.

For the test in figure 2.63 a second inductor was used to filter the injected DC current, and therefore two inductors would be contributing to the observed current noise. Interestingly if we model these two inductors as having coherent noise then we get a fairly good approximation for the noise floor increase as shown in figure 2.68.

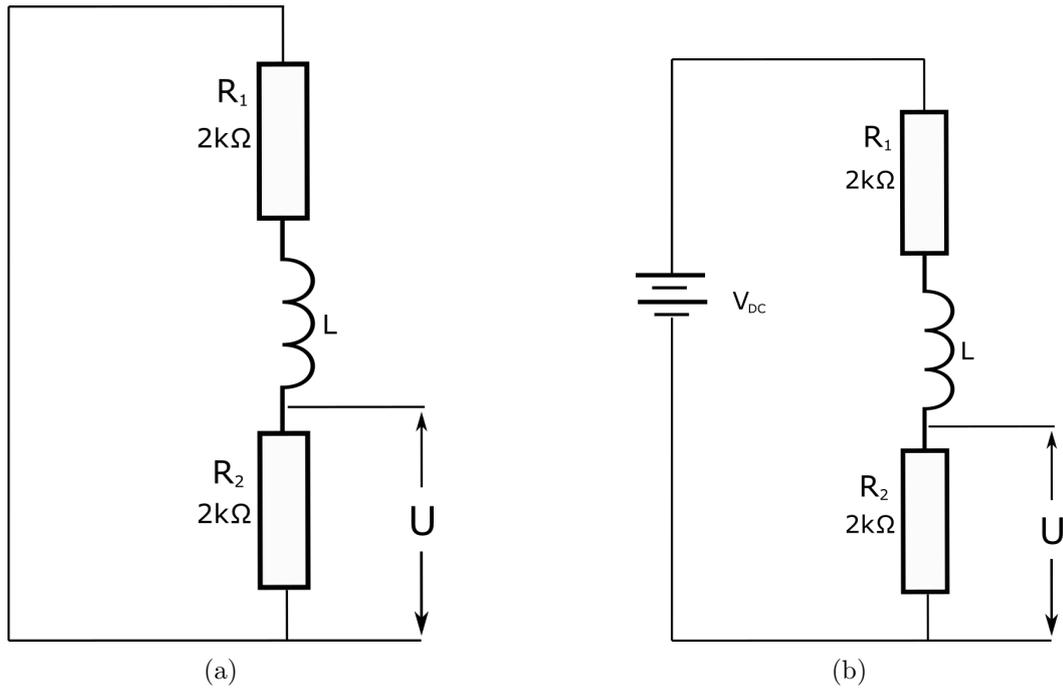


Figure 2.65: Two test setups used to measure the inductor current noise. **a)** the reference measurement, shorting the wires means only the Johnson noise of R_2 is being measured. **b)** $4 \times 9\text{ V}$ batteries were used to inject a DC current through the setup. The measurement now contains the Johnson noise of R_2 and the additional voltage noise from the induced current noise of L_1 passing through R_2 . For **b)** standard 9 V batteries were used since they exhibited low enough voltage noise to not dominate the setup.

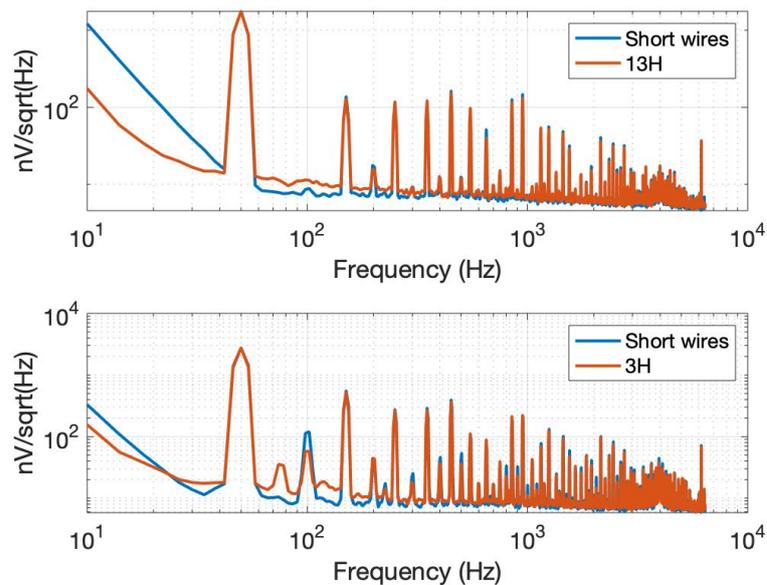


Figure 2.66: Voltage measured across R_2 for the shorted setup and 10 mA DC injection.

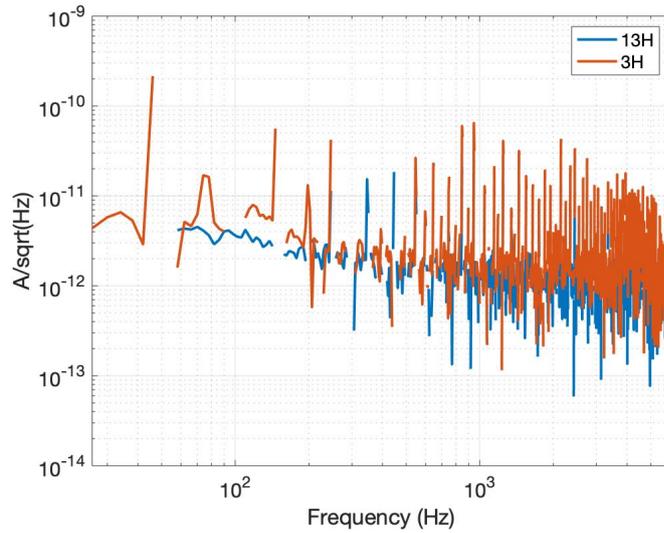


Figure 2.67: The current noise from the 13 H inductor and 3 H inductor induced by a 10 mA DC current. This plot was made by taking the difference in quadrature between the measurement with the battery connected and the setup with the battery shorted. The gaps in the plot are due to the noise in the shorted setup being greater than the noise in the test setup so the quadrature noise becomes imaginary and is not plotted. The noise is higher is due to 50 Hz mains pickup.

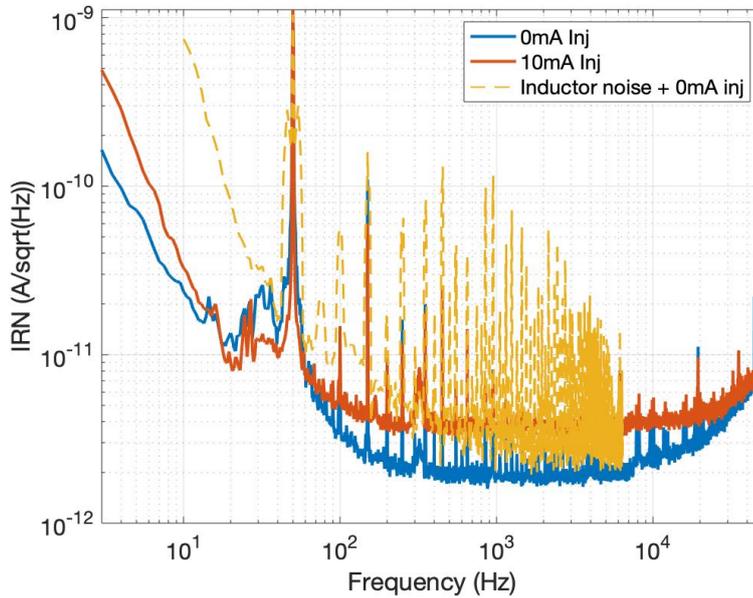


Figure 2.68: Converting the voltage noise for the 0 mA and 10 mA setup from figure 2.63 into an input referred plot. Also shown is the sum in quadrature of the 0 mA plot and the two inductor noises from figure 2.67 that were summed coherently together.

The other possibility is that there is residual noise coming from the injection system. Based on the result below it seems that the latter is the more likely explanation but

this should be the subject of further investigation.

Magnetic DC Saturation

Further testing was carried out to see if this noise floor increase was due to the magnetic properties of the mu-metal core or due to the test set-up used to inject the DC current. A Helmholtz coil was used to inject a static magnetic field into the inductor to try to replicate the result seen when using a DC current.

Due to time constraints it was not possible to remove all sources of noise coupling into the inductor so the baseline noise measurement in figure 2.69 is not particularly clean, however the noise floor does reach the SR pre-amp noise floor at ~ 3 kHz so it was expected that any increase in the inductor noise that is comparable to those in figure 2.63 would be observed.

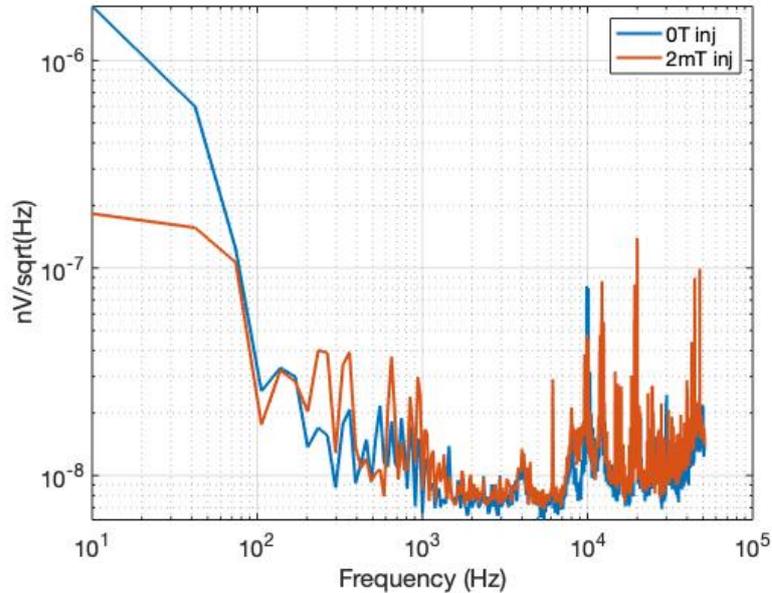


Figure 2.69: A Helmholtz coil was used to inject a DC magnetic field through the inductor aligned with its core. The shield was left on for this test because there were concerns that ambient fields or noise on the Helmholtz field would couple into the measurement. A magnetometer was used to measure the magnetic field inside the mu-metal shield and the field was found to be ~ 2 mT for a 1.5 A current through the coil. This is fairly close to the induced magnetic field from applying a 10 mA DC current to the inductor.

No significant change is observed in the noise levels when a magnetic field is applied, which implies that the noise seen on the inductor is not magnetic in origin.

Ultimately the reason for this noise floor increase is not fully understood at the time of writing. It was thought that this noise could be the Barkhausen noise induced within the mu-metal inductors however the tests using a magnetic field

imply otherwise. Perhaps there is noise being introduced within the readout by another component that has not been considered yet. Alternatively it may be that the current injection setup is not effective enough at filtering the AC noise when injecting a DC voltage and we are simply observing an increase in injection noise in figure 2.63.

2.3.8 LIGO Results

The DCPD readouts were sent to the Livingston and Hanford sites and installed in February 2022.

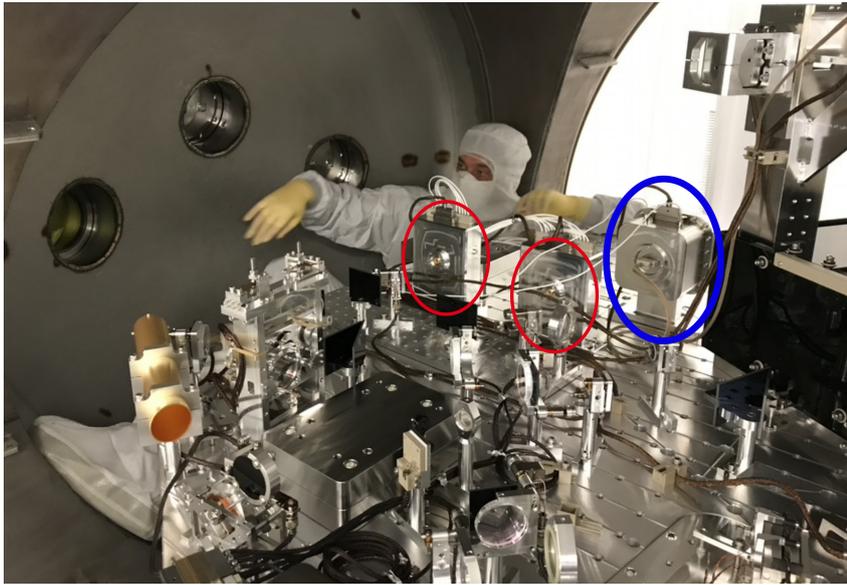


Figure 2.70: A colleague taking a well-earned rest inside the HAM-6 chamber after installing the DCPD readouts at LIGO Livingston. The vacuum case containing the new DCPD and inductors is circled in **blue**. The two photodiodes sending signals to the readout electronics are circled in **red**. On the far right side of the image is the output mode cleaner on its suspension system. It is encased in black glass to minimise outside light coupling into the beam.

Precise and consistent measurements of the DCPD transfer functions to the sub-1% level are required in order to calibrate the inverse transfer functions that are used to reconstruct the gravitational wave signal. Various measurements of the DCPD readouts were made upon installation at the two sites and some time-dependent discrepancies were discovered. This section details and explains the discrepancies.

LIGO transfer function Discrepancies

The most egregious difference is shown in figure 2.71 which shows a comparison between the transfer function taken using a bench-top setup in Cardiff in June 2021 and then a measurement taken upon installation at the LIGO HAM-6 end-station in February 2022 with the chamber at atmospheric pressure.

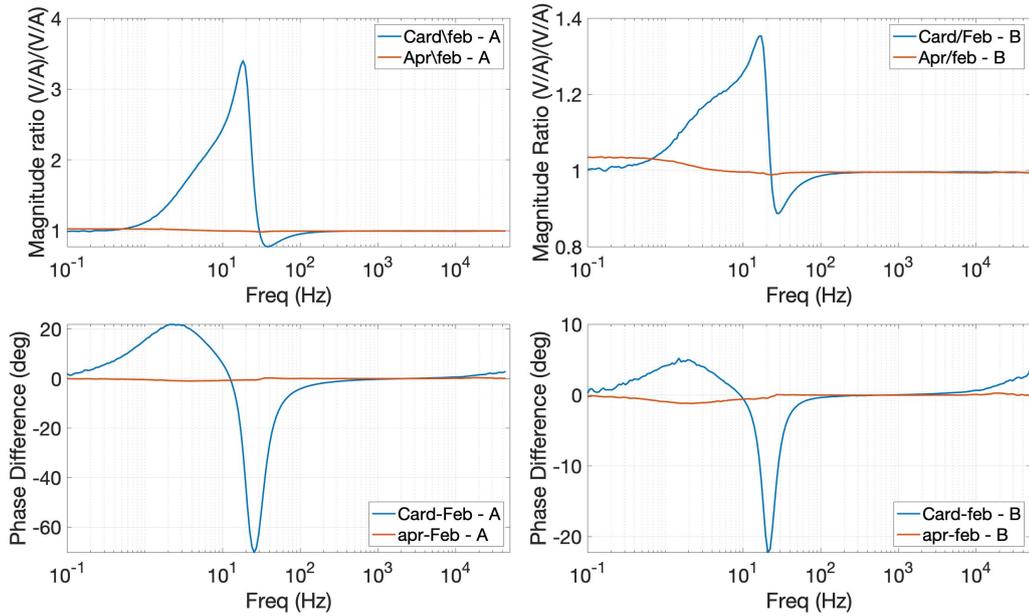


Figure 2.71: **Top:** Plots showing the ratio of the transfer functions between the Cardiff to February measurement (**blue**) and the February to April measurement (**orange**). **Bottom:** the phase difference between these measurements

This plot shows significant changes to the response of the circuit in the region of 25 Hz which contains the complex zero/pole formed by the inductor and capacitors. The most significant change is seen in channel A of the SN01 board which shows the complex pole of the circuit has shifted by almost 10 Hz.

It was proposed that these changes could be the result of the air gap of the inductors shifting during transit (see section 2.3.1 and appendix 2.2.3). In the case of channel A (the larger shift) this would correspond to a decrease in inductance value of $\sim 50\%$. In order to produce this shift the air gap would need to increase by roughly 0.1 mm. However the testing outlined in 2.2.4 implies that a slightly larger air gap adjustment would be required.

Subsequently, comprehensive tests were carried out to investigate the time dependence features of the inductor. The inductor was left with a constant DC current passing through it for 48 hours and sub 0.1% changes in inductance were observed (0.1% is the statistical error in the measurement). However even light handling of the inductor caused the value to change on the order of a few %. This is presumed to be the result of the air gap changing by a few % of 0.1 mm when the inductor is handled.

Since the noise performance also depends on the inductance value, tests were also performed on the time dependence of the DCPD noise. A circuit in Cardiff was left powered up for several days and the noise floor was remeasured as shown in figure

2.72.

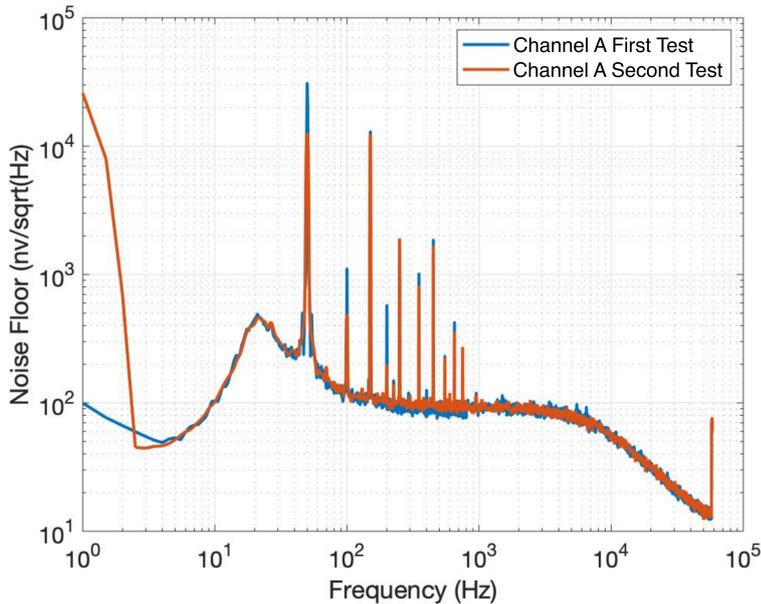


Figure 2.72: Noise floor of the audio output of the DCPD with measurements taken 3 days apart. No time-dependence in the measurement is seen.

If the inductance had changed value there would be a noticeable discrepancy around the complex pole at 20 Hz but figure 2.72 shows that the only significant change is in the pickup from the 50 Hz harmonics. These results seem to agree with the hypothesis that the inductors' mechanical properties were affected during transit to LIGO and that caused a significant change to the frequency of the complex pole in the circuit. Given that it was possible to change the inductance by 10% by applying physical force to the inductor in the lab, it seems plausible that an inductance shift of 50% could have occurred during transit from Cardiff to CALTECH. Now that the inductors are installed at the LIGO sites there should not be any more changes like this.

Another discrepancy was highlighted by comparing a measurement taken in April to the measurement taken in February. This comparison shows around a 3% change in the transfer function at low frequencies, $\sim 1\%$ change at the complex pole at 25 Hz and $\sim 0.5\%$ from 100 to 10 kHz as shown in figure 2.73. This is a smaller change than the previous measurement but was still a cause of concern for LIGO commissioning.

These discrepancies can be explained by the fact that the HAM-6 chamber began pumping down to vacuum pressure in March 2022, soon after the February data was taken. Section 2.2.8 shows a temperature model of the electronics when in vacuum. It should be noted that the modelled temperature is an overestimate as it only takes heat transfer via radiation into account and does not consider heat flow via

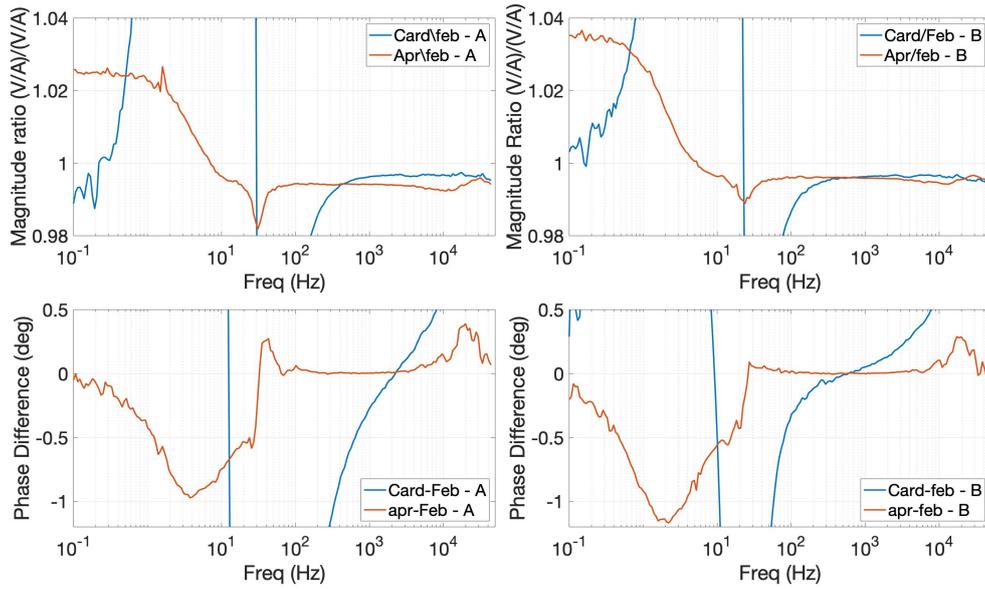


Figure 2.73: Same as figure 2.71 but scaled the y-axis to observe the change in the measurement from February to April (**Orange**).

conduction through the base. Furthermore it has assumed a 3 W power dissipation in the circuit when in reality the power drawn by the op amps is 2.7 W. Figure 2.74 shows a model of the DCPD frequency response at various temperatures.

For the model in figure 2.74 the resistors in the audio path all have a temperature coefficient of 25 ppm/°C and the inductor L_4 was changed following the data in figure 2.43 that shows a roughly 0.01 %/°C change in the inductance value. The film capacitors temperature coefficient was taken as 30 ppm/°C. At low frequencies the fraction of the photocurrent is set by the ratio of the resistance of R_8 to the ohmic DC resistance of the copper wire in inductor L_4 . The change in the copper resistance can be determined using:

$$R = R_{ref} (1 + \alpha_{cu} (T - T_{ref})) , \quad (2.46)$$

where R_{ref} is the resistance at some reference temperature, T_{ref} and α_{cu} is the temperature coefficient of copper which is 0.00393 Ω /°C.

If the DC response is used as a temperature sensor then it can be estimated that the temperature of the readouts has increased by 6.4°C. The change around the complex pole is slightly larger in the measurement compared to the model and there appears to be a 0.5% drop in measurement across all frequencies. The likely explanation is that this is due to some difference in the setup when taking the measurement at LIGO. Now that the HAM-6 chamber has pumped down to vacuum the temperature of the readouts is stable and no more changes to the transfer function should be seen.

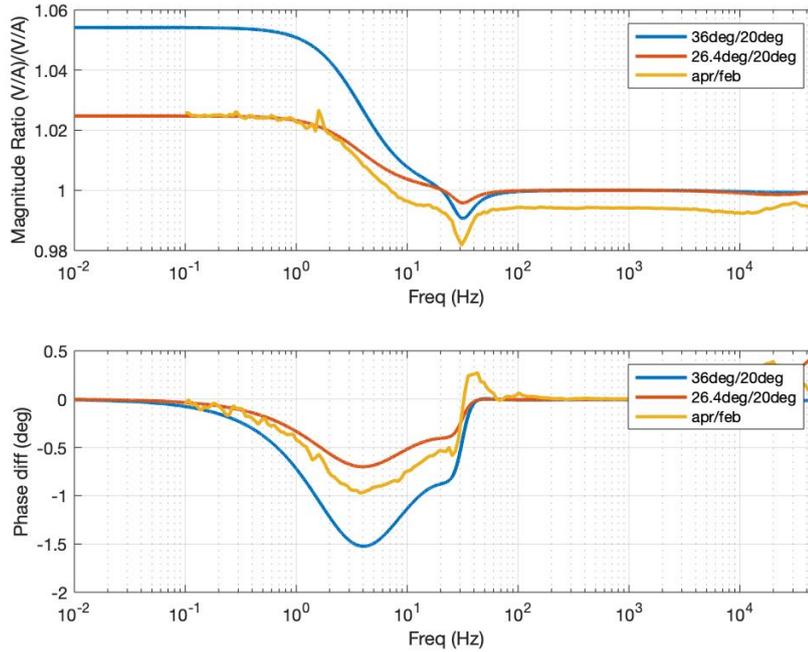


Figure 2.74: **Yellow**: observed change in circuit response from the LIGO measurement of Channel A in SN01. This is the same data shown in the left plots in figure 2.73. **Blue**: LISO model of the effect of a 16°C temperature change of the DCPD readouts. **Orange**: LISO model of a 6.4°C temperature change.

LIGO Noise Measurements

The noise levels of the DCPDs were measured several times after their installation. The most direct measurement was taken by using a breakout box at the DCPD output just before the signals were sent onto a separate whitening stage. This measurement is shown in figure 2.75.

On both channel A and B the new dark noise is a factor of 6.5 below the previous dark noise measurement at 1 kHz. At 100 Hz channel B is a factor of 5 lower whereas channel A is only a factor of 3.2 lower due to the aforementioned change of the inductance value of L_4 .

Measurements of the dark noise were also carried out using the LIGO CDS system which records the noise values after a separate whitening stage and digital filtering has been applied. A comparison of the CDS measurement and the signal analyser measurements is shown in figure 2.76:

In order to convert the CDS data into an input referred current noise, several inverting filters are applied and are shown in appendix 3.4.4. These filters are: counts to volts (ADC have units of ‘counts’), anti-whitening (accounting for the whitening stage that comes after the DCPD) and volts to amps (this is simply the transfer function of the DCPD electronics).

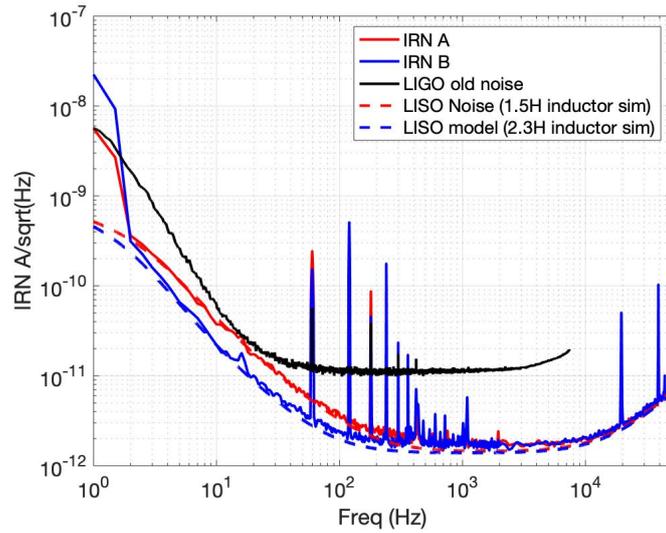


Figure 2.75: **Red solid:** input referred noise on DCPD SN01 channel A. **Red dashed:** LISO model of the IRN using a 1.5 H inductor. **Blue solid:** input referred noise on DCPD SN01 channel B. **Blue dashed:** LISO model of the IRN using a 2.3H inductor. **Black:** input referred noise of the previous DCPD readout. The noise of channel B closely matches the expected noise for a 2.3 H inductor. The noise of channel A is slightly worse at frequencies < 100 Hz and is explained by a shift of the inductance from ~ 2.3 H to 1.5 H.

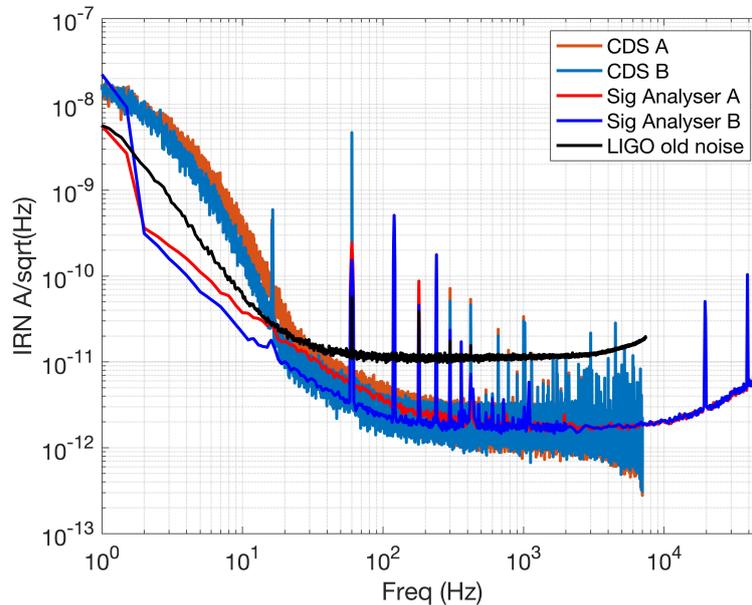


Figure 2.76: **Dark red:** channel A of SN01 DCPD measured using the signal analyser. **Dark blue:** channel B of SN01 DCPD measured using the signal analyser. **Orange:** Channel A of SN01 measured using the LIGO CDS. **Light blue:** Channel B of SN01 measured using the LIGO CDS. **Black:** the dark noise of the old LIGO readout for reference.

The data taken using the CDS system does not have as much averaging as the data taken using the signal analysers but the noise floor at 1 kHz for both channels is consistent between the CDS and signal analyser measurements. At low frequency (< 5 Hz) however, there appears to be more than a factor 10 in the noise recorded using the CDS measurements. It is suspected that this noise is the result of either additional low-frequency noise in the whitening filter (the whitening filter is an additional electronics stage after the DCPD and before the ADC) or from incorrect inverse filters definitions. The measurements from the signal analyser were taken using a breakout box directly at the output of the DCPD readouts and did not include the whitening filter or the inverse filter definitions. For this reason they are a more reliable representation of the dark noise of the readouts and are therefore used instead of the CDS measurements in subsequent discussions. I am currently in contact with LIGO commissioners and am sure the discrepancy in noise will be swiftly resolved so that the CDS system can take reliable measurements of the low-frequency noise.

Section 2.4 shows a projection of the new LIGO DARM sensitivity now that these DCPD readouts are installed.

2.3.9 VIRGO Prototype

The DCPD readout is not only applicable to LIGO, but can be adapted to a wide variety of experiments with ultra-low noise requirements. One such example illustrated here is a prototype design for the VIRGO detector.

The VIRGO detector uses a lower DC photocurrent of 2 mA rather than the 20 mA used at LIGO. This reduces the shot noise level in the readout by a factor of $\sqrt{10}$ according to equation 2.36 which brings the shot noise level closer to the dark noise. However, due to the lower photocurrent it is possible to increase the component values to reduce the dark noise even further than the LIGO readout. A 2 mA photocurrent through the load resistor means the resistor can be increased from 10 k Ω to 100 k Ω which would reduce the dark noise from this source by a factor $\sqrt{10}$. Increasing the load resistor will of course increase the gain on other noise sources, specifically the voltage noise from the first transimpedance stage will play a significant role at low frequencies. To counter this, the inductance of L_4 can be increased from 2.6 H to 10.4 H to reduce the gain on the op amp voltage noise. In some setups, this would require a significantly larger and heavier inductor, however the benefit of a low photocurrent at VIRGO is that the wire windings can simply be doubled around the mu-metal core with no significant risk of core saturation. The inductor used in the VIRGO-style readout is shown in section 5.5.

Whilst the LIGO sensitivity region is limited by low frequency alignment noise below ~ 20 Hz, the VIRGO suspension system exhibits less alignment noise and is only

limited at frequencies below 10 Hz; therefore the complex zero/pole can be formed at 11 Hz without risk of saturation of the readout from low frequency noise. Other components can then be adjusted appropriately to keep the 10 kHz double pole as before. The frequency response of the VIRGO-style readout is shown in figure 2.77.

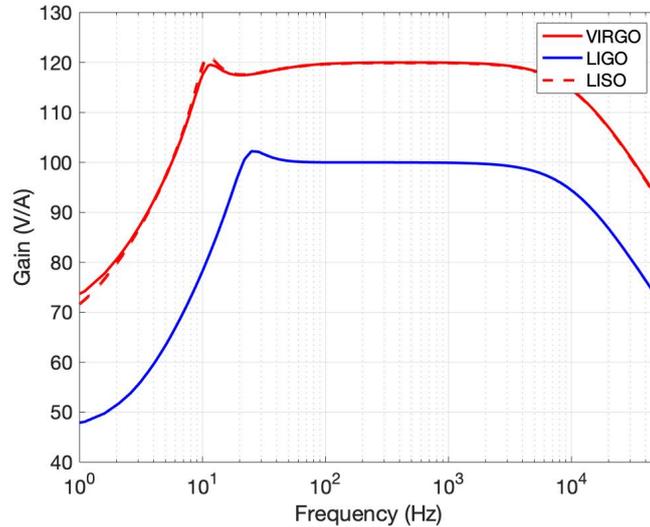


Figure 2.77: **Red solid:** Measurement of the proposed transfer function for a VIRGO-style experiment. **Red dashed:** LISO model. **Blue:** The measurement from the LIGO readout now installed at Livingston. The increased gain in the VIRGO readout is due to the transimpedance resistor increasing from 10 k Ω to 100 k Ω . The lower frequency complex pole is due to the inductor L_4 increasing from ~ 3 H to ~ 11 H.

Measurement of the voltage noise floor from the circuit in figure 2.78 showed the noise floor between 100 Hz – 10 kHz was 200 nV/ $\sqrt{\text{Hz}}$ higher than expected from the LISO model.

Initially it was suspected the higher noise level was the result of the 11 H inductor having additional Barkhausen noise, however the purple trace shows a measurement of the circuit noise with the inductor removed. The voltage noise floor below 100 Hz is of course lower because the inductor no longer affects the low frequency gain of the readout, however the noise at 100 Hz – 10 kHz is the same as when the inductor is connected, implying that the additional noise is not coming from the inductor.

Figure 2.79 shows the LISO model of the key components contributing to the noise floor from figure 2.78.

Due to the transimpedance resistor increasing from 10 k Ω to 100 k Ω the dominant noise contribution between 100 Hz – 10 kHz is now the current noise from the OPA2210. It is suspected that the manufacturers specification of 0.5 pA/ $\sqrt{\text{Hz}}$ is an underestimate and that a higher current noise is causing the measured noise floor increase seen in figure 2.78. If the current noise is actually a factor of 2 higher

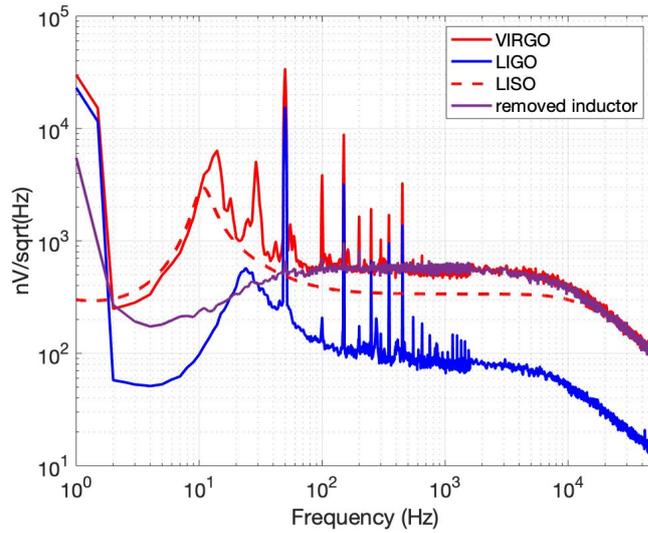


Figure 2.78: **Red solid:** Measurement of the voltage noise at the audio output of the VIRGO style readout. **Red dashed:** A LISO model of this circuit that shows a slightly lower noise floor from 100 Hz to 10 kHz. **Purple:** A measurement of the VIRGO-style readout with the 11 H inductor removed. The higher noise floor is still visible at 100 Hz to 10 kHz. **Blue:** A measurement of the LIGO voltage noise for comparison. Note that for the LISO model the op amp current noise was taken to be $0.5 \text{ pA}/\sqrt{\text{Hz}}$ as per the manufacturers data sheet.

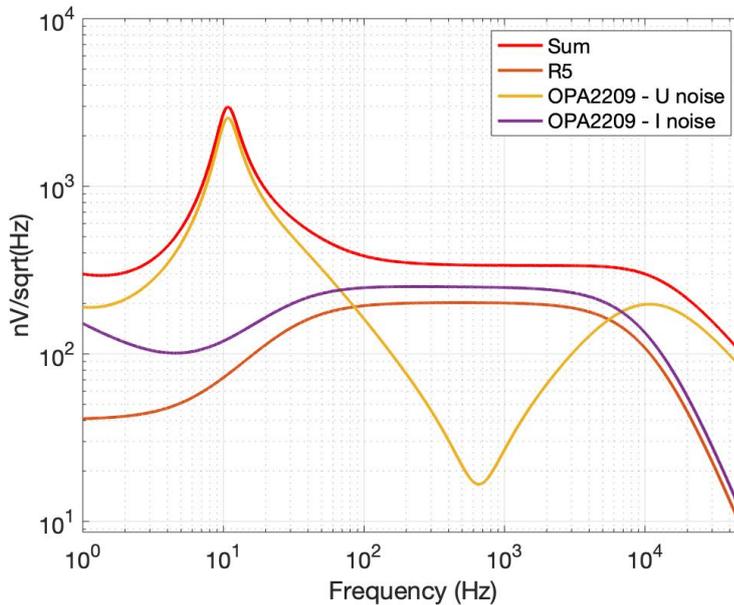


Figure 2.79: **Red:** LISO model of the sum of all noise sources in the readout. Also plotted are the noise contributions from the dominant sources: the current (**Purple**) and voltage noise (**Yellow**) from the op amp in the first transimpedance stage and the transimpedance resistor (**Orange**). For this plot the current noise is taken from the manufacturers data sheet to be $0.5 \text{ pA}/\sqrt{\text{Hz}}$.

than the data sheet value for the OPA2210 then the LISO model closely matches the measurement. The IRN of the measurement from figure 2.78 is shown below alongside a LISO model that uses a current noise value of $1 \text{ pA}/\sqrt{\text{Hz}}$ rather than the $0.5 \text{ p}/\sqrt{\text{Hz}}$ value used in figure 2.79.

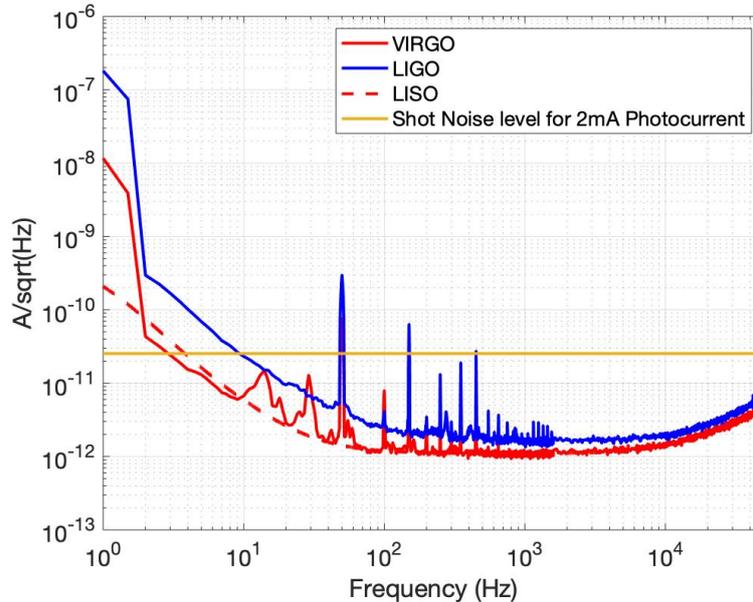


Figure 2.80: **Red solid:** IRN measurement of the VIRGO-style readout. **Red dashed:** LISO model of the IRN using a higher \tilde{I}_n level for the OPA2210 op amp of $1 \text{ pA}/\sqrt{\text{Hz}}$. **Yellow:** un-squeezed shot noise level for a 2 mA photocurrent. **Blue:** IRN of the LIGO readout for comparison

By modelling the op amp with a higher current noise the measurement of the VIRGO-style readout closely matches the model. Even with the higher current noise from the op amp, an improvement is seen in the IRN across all frequencies compared to the LIGO readout, including a reduction at the lower complex pole at 10 Hz. The overall IRN of the VIRGO-style prototype is a factor 22 lower than the shot noise level for a 2 mA photocurrent whereas the LIGO readout is a factor 16 below the 2 mA shot noise level.

The feature at 30 Hz is not currently understood but it is suspected that this is some resonance feature coming from the 11 H inductor. This version of the inductor has a slightly different construction from the other inductors and may result in a resonance at this lower frequency.

A possible extension for this project would be to investigate alternative op amps for the first transimpedance stage that have a lower current noise density. One potential is to replace the OPA2210 with an OPA165x. This FET-input op amp has a slightly higher voltage noise at $2.9 \text{ nV}/\sqrt{\text{Hz}}$ so would have increased noise around 10 Hz, however the current noise density is only $5 \text{ fA}/\sqrt{\text{Hz}}$ which would

result in about $2\text{ nV}/\sqrt{\text{Hz}}$ at the circuit output between $100\text{ Hz} - 10\text{ kHz}$, far below the contribution of the transimpedance resistor at $200\text{ nV}/\sqrt{\text{Hz}}$. It is also worth considering whether the second amplification stage could be removed since there is the additional gain coming from the higher transimpedance resistor in the VIRGO-style readout.

This prototype has shown a significant reduction in the dark noise of the readout and is a factor 22 below the shot noise level seen at VIRGO. Such a design may be of interest to the VIRGO collaboration in the future, however for the upcoming observing run the VIRGO collaboration have installed a different readout scheme that simply increases the impedance of the transimpedance resistor. The issue with this is that, due to the voltage limitations outlined in section 2.2, there is a limit to how high the impedance of the resistor can be increased which then limits the overall noise reduction. If further noise reduction is desired by the collaboration then the design presented here has clear improvements on the current design, with room for further improvements in the future.

2.4 New LIGO Sensitivity Projection

As shown in figure 2.75 in section 2.3.8 there is a factor 6.5 improvement in the dark noise going from the old readout to the new readout. Due to other noise sources this of course does not translate directly to a factor 6.5 improvement in DARM sensitivity as shown in figure 2.81.

The noise budget modelling tool GWINC [88] was used to model the noise sources for advanced LIGO in figure 2.81. The noise for the GWINC model is in some cases slightly lower than the noise of the measurement shown in figure 2.10 due to GWINC being a model of the best-case noise floor rather than a real measurement but is still a good estimate for the noise seen in O3.

Figure 2.81 shows a model of the DARM sensitivity of the LIGO detector during O3 with and without the new DCPD readout. The contribution to the total sensitivity is most significant over 600 Hz and improves the DARM sensitivity by about 4% at 1 kHz (see figure 2.83). Note that the laser amplitude noise is not included in this noise budget and based on figure 2.10 it is expected that the laser amplitude noise will exceed the dark noise at around 3 kHz .

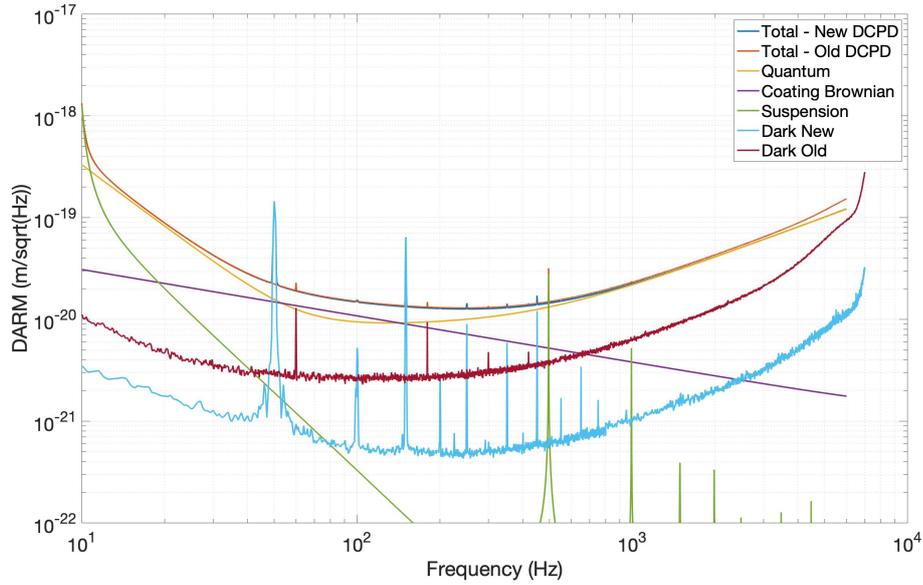


Figure 2.81: The DARM sensitivity based on GWINC models for the dominant noise sources. **Green:** suspension noise. **Purple:** Brownian noise on the mirror coating. **Yellow:** quantum noise including radiation pressure noise and shot noise. **Maroon:** the dark noise of the DCPD readout used in O3. **Cyan:** the dark noise of the new DCPD readout. **Blue:** The sum of all noise sources using the new DCPD dark noise. **Orange:** the sum of all noise sources using the old DCPD dark noise. For simplicity some noise sources from GWINC are omitted from the plot but they were still used in the calculation of the sum of all noise sources. Note that laser amplitude noise is not included on the GWINC model. This is because the plot of laser amplitude noise from figure 2.10 comes from injection measurements rather than modelling.

One of the main motivations in reducing the dark noise was to provide further headroom for the LIGO detector in the case of improved squeezing levels. Figure 2.82 below shows a projection of the improvement to the LIGO detector as squeezing levels increase.

An easier way to visualise the improvement in the DARM sensitivity due to the new DCPD readout is to plot the ratio of the DARM sensitivity with the old readout over the DARM sensitivity with the new readout as shown in figure 2.83.

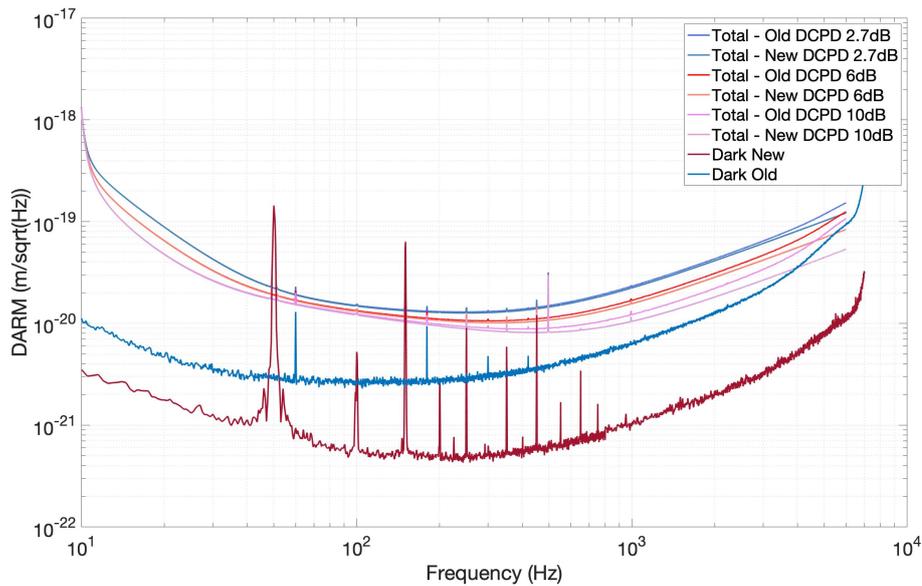


Figure 2.82: Model of the DARM sensitivity for various levels of observed squeezing with the old DCPD and new DCPD. **Blue**: 2.7 dB, old DCPD. **Blue Steel**: 2.7 dB, new DCPD. **Red**: 6 dB, old DCPD. **Salmon**: 6 dB, new DCPD. **Violet**: 10 dB, old DCPD. **Plum**: 10 dB, new DCPD. Other noise sources were omitted from the plot for simplicity but were still used in calculating the overall noise floor. The GWINC noise models of all other sources were kept the same in each projection.

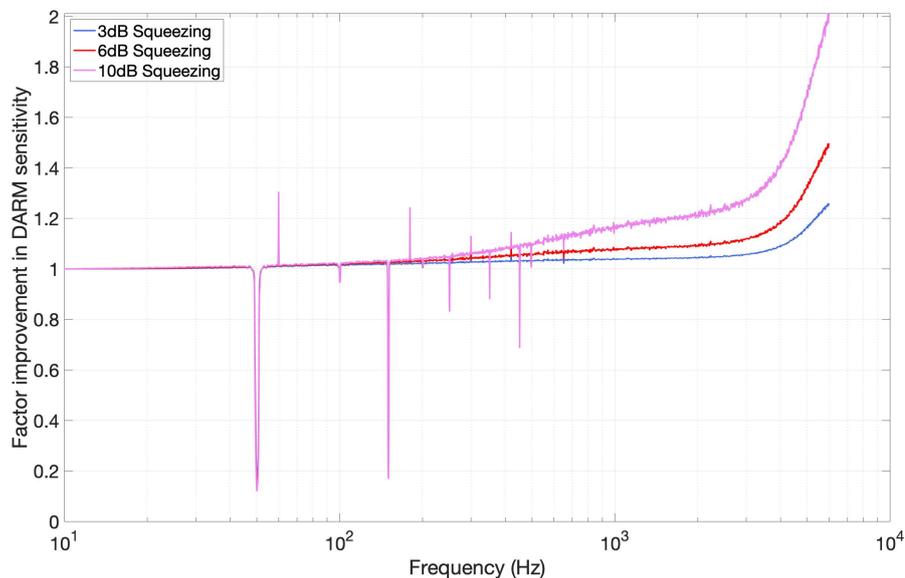


Figure 2.83: **Blue**: The ratio of DARM sensitivity for 2.7dB squeezing. **Red**: 6dB squeezing - the squeezing level target for O4. **Pink**: 10dB squeezing - the squeezing level target for O5. The spikes are due to the pickup from mains frequency, this has been shown to be reduced when the sum of the two readouts are considered. The improvement beyond 3 kHz will not be as significant because laser amplitude noise will likely exceed the dark noise level (see figure 2.10).

The benefits that were outlined back in section 2.1.6 are now readily apparant. In the frequency region 600 Hz - 3 kHz where the dark noise is the second most dominant noise source, the DARM sensitivity is shown to significantly improve when the new detector readout is used and the improvement becomes more significant as squeezing levels improve. Table 2.3 shows the percentage improvement in DARM sensitivity at 2 kHz.

Squeezing Level	Projected DARM Improvement
2.7 dB	4 %
6 dB	10 %
10 dB	20 %

Table 2.3: The factor improvement in DARM sensitivity at 2 kHz from figure 2.83

The DARM improvement will potentially be even larger than we have shown in figure 2.83 since a model for the noise budget during O3 was used to make all projections and everything but the shot noise and dark noise was kept constant. In reality other noise sources will be reduced during the A+ upgrade. As these other noise sources get reduced the role of the dark noise also becomes more significant. Consider for instance if the new readouts were not implemented in time for O4 but the thermal and residual gas noise (see figure 2.10) were significantly reduced. The dark noise would then be the second most dominant noise source at frequencies lower than 600 Hz and reducing it would improve DARM sensitives across a broader range of frequencies.

2.5 Future Work

The detectors have been installed at LIGO Hanford and Livingston in time for the O4 detection run and demonstrate a reduction in the detector dark noise of a factor 6.5. There are however, areas of improvement that could be implemented for a future version of the readout that could potentially be installed for an O5 version of the DCPD.

2.5.1 Test Relay Improvements

There is a 100 k Ω resistor used in the test input relay to convert an input voltage into an approximate input current. This particular resistor was never measured independently of the other components in the readout which makes calibration of the audio readout to the sub – 1 % level slightly complex. It may be possible to use a stable injection of DC light to measure the DCPD audio transfer function but then one is reliant on the noise on the light injection being below 1 %. Future versions of the readout should have the input resistors measured to a sufficient accuracy that

audio calibration is a straightforward task.

In a similar vein there should be an alternate way to take the transfer function of the RF output. This could be done either using a higher bandwidth test relay or having an independent test relay for the RF transfer function.

2.5.2 DC Signal Injection

Further investigation is required into the problem outlined in section 2.3.7. Magnetic tests indicate that the increase in audio noise upon DC signal injection is due to the power supply noise used as the DC current source, but this needs to be verified.

2.5.3 Inductor Changes

The mu-metal core inductor produced by Sowter was successfully used in the DCPD readouts; however there are elements of the design that could be improved for a potential O5 version of the readout.

Core Material: As mentioned above it is currently still an open question as to where the increase in noise originates from. In the case that it is due to the core material then a new material should be considered. This was the subject of some investigation by Grote et al. in 2016 [69] and they were unable to observe any excess noise in the mu-metal core. This conclusion was the motivation for selecting mu-metal as the inductor core; however Barkhausen noise is strongly dependent on the manufacturing process and the impurity composition will vary between batches.

It is possible that more impurities are present within this batch of inductors which results in a higher noise. Alternative core materials were not investigated during this project.

Inductor Assembly: The assembly of the inductor was shown in section 2.2.4 and it was mentioned that it was not possible to increase the air gap beyond approximately 0.3 mm and that there was some variation of the individual sheets across the air gap. The adjustment of the air gap is also time-consuming as the whole inductor must be taken apart and then reassembled. A convenient feature for a future design would be one that allows easy adjustment of the air gap for example by the tightening of a screw on top of the inductor. Ideally this feature would also better secure the air gap so that shifts like those outlined in section 2.3.8 are not seen.

2.5.4 Moving the Electronics and Introducing Cryogenics

A potential improvement for a future design for O5 design would be to move the readout electronics outside the vacuum chamber. This would provide benefits to heat dissipation in the readout (see section 2.2.8 and more flexibility in accessing and modifying the electronics. The downside is that the photocurrent signal from

the photodiodes would have to pass through long lines to exit the vacuum chamber and there could be significant magnetic coupling. Sufficient magnetic shielding on these lines would therefore be a requirement for this potential upgrade.

If the electronics were migrated outside the vacuum chamber there is a further - more ambitious - upgrade that would then become more feasible which is to introduce cryogenic cooling to the readout. As mentioned in section 2.1.7 it is non-trivial to introduce cryogenics to cool systems inside the vacuum chamber. However, if the electronics were moved outside the chamber, cryogenically cooling the electronics to $\sim 5\text{ K}$ could provide a further reduction by a factor of $\sqrt{\frac{300}{5}} \approx 8$ in the dark noise, following equation 2.39. Cooling the electronics to $\sim 5\text{ K}$ would also open up the possibility of using alternative inductor designs such as superconducting air core coils which could address the aforementioned mu metal inductor noise. It would still be a costly and time consuming investment for the collaboration to make and the additional improvement in dark noise may be unnecessary at this point in time. Perhaps when O5 or O6 eventually arrives, the sensitivity of the detector will be high enough that inclusion of cryogenics will not seem like such a distant a possibility.

2.6 LIGO Concluding Remarks

New photodiode readouts have been designed, built and installed at both the LIGO Livingston and LIGO Hanford sites. The new readouts exhibit a dark noise that is a factor 47 below the un-squeezed shot noise level (20 mA photocurrent) providing sufficient headroom for up to 13 dB of squeezing before the dark noise is even a factor of 10 below the shot noise. The new readout noise levels are also a factor of 6.5 lower than the previous DCPD readouts at 1 kHz. A projection of the new DARM sensitivity shows that the new readouts will result in at least a 10 % improvement in DARM sensitivity at 2 kHz during the next LIGO observing run (assuming 6 dB of squeezing). The actual improvement could be even higher depending on the other noise sources that may be reduced. This frequency region is especially interesting for binary neutron star systems as the merger frequency typically occurs in the 1-5 kHz range. Hopefully the new readouts will help the LIGO detectors probe this frequency region and uncover new insights into the mergers of binary neutron stars and the post-merger effects.

In section 2.3.9 a potential alternative to the LIGO readout was shown with a proposal for a VIRGO-style prototype. This version of the readout exhibits a factor 1.5 less noise at 1 kHz than the LIGO-style readout with a trade off of a factor 10 lower dynamic range. This noise floor is limited by the current noise of the op amp used in the transimpedance stage, and therefore a lower noise op amp could be the subject of future investigation. This prototype already demonstrates a noise reduction that may be of interest to the VIRGO collaboration as they also introduce increased levels of squeezing into the detector.

Chapter 3

ALPS

3.1 ALPS Theory

3.1.1 The Standard Model and Dark Matter

The standard model of physics has been developed over the course of many decades and is one of the most widely accepted models that describes three of the four fundamental forces, those being electromagnetic, weak and strong interactions, omitting gravity. Forces in the standard model are described via the exchange of bosonic particles. The way that this model has been tested has been the same way that any model has been tested, first by theoretical prediction and then by experimental verification. This process was used in the late 20th century with the prediction [89] and subsequent discovery of the top quark [90] [91] and most recently with the 1964 prediction[92][93][94] and 2012 discovery of the Higgs boson [95][96].

Despite its successes however, the SM of physics is incomplete and falls short in many key areas. It does not provide a full description of gravity, with physicists forced to use separate theories such as General Relativity (GR) [97] or to combine quantum mechanics and general relativity to form Quantum Gravity (QG) [98]. The standard model also cannot account for the relative abundance of Dark Matter (DM) and dark energy that exist in the universe based on indirect observations. According to the most up-to-date estimates only 5% of the matter/energy content of the universe is constituted by ordinary matter, 27% is so called ‘dark matter’ and the remaining 68% is made up the mysterious ‘dark energy’ that is inferred from observations on the expansion of the universe [99]. Predictions for dark matter significantly predate any prediction for dark energy with observations of the rotational speeds of galaxies leading many to suggest that there is some additional unseen mass exerting a higher gravitational force to hold stars in orbit [100][101]. In 1933 F. Zwicky observed that galaxies in the Coma cluster had a higher average velocity than predicted and attributed the effect to ‘Dunkle Materie’ or dark matter [102].

There have been various other sources of evidence for dark matter over the years including:

- Gravitational lensing [103] where the distortion of light, either magnification or deflection, points to additional mass coming from a DM candidate.
- Measurements of the Cosmic Microwave Background (CMB) from the WMAP [104], Planck [105] and ACT [106] experiments are consistent with the Λ CDM theory which proposes that the structure of the CMB has been determined by a Λ term for dark energy, a ‘CDM’ term for cold dark matter and the remaining terms being described by ordinary matter. Other theories suggest that an imprint was left on the CMB by DM annihilation in the early universe [107].

The term ‘dark matter’ now refers to any candidate that can explain one or more of these cosmological and astrophysical observations. Potential dark matter candidates span an astonishing 90 orders of magnitude in mass [108] from theories involving ultralight bosons [109] to primordial black holes [110]. The theories behind the various DM candidates are equally varied and as DM searches struggled to detect any evidence for a particular candidate (see the search in vain for WIMPS [111]) the search efforts turned to previously less fashionable sources. These include but are by no means limited to: various types of dark matter described by scalar fields (Dilaton/Modulus, Relaxion Halo) [112], coherent relaxation dark matter [113] and finally the axion and axion-like particles [114] around which this section of the thesis is focused upon.

Whilst discussing dark matter, it is also worth giving an honorable mention to the Modified Newtonian Dynamics theory (MOND) which attempts to explain the aforementioned phenomena without invoking dark matter at all. The theory was conceived 30 years ago and attempted to explain the problem around galaxy rotation speed by modifying the theory of Newtonian gravity [115]. Previously the work on adapting Newtonian mechanics struggled to explain the anisotropies in temperature seen in the CMB (this is explained in the Λ CDM model by the dark matter contribution), however recently a relativistic model for MOND was put forward that explains not only gravitational lensing and gravitational wave-speed measurements but also is able to account for the structure of the CMB [116]. Although some versions of MOND have been ruled out by LIGO observations [117] and the theory of dark matter is generally well accepted, this new MOND theory is exciting in how it attempts to break the status quo in the field of dark matter.

Before we continue on this discussion it is worth noting that any one individual theory for dark matter may not be singly correct and in fact it seems unlikely that that would be the case. In several of the theories we explore in section 3.1.2, it is assumed that their candidate particle is solely responsible for the dark matter

phenomena they attempt to describe. They generally make this assumption so that they can place the most optimistic limits on the particle's abundance and hence the most optimistic detection probability to try to entice experimental physicists to search for them. However, there is no reason that several DM candidate particles could not exist in parallel and it may well be that dark matter is more like a cocktail of stealthy particles waiting to be discovered.

3.1.2 Axion-like Particles

It helps a DM candidate immensely if the theory of its existence also explains some other ongoing mystery in physics. Interestingly, there are some potential dark matter candidates that arise from trying to resolve issues within the standard model itself. One such issue is known as the strong Charge-Parity (CP) problem.

Parity describes the spatial coordinate of a physical system; an inversion of the parity would transform the spatial coordinate \vec{x} into $-\vec{x}$. An inversion of a particle's charge transforms it into its antiparticle e.g. e^- into e^+ . If an interaction is symmetric in charge and parity it means the charge and parity of the system can be inverted and one will obtain the same outcome. There are cases involving Weak force interactions where it is predicted that CP symmetry will not hold [118] and indeed, experimental observation has shown this to be the case. It is also expected that strong interactions - those being the interactions underpinning Quantum Chromodynamics (QCD) - should also violate CP symmetry but experimentally this has not been observed. Although one cannot say for certain that there will *never* be a measurement of such a violation, we can say with some certainty that to date there has never been an observation of any violation of CP symmetry in strong force interactions, leading to the obvious question of why that should be.

There are several proposed solutions to the CP problem [119]; however the most famous solution comes from the Peccei-Quinn theory in 1977 [120] [121]. This theory solves the CP problem by introducing a new chiral symmetry term in the Lagrangian, $U(1)$. A consequence of this new term is the prediction of a light pseudoscalar particle, the axion [122]. The properties of the predicted axion include a very low mass ($< \text{meV}$) and a very weak interaction term with matter and light which prompts it to be a potential dark matter candidate [123]. The interaction with light shows that an axion can couple into two photons and vice versa. This is the key process that all of the axion-like particle experiments exploit in their search for this particle. In the following discussion however, one of these photons is real and the other is virtual and is provided by the magnetic field that the photon propagates through.

For the case where $\hbar\omega \gg mc^2$, the probability of axion-photon interactions in the presence of a magnetic field is symmetric [124], i.e. the probability of photon to

axion conversion is the same as axion to photon interaction and is given by:

$$P_{\gamma \rightarrow \alpha} = P_{\alpha \rightarrow \gamma} = \frac{1}{4} \frac{\omega}{k_\alpha} (g_{\alpha\gamma\gamma} BL)^2 |F(qL)|^2, \quad (3.1)$$

where ω is the angular frequency of the photon, $k_\alpha = \sqrt{\omega^2 - m_\alpha^2}$ where m_α is the mass of the axion or axion-like particle, $g_{\alpha\gamma\gamma}$ is the coupling constant between photons and axions, B is the magnetic field amplitude, L is the path length of the photons through the magnetic field and $|F(qL)|$ is the form factor which, for the cases considered in this thesis where we encounter a single homogeneous magnetic region, is defined as:

$$|F(qL)| = \left| \frac{2}{qL} \sin\left(\frac{qL}{2}\right) \right|, \quad (3.2)$$

where q is the momentum transfer to the magnetic field:

$$q = n\omega - \sqrt{\omega^2 - m_\alpha^2} \approx \omega(n-1) + \frac{m_\alpha^2}{2\omega}. \quad (3.3)$$

However this all simplifies for low mass axions in a vacuum where $n = 1$. Low values of q correspond to a low axion mass, $m_\alpha \ll 8.9 \times 10^{-4} \text{ eV}$, where we then find: $\sin\left(\frac{qL}{2}\right) \approx \frac{qL}{2}$ and $k_\alpha = \sqrt{\omega^2 - m_\alpha^2} \approx \omega$. This causes equation 3.1 to simplify to:

$$P_{\gamma \rightarrow \alpha} = P_{\alpha \rightarrow \gamma} = \frac{1}{4} (g_{\alpha\gamma\gamma} BL)^2, \quad (3.4)$$

where we can now see that the probability of a single axion \rightarrow photon conversion or vice versa is dependent only on the magnetic field strength, the distance that the photons or axion propagates through the magnetic field, and the coupling constant, $g_{\alpha\gamma\gamma}$ predicted by a particular axion model.

There are several experiments that have tried to search for axion-like particles and employed various methods to do so. The first are the helioscopes. These experiments look for axions or axion-like particles that are generated by X-ray photons from the sun passing through the solar magnetic field. By pointing an optically opaque telescope towards the Sun and having a long magnet string inside the telescope, one would have a chance to convert axions from the sun back into photons where they could be detected. Two significant helioscope experiments are the International Axion Observatory (IAXO) proposal [125] and the CERN Axion Solar Telescope (CAST) experiment [126].

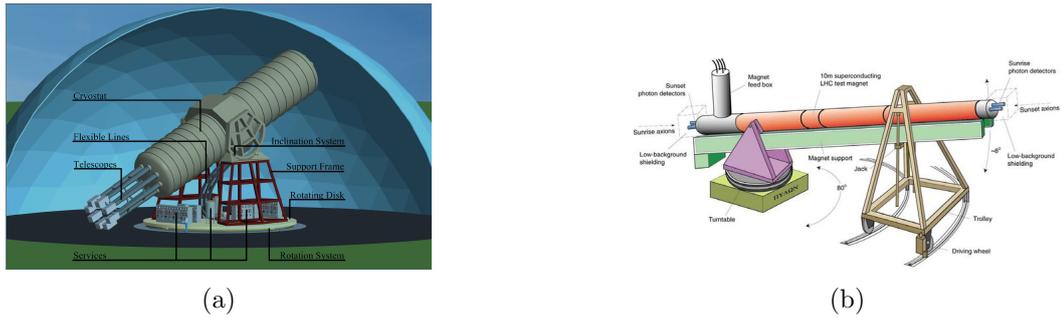


Figure 3.1: **a)** Simplified design for the IAXO proposal [125]. **b)** Simplified design for the CAST experiment [127]. These experiments rely on pointing optically opaque telescopes towards the Sun and searching for solar axions converted into photons using strong magnets within the telescope.

The next group of experiments are the haloscopes. These experiments search for axions or axion-like particles coming from the suspected dark matter halo around the Milky Way. As before, a strong magnetic field is used to convert any axions into photons where they can be detected with a suitable detector. The energy expected from these DM halo axions is $40 \mu\text{eV} - 400 \mu\text{eV}$ which corresponds to photons in the microwave region of the EM spectrum. The haloscopes therefore employ microwave resonators within the magnetic strings to increase the probability of an axion to photon conversion. Some examples of prominent haloscope searches include the proposal for the Magnetized Disc and Mirror Axion experiment (MADMAX) [128], the Axion Dark Matter Experiment (ADMX) [129] and the impressively named: A Broadband/Resonant Approach to Cosmic Axion Detection with an Amplifying B-field Ring Apparatus (ABRACADABRA) [130].

Although the sensitivities of these experiments are very high and they can map out a large section of potential ALP parameter space (see figure 3.4) there is a downside to them in that they are all model dependent. This means a successful detection relies on the axion-like particle actually being present in the region they are looking at, be that the galactic halo or the core of the sun. There is a third type of axion search which does not have this restriction however - the light-shining-through-walls (LSW) experiments.

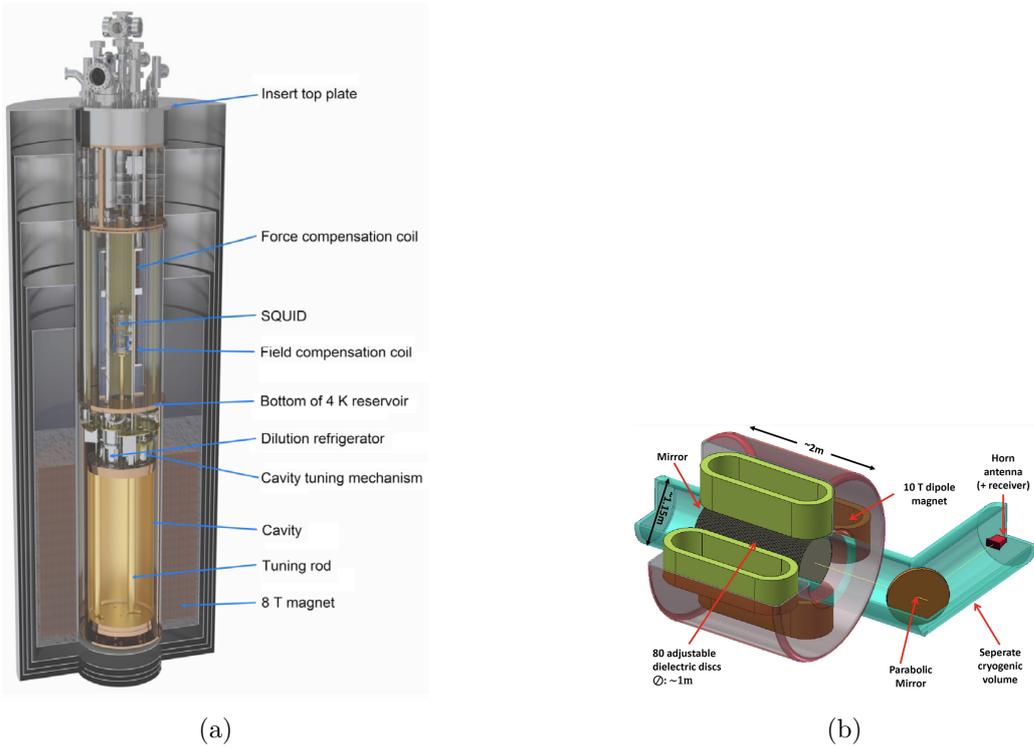


Figure 3.2: **a)** Simplified design for the ADMX experiment [131]. The microwave cavity and 8 T magnet string are shown in the lower part of the diagram. **b)** Simplified design for the MADMAX proposal [128]. The microwave cavity and 10 T magnet string are shown on the left side of the diagram. These experiments use microwave resonance cavities within strong magnetic fields to try to detect axions or axion-like particles coming from the Milky Way galactic halo.

LSW experiments are fairly easy to understand and - as the name suggests - involve shining a powerful laser beam at an opaque wall and trying to detect photons on the other side. A string of magnets line the photon path both before and after the wall. Since axions are theorised to couple to two photons in a strong magnetic field [132], for a sufficiently strong magnetic field and a high enough flux of photons on the wall it should be possible for some photons to convert into axions before the wall. The axions would then propagate through the wall and pass through the second magnet string which would cause some of them to reconvert into photons on the far side of the wall where they can be detected using a suitably sensitive detector. The photons therefore appear to have propagated straight through the opaque barrier.

Examples of LSW experiments include the Optical Search for QED Vacuum Birefringence, Axions, and Photon Regeneration (OSQAR) [133] and the Any Light Particle Search (ALPS) [134], the latter of which, of course, is the focus of this chapter of the thesis.

Equation 3.4 shows the probability of a single axion \rightarrow photon interaction or vice versa. Since the probability for the inverse interaction is identical, the probability

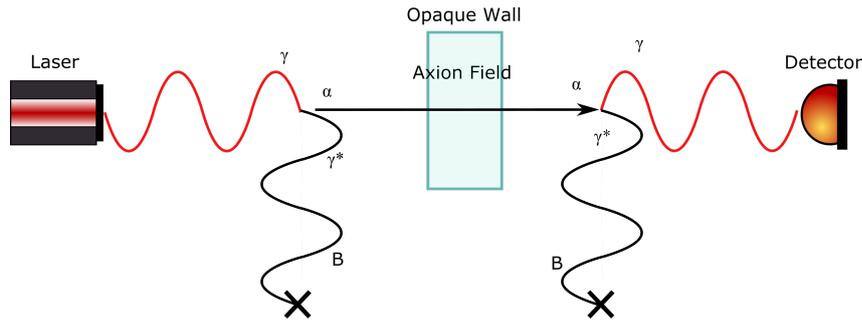


Figure 3.3: Simplified diagram of a typical LSW experiment. Real photons, γ , are provided by a high power laser and propagate towards an optically opaque barrier. Strong magnetic fields perpendicular to the propagation direction of the photons provide virtual photons γ^* . The interaction of the real and virtual photons generates axions which propagate through the opaque wall and are reconverted to photons in the magnetic field on the other side of the wall. The regenerated photons are then detected with a suitable detection method.

of a single regenerated photon in a LSW experiment is:

$$P_{\gamma \rightarrow \alpha \rightarrow \gamma} = P_{\gamma \rightarrow \alpha} P_{\alpha \rightarrow \gamma} = P_{\gamma \rightarrow \alpha}^2. \quad (3.5)$$

The rate of photons reconverted on the far side of the wall is simply the probability of a single regenerated photon $P_{\gamma \rightarrow \alpha \rightarrow \gamma}$ multiplied by the rate of photons incident on the left hand side of the barrier:

$$N_{reg} = N_{laser} \frac{1}{16} (g_{\alpha\gamma\gamma} BL)^4. \quad (3.6)$$

In reality, however, such a setup would be far too insensitive to test for axions. For an incident power of 50 W, a coupling factor of $g_{\alpha\gamma\gamma} = 2 \times 10^{-11}$, a magnetic field $B = 5.3$ T and a magnetic path length of $L = 120$ m the regenerated photons would be on the order of 1 photon/70,000 years!

Before delving deeper into the more detailed design of the experiment to explain how the sensitivity of the ALPS IIc detector is increased it is worth pausing to consider the parameter space covered by ALPS and the other significant axion searches mentioned in 3.1.2 as well as the axion models being tested. Although the original prediction for axions stems from solving this strong CP problem observed in QCD there are more recent models for axions stemming from astrophysical observations, specifically the transparency of the universe to gamma ray radiation [135]. The amount of TeV radiation we observe from certain galactic nuclei is higher than expected and could be explained by the conversion of these high energy photons into axions in the presence of the strong magnetic fields within the galactic medium. These axions would then be able to pass through the universe mostly uninhibited

and be reconverted into gamma rays as they enter the magnetic fields surrounding the Milky Way galaxy where they are subsequently detected. These theories provide new areas of parameter space that axion-like particles could be hiding in and that experiments are trying to probe as shown in figure 3.4.

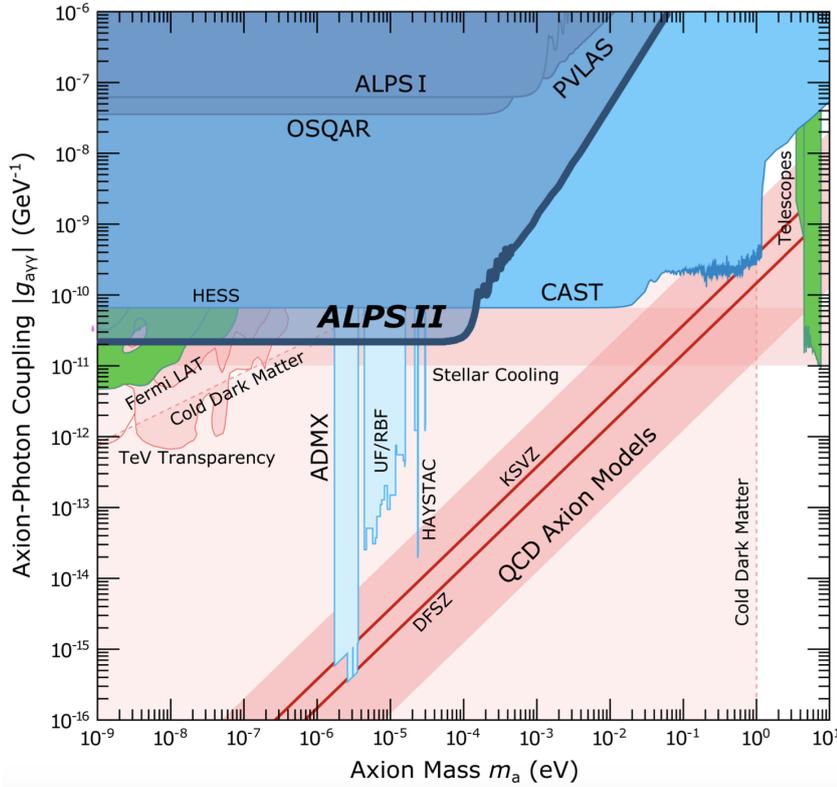


Figure 3.4: A plot of the coupling constant $g_{\alpha\gamma\gamma}$ plotted against axion mass m_α for models predicting ALPS from various sources (**red**) and the experiments' limits on this parameter space (**blue**). This plot was taken from the design document of the ALPS II experiment [134]. These axion models are not explored further in this thesis but are shown in [136–146]. As mentioned in [134], exclusion limits rely on untested assumptions such as axions or axion-like particles making up 100% of the dark matter in the milky way and structures of galactic and intergalactic magnetic fields.

From figure 3.4, it appears that the ALPS IIc experiment will not be sensitive enough to probe the traditional model for QCD axions but will instead test models involving astrophysical observations.

That being said there is one final model for the QCD axion which is not shown in figure 3.4 (due to how recently it has been introduced) but that predicts significantly higher axion coupling than the standard QCD model [147]. In this paper Sokolov and Ringwald outline a new model for the QCD axion that predicts a $g_{\alpha\gamma\gamma}$ coupling constant of around 2×10^{-11} for axion masses of m_α of 10^{-4} eV, just inside the sensitivity of the ALPS II detector and meaning that the ALPS II sensitivity may be able to at least somewhat probe the QCD model for the axion.

3.2 ALPS Experiment Design

The ALPS experiment utilises two high finesse cavities both before and after the opaque wall. In the Production Cavity (PC) the circulating power of the input beam is amplified which increases the number of potential photon \rightarrow axion interactions. In the case of the Regeneration Cavity (RC) the optical cavity enhances the EM field which in this case is the vacuum state. This in turn increases the re-conversion probability of axions \rightarrow photons on the far side of the barrier based on the optical cavity build up parameter β_{RC} .

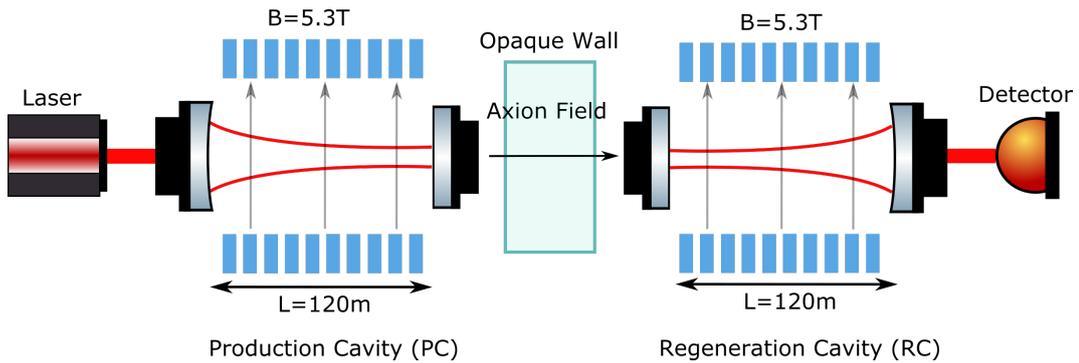


Figure 3.5: Diagram of the ALPS IIc experiment. The feed laser is resonant within the production Cavity (PC) which increases the number of photon \rightarrow axion interactions. The regeneration cavity increases the likelihood of axion \rightarrow photon conversion.

The equation for the photon rate on the detector, equation 3.6, is now slightly more complicated as the number of photons circulating in the magnetic field in front of the opaque wall is dependent on the finesse of the PC; correspondingly the enhancement of the regenerated signal is dependent on the power build up factor of the RC. As shown in figure 3.5 a mismatch in eigenmode alignment will prevent the regenerated photons from resonating within the regeneration cavity, and therefore there is also a new dependence on the relative alignment of the cavities' eigenmodes with respect to each other. The equation for the regenerated photon rate therefore becomes:

$$N_{reg} = \frac{\eta}{16} (g_{\alpha\gamma\gamma} F(qL) BL)^4 \frac{P_{PC}}{h\nu} \beta_{RC}, \quad (3.7)$$

where η is the factor for the relative cavity eigenmode overlap, h is Planck's constant, ν is the optical frequency, P_{PC} is the circulating power in the production cavity and β_{RC} is the cavity build up parameter for the regeneration cavity.

The initial science goal for ALPS IIc is to have 150 kW power in the PC, a power build up of $\beta_{RC} = 40,000$ and an eigenmode overlap of at least 95% [134]. This gives a sensitivity to coupling constants as low as $g_{\alpha\gamma\gamma} = 2.45 \times 10^{-11}$ which corresponds to a regenerated photon rate of $2.8 \times 10^{-5}/\text{s}$. In the final configuration, ALPS IIc is designed to test coupling constants as low as $g_{\alpha\gamma\gamma} = 2 \times 10^{-11}$.

In the ALPS IIc experiment long observation times on the order of 1.7×10^6 s (20 days) are required in order to reach this design sensitivity. During these timescales it is not unusual for drifts in the pointing angle of the input beam to occur due to thermal effects and from acoustic motion of the input optics (specifically the two large 3-inch pico-motor steering mirrors used in the input optics). These effects will both contribute to power noise within the cavity, the latter at higher frequencies and the former at very low, approximately DC, frequencies. As of writing this thesis it is still unclear exactly how much DC drift and how much higher frequency noise is expected but this project set about designing and implementing an active feedback alignment system to suppress the pointing noise of each beam into their respective cavity and which is described in section 3.2.4.

3.2.1 ALPS IIc Detection Schemes

A photon rate of 2.8×10^{-5} /s is roughly 2 photons/day and is too low for a traditional photodetector to detect; instead a specific readout is required that can measure individual photons. Two methods of detecting the regenerated photons have been proposed for use in the ALPS IIc experiment. The first is a heterodyne technique which locks the regeneration cavity using a laser that has a frequency offset to the laser resonant within the production cavity. Any regenerated photons within the regeneration cavity will form a beat note with the frequency offset carrier in the regeneration cavity which can be detected via demodulation [148].

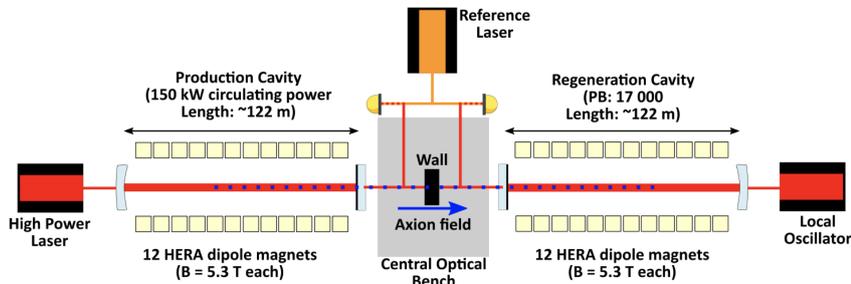


Figure 3.6: Heterodyne detection scheme for the ALPS IIc experiment [148]. A LO with a frequency offset from the PC carrier is used to lock the RC. A Phase Lock Loop (PLL) is maintained between the two cavities by locking the PC carrier to a reference laser (RL) that is in turn locked to the LO. The beat note between regenerated photons (at the PC carrier frequency) and the LO is detected via demodulation (demodulation scheme not shown in the diagram).

The other detection method is to use a superconducting single photon detector known as a Transition Edge Sensor (TES) [149]. The TES detector at its core is a superconducting metal that is held at the boundary between its superconducting and normal states. Any photon that is incident upon the TES detector will be absorbed and cause a slight change in temperature of the metal. This temperature change will cause a corresponding change in the impedance of the metal which can

be converted into a detectable voltage. In the TES readout a laser that has half the wavelength of the PC cavity light is used to keep the RC on resonance.

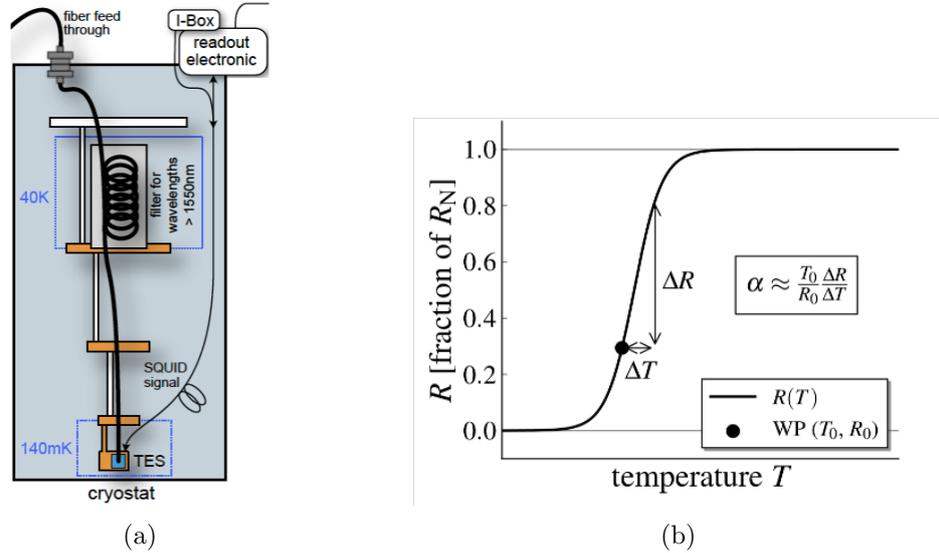


Figure 3.7: **a)** Schematic of the TES detector [149]. Regenerated photons are sent to the detector via fibre optic cable. **b)** Model of the TES detector resistance vs. temperature. A photon incident on the detector causes a temperature change which in turn causes a change in the resistance of the superconductor.

3.2.2 Optical Table Design

This section goes into detail on the optics of the ALPS IIc experiment's end tables and introduces some nomenclature commonly used in the ALPS experiment. The optical table that controls the input beam to the Production Cavity (PC) is to the North Left of the experiment and so is typically referred as the NL table. Similarly the optical table containing the input optics to the Regeneration Cavity (RC) is to the North Right of the experiment and is referred to as the NR table. The optical table between the PC and RC is referred to as the Central Optical Bench (COB).

The COB serves not only to align the eigenmodes of the cavities with respect to each other but it also contains the optics required to set up a phase-lock-loop (PLL) between the high power laser (HPL), the Reference Laser (RL) and the local-oscillator (LO) that is resonant within the RC and has the optics used to measure the signal from any regenerated photons. This is too deep of a discussion for the purposes of this thesis and since it was outside of my own personal contribution to the ALPS II experiment I instead refer the reader to [150] for details on the COB.

Due to the symmetry of the ALPS experiment, the NL and NR table designs are almost identical and in this section I will cover the overall design of the NL table. The specifics of each design can be found in the final design documents for each table [151] [150]. My personal contribution to the optics design begins just after the PMC;

however a basic overview of the other systems is provided here for completeness. The table layout for the NL table is shown in figure 3.8.

With reference to figure 3.8: **Top right**: the beam from the input laser passes through the amplifier and out the light tight box going left in the diagram. **Top middle**: it travels through a mode matching telescope LT L24 and LT L4 which mode match it to a PMC cavity (**top left**). There is then a collimating telescope to pass it through a Faraday isolator LT FI1 followed by a Gouy phase telescope to ensure 90° phase separation between two auto-alignment steering mirrors (**middle right**) (see section 3.2.4 for details) and then the mode matching telescope (**middle**) to ensure the mode matching of the input beam to the optical cavity (**right**). **Bottom middle**: a pick-off from the cavity reflection/cavity leak power is used to generate the auto-alignment error signal (detailed in section 3.2.4). **Bottom right**: a HeNe green laser is used for the initial alignment of the cavity since it is easier to roughly align the beam using visible light rather than 1064 nm.

The initial beam comes from a Mephisto2000 (coherent) laser with a 2 W power output. Details on the laser are shown in table 3.1. Phase Modulated (PM) sidebands are applied using two EOMs (LA EOM1 and LA EOM2). Only one pair of sidebands are needed for the length of the cavity and the auto-alignment error signals and are chosen to be at 9 MHz. The other pair of sidebands are optional in locking to the optional Pre-Mode Cleaner (PMC). A single lens (LA L1) is used to ensure there is minimal clipping on the two EOM's apertures. A Faraday isolator ensures there are no back reflections coupling back into the master laser. Precise mode matching into the neoLASE Nd:YVO₄ amplifier is achieved using the telescope formed by LA L3 and LA L4.

Parameter	Specification
1. Type of laser	Nd:YAG NPRO, Nd:YVO ₄ amplifier
2. wavelength	1064 nm
3. output power (after amplification)	> 50W (70W max)
4. power in higher order modes	< 5W
5. polarization extinction ratio	100:1 in the vertical plane
6. relative power fluctuations	< $10^{-3}/\sqrt{\text{Hz}}$ between 1-10 Hz < $10^{-3}/\sqrt{\text{Hz}}$ between 10-10 kHz
7. frequency fluctuations	< $10^4/\sqrt{\text{Hz}} \times [1\text{Hz}/f]$ between 1-10kHz (same as free running npro)

Table 3.1: Laser specifications [151]

A laser shutter (LA SH2) is placed at the aperture to the HPL light tight box and is connected to the ALPS interlock system. After the HPL box, half-wave plates and PBS are used to control the power injected into the rest of the setup.

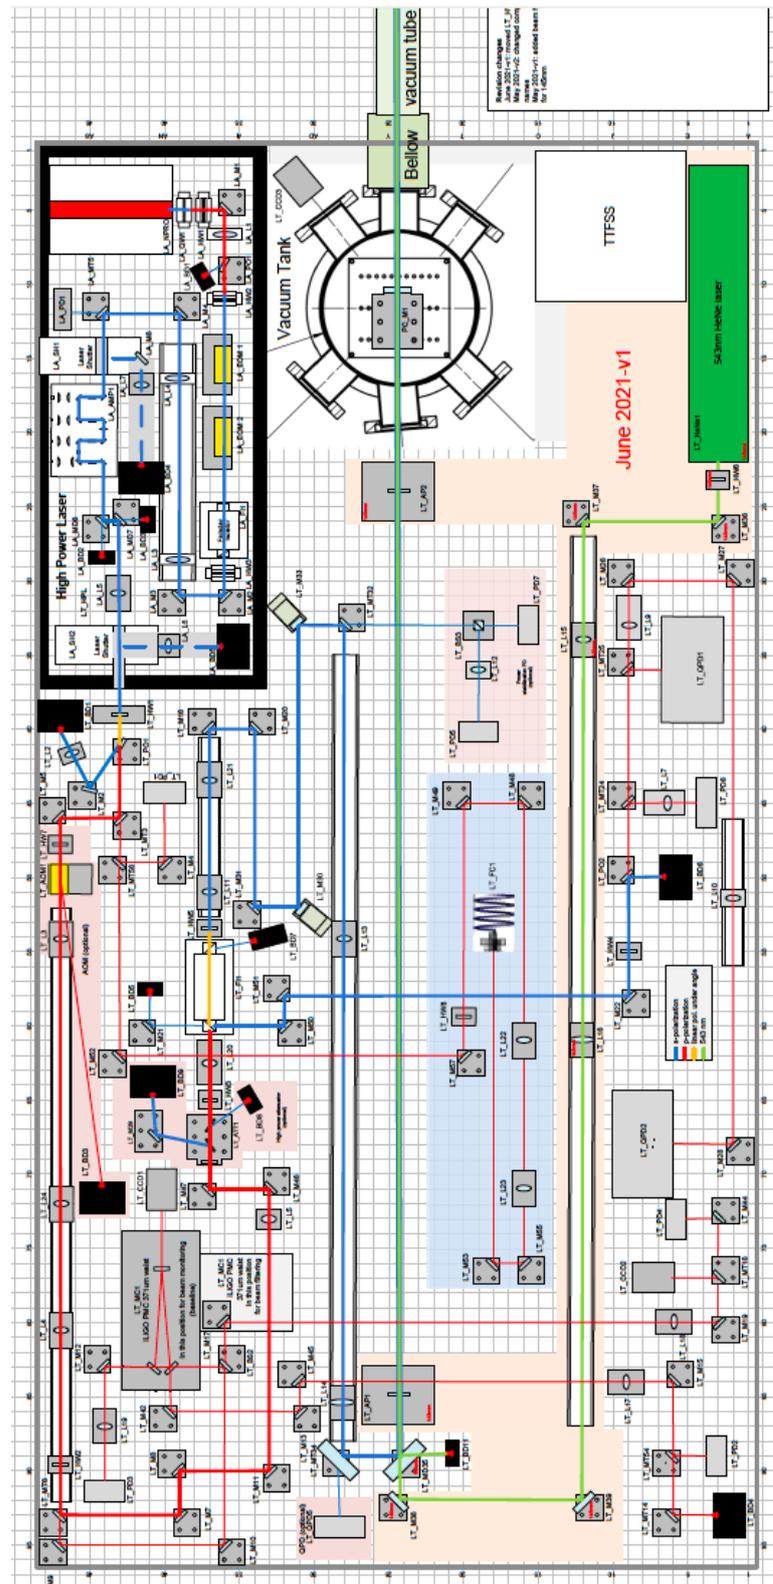


Figure 3.8: Table layout of the NL breadboard [151]. Due to the symmetry of the ALPS experiment the NR table layout uses a nearly identical layout [150].

LTL24 and LTL4 are used to mode match the HPL to a pre-mode cleaner (PMC) that is used to filter any higher order mode content present in the beam after amplification.

I designed the layout for the various mode matching and Gouy phase telescopes that are required after the PMC and decided on appropriate mirror placement based on the space constraints on the optical table. An explanation of each design choice is outlined below. Some basic theory on Gaussian beams and cavity mode matching is shown in appendix 5.1 and 5.2.

A Faraday isolator is placed to prevent back reflections from the cavity and to pick off the back reflected beam for the DWS system (see section 3.2.4). A thermal effect in the Faraday isolator was discovered by colleagues from AEI so a telescope (LT5 and LTL20) is placed before it in order to collimate the beam with a 1.7 mm beam radius. This radius is large enough to mitigate the thermal effect but small enough such that no significant clipping should occur on the Faraday aperture.

After the Faraday isolator another lens telescope (LTL11 and LTL21) is used to form a $290\ \mu\text{m}$ waist at a point equidistant between two steering mirrors (LTM30 and LTM33) that are used in the auto-alignment loop (see section 3.2.6). The mirrors are positioned far enough away from each other to ensure a 90° Gouy phase between the two steering mirrors as the beam propagates through the waist.

Lenses LTL13 and LTL14 are used to mode match the input beam to the production cavity, the curved input mirror being PCM1. Following equation 3.7, the coupling sensitivity scales as $(BL)^4$, and therefore the longest magnetic string length possible is desired. Powerful magnets were repurposed from the disused HERA particle accelerator [152] and installed in a section of the old HERA injection tunnel that stretches roughly 250 m. The number of repurposed magnets and the total length of the tunnel mean that each optical cavity is 120 m long. Due to the symmetry of the experiment the beam waist must be located inbetween the PC and RC, i.e. at the COB. The beam size was selected such that the Rayleigh range, $z_R = \frac{\pi w_0^2}{\lambda}$, is roughly the length of the cavity (~ 120 m) such that minimal clipping occurs within the magnet string. This gives the desired beam waist inside the cavity of ~ 6 mm. Following basic Gaussian beam propagation formulae, the beam has a radius of curvature of $R(z) = 222$ m and a radius of ~ 9 mm at the input to the cavity. These beam parameters then define the mirror properties: one flat end mirror and one curved input mirror with radius of curvature $R_c = 222$ m. The aforementioned mode matching telescope was designed to transform the input beam size and shape to match these conditions in order for resonance to occur.

The elements of the optical layout concerning the differential wavefront sensing are outlined in sections 3.2.3 and 3.2.4.

3.2.3 Differential Wavefront Sensing Theory

As mentioned previously, the main goal of this project was to design and implement an auto-alignment system for the ALPS IIc experiment. Before outlining the design of the system, some background theory on the technique of Differential Wavefront Sensing (DWS) is required.

Differential Wavefront Sensing is a technique used to detect a difference in phase-fronts of a measurement beam with respect to a reference beam. In some senses it is similar to the method of Pound Drevor Hall (PDH) length sensing [53] except that it utilises several photodiode elements (QPDs) to measure phase gradients across the spatial components of the beam rather than a single photodiode element to detect length differences.

A diagram illustrating the basic premise of DWS is shown in figure 3.9.

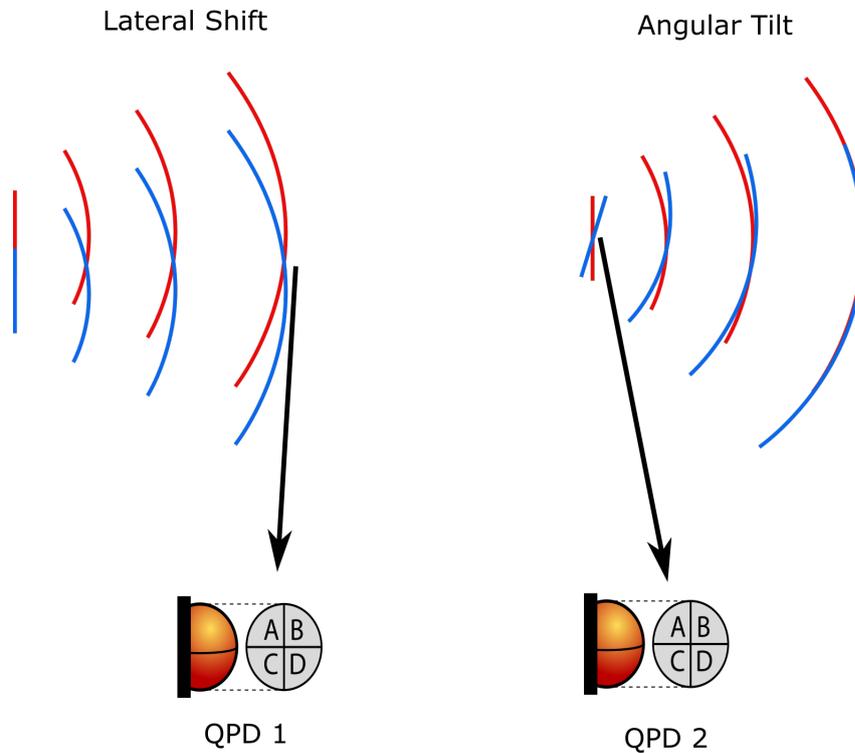


Figure 3.9: A model illustrating the premise of DWS. **Red**: The phase-fronts of the reference beam. **Blue**: the phase-fronts of the measurement beam as they propagate from left to right. **Left**: illustrates how the two phase-fronts interfere for a lateral shift of the measurement beam with respect to the reference. In the near field there is no phase gradient seen, however in the far field there is a phase gradient across the profile of the beams. **Right**: illustrates how the two phase-fronts interfere for an angular tilt of the measurement beam with respect to the reference. Now in the far field there is no phase gradient seen, however in the near field there is a phase gradient across the profile of the beams. This signal can be extracted by demodulating the signals on appropriately placed quadrant photodiodes (QPDs) 1 and 2 and taking the appropriate sums and differences of the quadrants.

In theory DWS can be used to detect misalignments between any two beams; however in this thesis we will be focusing on the misalignment of a beam with respect to the eigenmode of a Fabry-Perot cavity. It is also assumed that there is sufficient mode matching into the cavity such that coupling from second order modes can be ignored.

Although figure 3.9 provides a good illustration of how these signals arise and can be extracted, in reality the signals appear as a result of higher order mode coupling. The paper by D. Anderson outlines how higher order modes couple into a beam that is misaligned into a Fabry-Perot cavity [153] and an overview is provided here.

For a Gaussian beam, the fundamental and first higher order mode can be described respectively by:

$$U_0(x) = \left(\frac{2}{\pi x_0^2}\right)^{\frac{1}{4}} e^{-\left(\frac{x}{x_0}\right)^2}, \quad (3.8)$$

$$U_1(x) = \left(\frac{2}{\pi x_0^2}\right)^{\frac{1}{4}} \frac{2x}{x_0} e^{-\left(\frac{x}{x_0}\right)^2}, \quad (3.9)$$

for a waist radius of x_0 and a lateral distance from the centre of the beam waist x . For the case where the input beam is well aligned to a Fabry-Perot cavity the field distribution is described only by the fundamental mode i.e. $\Psi(x) = AU_0(x)$. A translation of the input beam by a distance a_x with respect to the aligned case is then:

$$\Psi(x) = AU_0(x - a_x), \quad (3.10)$$

which after some algebra simplifies to:

$$\Psi(x) = A \left[U_0(x) + \frac{a_x}{x_0} U_1(x) \right]. \quad (3.11)$$

Now let us consider the other type of misalignment that can occur which is a tilt of the input beam with respect to the cavity. Equation 3.12 below, seems to show that small tilts do not couple to the field distribution to the first order:

$$|\Psi(x)| = |\Psi(x')| (\cos \alpha)^{-1} = |\Psi(x')| (1 + \alpha^2 + \dots), \quad (3.12)$$

where α is the tilt of the beam with respect to the cavity eigenmode. However, due to the tilt of the wave, the phase now varies with x which results in:

$$\Psi(x) = AU_0(x) e^{i\frac{2\pi\alpha x}{\lambda}}, \quad (3.13)$$

which can be expanded for small misalignments $\frac{2\pi\alpha x}{\lambda} \ll 1$ and simplified to give:

$$\Psi(x) = A \left[U_0(x) + i\frac{\pi\alpha x_0}{\lambda} U_1(x) \right], \quad (3.14)$$

which shows that small tilts of the input beam also result in coupling of the fundamental mode to the first HOM. Comparing equations 3.11 and 3.14 we can see that misalignments caused by lateral shifts of the beam and tilts of the beam both cause coupling of the fundamental to the first higher order mode. However, in the case of 3.11 the coupling is purely real and in the case of 3.14 the coupling is imaginary; i.e. the signal is in the orthogonal quadrature. This is a crucial point in using DWS to form an auto-alignment system as one detector can be placed such that it is sensitive to lateral misalignments of the beam to the cavity and the other at a 90° phase difference such that it is sensitive to tilt misalignments, and the two can be used together to adjust two steering mirrors that have a similar 90° Gouy phase separation. The phase difference between the components can be established by forming an appropriate beam waist (see [51] for a description of the accumulation of Gouy phase as a beam passes through a waist). The signal can be extracted by using a quadrant photodetector and taking the appropriate sums and differences of the signals seen on each quadrant to reconstruct error signals for x and y in the lateral and tilt degrees of freedom.

In reality however, one does not need to place one sensor in the near field and one in the far field, nor does one steering mirror need to be in the near field and one in the far field. As long as there is a “sufficient” Gouy phase separation between the sensors and there is a similarly “sufficient” separation between the steering mirrors one can use a matrix to rotate the signals from the sensor basis into the mirror basis. Exactly how much Gouy phase separation is sufficient depends on the noise of the experiment; anything besides 90° separation and you begin to lose SNR in one quadrature. Another important point is that the Gouy phase of the light on the QPD sensors does not need to be identical to the phase of the light on the steering mirrors as it is possible to rotate the signals from any arbitrary sensor basis into any arbitrary steering mirror basis. The determination of the coefficients for the linear matrix combination is described in section 3.2.5.

In some ways this approach using a linear combination matrix may appear more complicated than is required. Why not just place a sensor in the near field to act on a mirror that is also in the near field? Actually it grants greater flexibility in the optics design and is often even unavoidable. Consider trying to position one sensor at exactly 0° accumulated Gouy phase in order to actuate on a steering mirror

that is also at exactly 0° Gouy phase. It is a simple and elegant design in theory but in reality the position (and hence the Gouy phase) of the optical components will always be slightly wrong, resulting in some signal from the desired quadrature being lost and some signal from the other quadrature being coupled in. A linear combination matrix is therefore almost always required.

The placements of the sensors and steering mirrors and the Gouy phase telescopes are described in section 3.2.4.

3.2.4 Differential Wavefront Sensing Design

Section 3.2.3 described a measurement beam and a reference beam in order to detect eigenmode misalignments; however we mentioned back in section 2.1.4 that photodiodes are not fast enough to measure electric field oscillations from 1064 nm light. Instead sidebands can be imposed on the carrier and the phase information between the two can be extracted via demodulation of the photodiode signal. In the case of the ALPS experiment the measurement beam is the carrier beam that is input to the optical cavity (whether that is NL or NR) and the reference beam comes from 9 MHz sidebands that are imposed on the carrier by the EOM mentioned in section 3.2.2. The 9 MHz sidebands are promptly reflected from the input cavity mirror, whilst the carrier beam is resonant within the PC, incident upon the end mirror (LT M2) and then partly leaks back through the input mirror LT M1 where it beats with the promptly reflected sidebands. The carrier and sidebands propagate together back through the setup and are picked off at the Faraday Isolator LT FI1 and sent to the differential wavefront sensing (DWS) section of the optics board. Any signal related to the phase relationship of the carrier with respect to the sidebands can be extracted via demodulation as in [53].

As mentioned in section 3.2.3, this technique requires two sensors (QPDs) with $\sim 90^\circ$ Gouy phase separation between them. The sensors selected for the experiment also have a 1.5 mm optimum beam radius for their use which requires some fairly straightforward optical setup to ensure.

The mirror actuators also require a 90° Gouy phase separation. There is not a stringent requirement for the beam size incident on the steering mirrors as long as the beam is small enough such that significant clipping losses do not occur. The mode matching for the mirrors is shown in figure 3.12.

In figs. 3.14 to 3.18 it is shown how the error signals are sent from the DWS sensors shown in figures 3.10 and 3.11 to the steering mirrors shown in figure 3.12. It is simpler to illustrate the drift of optical components by fixing the position of the input optics and assuming that all motion comes from drift of the COB (The flat end mirrors are mounted on the COB). There are two types of misalignment that the steering mirrors are used to correct for: lateral shift of the COB and tilt of the

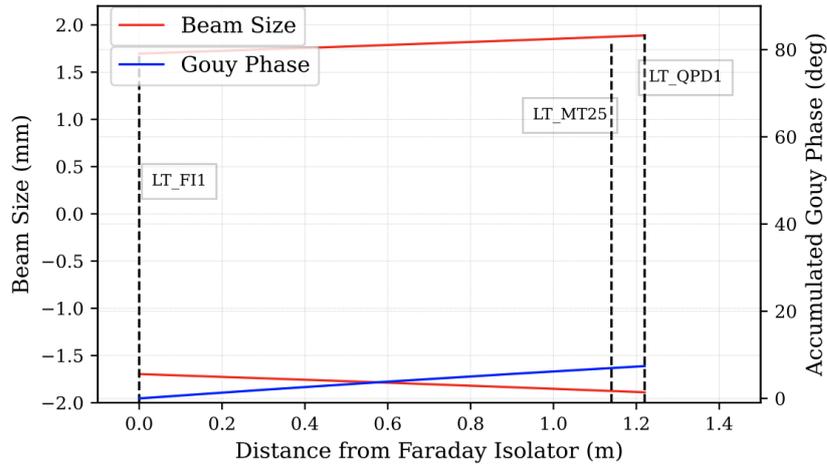


Figure 3.10: Mode matching from LT FI1 to LT QPD1. Component names reference the table layout in figure 3.8. Due to the size of the beam coming from the cavity and the size of the active area of the QPD, no further optics are required. For simplicity only the Gouy phase accumulated from the pick off at the Faraday isolator onwards is considered.

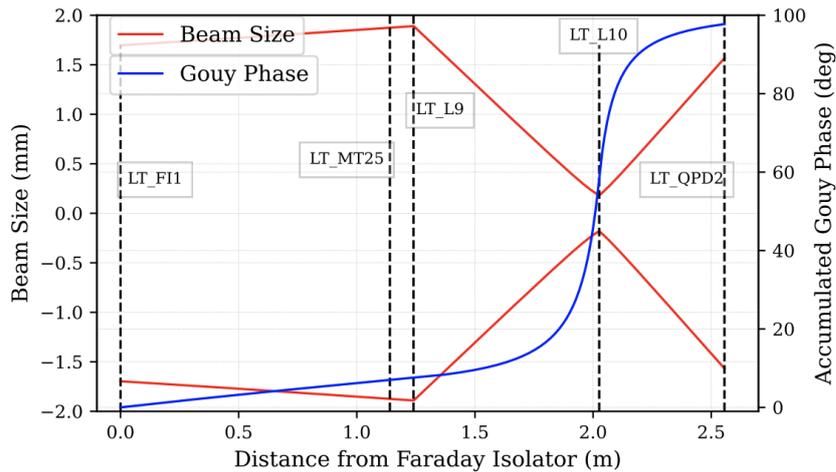


Figure 3.11: Mode matching from LT FI1 to LT QPD2. A Gouy phase telescope is required to ensure that there is a 90° Gouy phase difference between the two sensors and to ensure there is still a 1.5 mm beam radius incident on QPD2. Some careful positioning of LT L10 is required to achieve both of these requirements.

COB.

The coefficients a, b, c, d for the linear combination matrix are found by injecting a signal into one of the steering mirrors and measuring the amplitude of the error signals on each QPD. This process is repeated for each degree of freedom on each mirror.

In this case, the rotation of the curved cavity input mirror serves to shift the cavity eigenmode up and down. An important point to be aware of is if the curved mirror is actuated faster than the steering mirrors then the beam will not be properly aligned

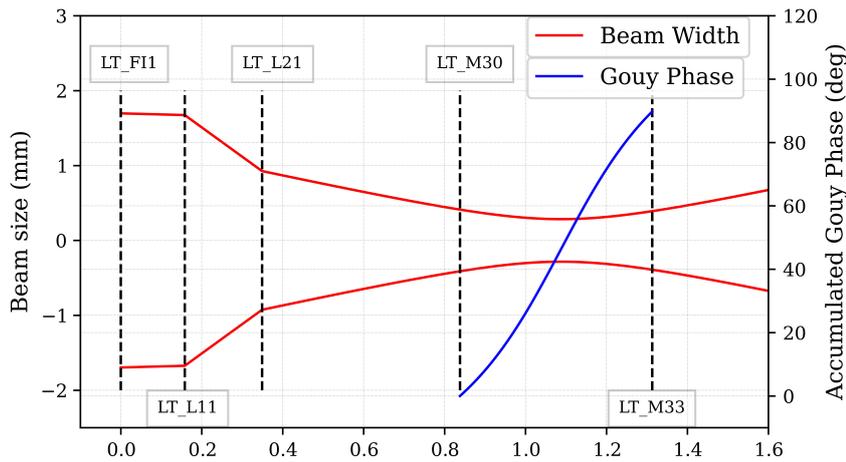


Figure 3.12: The Faraday isolator FI L1 is the same isolator that picks off the reflection from the cavity and sends it to the DWS setup in figure 3.10 and 3.11. Here the beam is propagating towards the cavity as it goes from left to right. A telescope from lenses $LT\ L11 = 443\text{ mm}$ and $LT\ L21 = -334\text{ mm}$ forms a beam waist of about $280\ \mu\text{m}$. For simplicity only the Gouy phase accumulated by the beam as it passes from Steering mirror 1 ($LT\ M30$) to Steering mirror 2 ($LT\ M33$) is plotted.

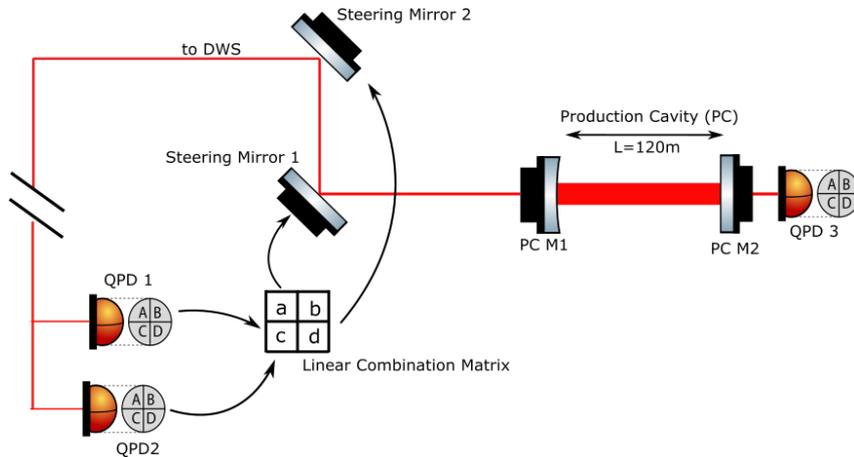


Figure 3.13: A simplified diagram illustrating the ‘aligned’ case for the beam into the Fabry-Perot Production Cavity. For simplicity most optical components have been omitted and it is assumed that lens positions (not shown) have been independently optimised to maximise coupling of light into the cavity. QPD 1 and QPD 2 represent the differential wavefront sensors mentioned in section 3.2.3. The black arrows represent a control signal. The linear matrix combination that rotates the error signals from the sensor basis to the mirror basis has coefficients a, b, c, d . These coefficients must be measured experimentally. QPD 3 is not strictly a part of the auto-alignment system but represents a QPD on the COB that is used to detect the spot position on the flat end cavity mirror and then actuate on the curved input mirror. Since the eigenmode of the cavity is always perpendicular to the flat mirror surface, this has the effect of shifting the eigenmode up and down within the cavity.

into the cavity. This has not been a focus of the commissioning work so far but will be an important factor when ALPS IIc finishes commissioning.

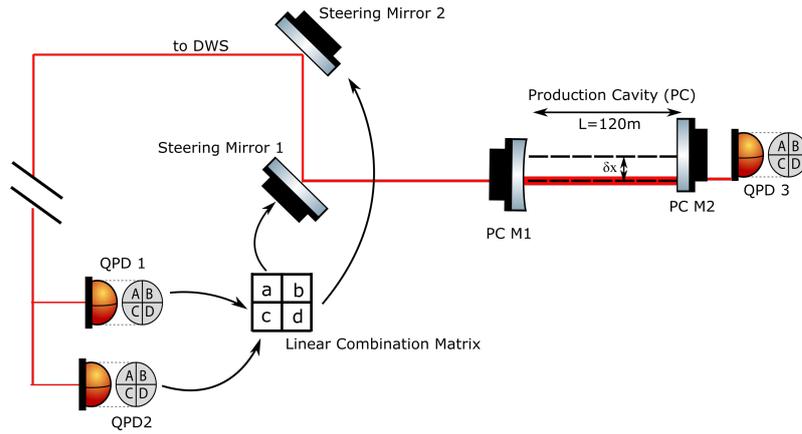


Figure 3.14: A lateral shift in the entire COB by a distance δx .

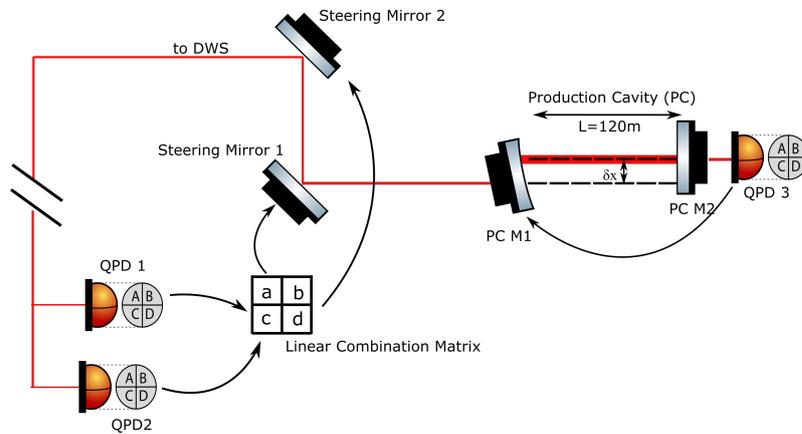


Figure 3.15: The signal from QPD 3 is used to actuate on the curved cavity input mirror and centre the eigenmode on the flat end mirror. This means the input beam is misaligned to the cavity which results in higher order mode coupling and reduced power circulating in the fundamental mode.

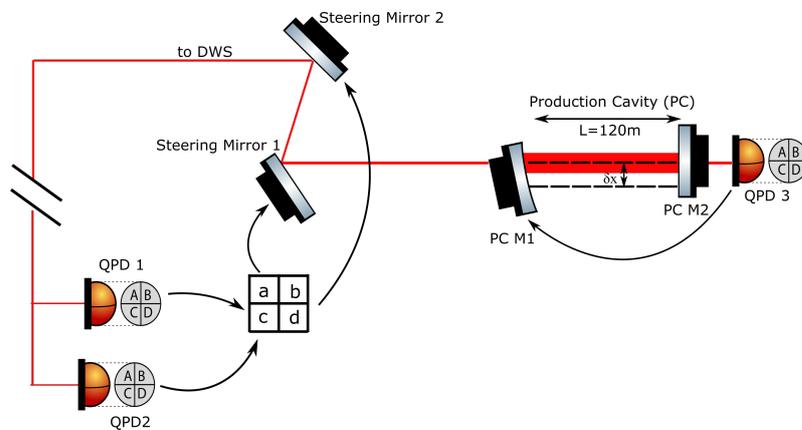


Figure 3.16: The error signals from QPD 1 and QPD 2 are used to actuate on steering mirrors 1 and 2 and optimise the input beam to the cavity.

The second type of misalignment is when the COB tilts with respect to the optical axis. It is worth mentioning that the scenario shown in figure 3.17 is really just an illustrative example. In reality only the lateral shift of the COB is corrected using the error signal from QPD 3, however it makes a clearer picture to consider only the motion of the COB. It is of course analogous to consider drift of the input optics, with respect to a stationary cavity eigenmode.

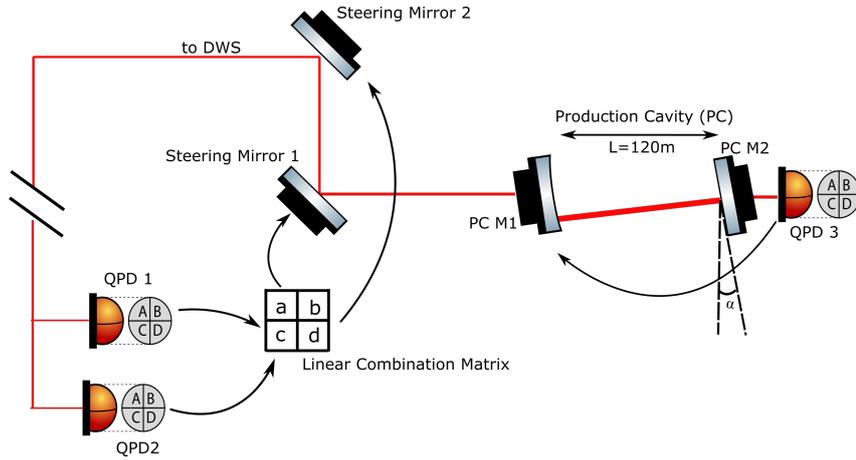


Figure 3.17: An angular tilt of the entire COB by an angle α . The curved input mirror tilts in order to keep the beam spot centered on the flat end mirror, however the input beam is now misaligned and will couple to higher order modes within the cavity, resulting in less light circulating in the fundamental mode.

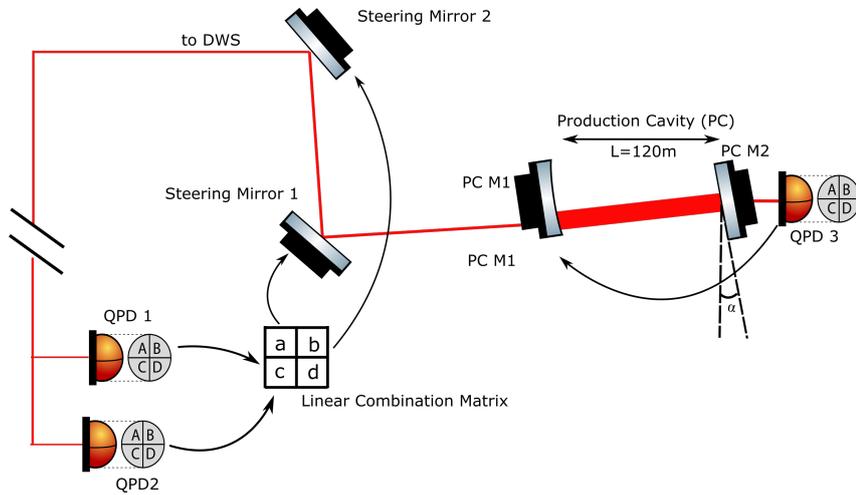


Figure 3.18: The error signals from QPD 1 and QPD 2 are used to optimise the input beam into the cavity and correct for the angular misalignment α of the COB.

Although this section has exclusively referred to the setup for the PC, the optics setup for the RC is almost an exact copy of this layout. There are some specific differences which are outlined in section 3.2.5 and 3.2.6.

3.2.5 Signal Processing

Although the optics layout is symmetric for both the PC and the RC there is some difference in the approach to how signals are processed in the two sides of the experiment. This was not a conscious decision but rather a by-product of having the design of the PC and RC sides of the experiment being carried out by mostly independent groups who each have their own preferences in experimental work. The PC AA signal processing uses all analogue electronics whilst the RC AA signal processing elects to use a mostly digital system. This section outlines each step in the signal processing chain and describes any differences between the two systems whether positive or negative. As we will discuss in later sections 3.3, the PC AA system was not installed at the time of writing this thesis which makes a comparison of the two systems quite tricky. For this reason it is (perhaps unfairly) assumed that the PC system would have worked perfectly when comparing it to the installed system for the RC.

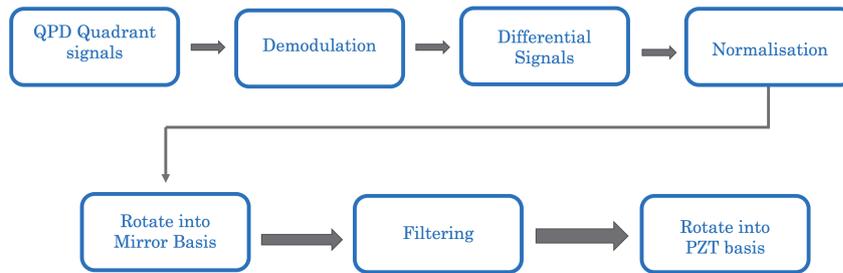


Figure 3.19: Flow chart showing each stage of signal processing required to turn the photocurrent signal into a correction signal that can be sent to the corresponding mirror actuator, referred to as a PZT - Lead (Pb) Zirconate (Zr) Titanate (Ti).

A misalignment of the input beam to the cavity eigenmode generates higher order modes that leak back out the cavity and co-propagate with the 9 MHz sidebands. The information on the phase between these two fields is extracted via demodulation with a 9 MHz Local-Oscillator (LO), the same as outlined in the PDH technique [53]. Both the PC and RC achieve this using analogue mixers with the LO sent to the onboard QPD electronics.

From this point onwards the RC experiment uses digital signal processing. The demodulated signals from each QPD quadrant are sent to an ADC and then a compactRIO (cRIO) system that contains a user-programmable FPGA. A downside of the digital approach is that the processing of the data is limited by the speed that the digital system can run. The cRIO is designed to run at 16 kHz, which is sufficiently fast for the auto-alignment; however, it was common in commissioning to see data being processed at rates as low as 8 kHz and even at the time of writing it seems that the system is running at a fraction of its design speed (see section 3.3.1).

The appropriate sum and difference of individual quadrant signals are then taken to form the x-y error signals. For a QPD arranged like in figure 3.13 the x signal is given by:

$$(A + C) - (B + D), \quad (3.15)$$

and the y error signal given by:

$$(A + B) - (C + D). \quad (3.16)$$

There is an optional normalisation of the error signal based on the sum signal seen on all four quadrants. This makes the magnitude of the error signal independent from fluctuations in light intensity.

The x-y error signals that are in the basis of the QPD sensors then need to be rotated into the mirror basis using a linear combination matrix as mentioned in section 3.2.4. For the RC digital setup the matrix is formed by adding values into the matrix cells on the digital interface. For the PC setup the matrix is established by adjusting the values of variable resistors on the auto-alignment servos.

The filtering required for the auto-alignment control depends on the experiment requirements and any specific noise features present in the error signal e.g. PZT resonances. Here the digital processing has some clear advantages as it is possible to easily adjust filter parameters, such as gain coefficients, completely remotely. During commissioning this proved especially useful as it was possible to remotely commission a large amount of the RC AA system (see section 3.3). Conversely using an analogue system means that filter parameters can only be adjusted by changing components on the physical board. There is space on the PC AA servos to mount additional components but it requires the boards to be uninstalled from the clean rooms and taken to an electronics lab for adjustments to be made. The only remote control that will be possible in the NL AA setup is an adjustment of the total gain of the feedback loop.

The control signals then need to be rotated from the mirror basis into the PZT basis. This is done using an appropriate matrix (this is different to the aforementioned linear combination matrix) determined by the placement of the PZTs on the steering mirror mounts. As seen in the next section there are different mirror actuators used in the PC and RC side which require different matrices to set up. For the RC this is another digital process carried out within the cRIO, whereas for the PC this is done within the High Voltage Amplifier (HVA) analogue electronics.

After the PZT rotation the control signals need to be amplified in a high voltage

amplifier before being sent to the PZTs. For the PC, the PZT matrix electronics are part of the HVA. For the RC, the signals are digital so are sent from the cRIO to a Digital Analogue Converter (DAC) which can output a $-10 \rightarrow 10$ V signal. The HVA can actually only accept $0 \rightarrow 10$ V input voltages, and therefore there is double the DAC noise present in the voltages going to the HVA than is necessary. Based on the results in 3.3 it does not appear that this is a limiting problem but it represents a small improvement that could be implemented in the RC system.

Before moving on to discuss the mirror actuators it is perhaps worth mentioning one final downside of using the digital system for the RC AA setup which is that, due to the group's unfamiliarity with digital signal processing, there were bugs and glitches in the digital system that required significant work to overcome, and in fact for some of these issues there is still no fix implemented. Although it is hard to say for sure, it is expected that commissioning of the PC AA system will be a more straightforward task since most of the group are familiar with analogue electronics and troubleshooting the typical problems seen with them. The various bugs and glitches are briefly outlined in section 3.4.7; however the problems were mainly solved by a fellow PhD student from the University of Florida, D. Brotherton. I therefore leave the detailed description of this work to his future thesis [154].

3.2.6 Mirror Actuators

The entire optics setup from the PMC onwards was simulated using the optics simulation software FINESSE [155] and the required angular tilt of each steering mirror was measured for increasing misalignments of the cavity. It was found that to correct for a 2 mm beam spot shift on the flat cavity end mirror (either due to angular tilt or lateral shift) the mirrors require a range of at least $\pm 75 \mu\text{rad}$ on mirror 1 and $\pm 112 \mu\text{rad}$ on mirror 2. This 2 mm range is a somewhat arbitrary number as it corresponds to the beam spot shift that would result in significant clipping within the cavity and does not refer to a specific amount of drift that is expected during a science run.

In the choice of mirror actuator we again have some asymmetry between the PC and RC setups. For the RC side of the experiment two KC1-PZ(/M) Thorlabs mirrors have been installed with an actuation range of only $\pm 73 \mu\text{rad}$. This is of course a smaller range than what was specified above. It was originally intended to install one KC1-P actuator ($\pm 275 \mu\text{rad}$ range) and one KC1-PZ(/M) actuator ($\pm 73 \mu\text{rad}$ range) [150] since one mirror is predicted to require less dynamic range. Unfortunately there was a mix up during commissioning which resulted in two of the low range KC1-PZ(/M) models being installed in the RC table. Removing them at this point would misalign the beam into the cavity and disrupt commissioning efforts during a crucial period. Therefore it was decided to leave them as is and only replace them if strictly necessary. Section 3.3.2 discusses the issues that may

be encountered as a result of these lower range mirror actuators.



Figure 3.20: **a)** picture of the Thorlabs KC1-PZ(/M) steering mirror. **b)** the arrangement of the PZTs to control the tip/tilt of the steering mirror.

A rotation matrix is required to rotate the X/Y control voltages into the individual PZT basis.

$$\begin{vmatrix} PZT_1 \\ PZT_2 \\ PZT_3 \end{vmatrix} = \begin{vmatrix} X - Y \\ X + Y \\ -X + Y \end{vmatrix} = \begin{vmatrix} 1 & -1 \\ 1 & 1 \\ -1 & 1 \end{vmatrix} \begin{vmatrix} X \\ Y \end{vmatrix}. \quad (3.17)$$

The steering mirrors purchased for the PC side are PIS-315.10 which have a range of $\pm 600 \mu\text{rad}$. This of course is significantly more range than is needed based on the simulations and are more expensive than other mirror actuators but the increased range provides more headroom in case more drift than expected is encountered.

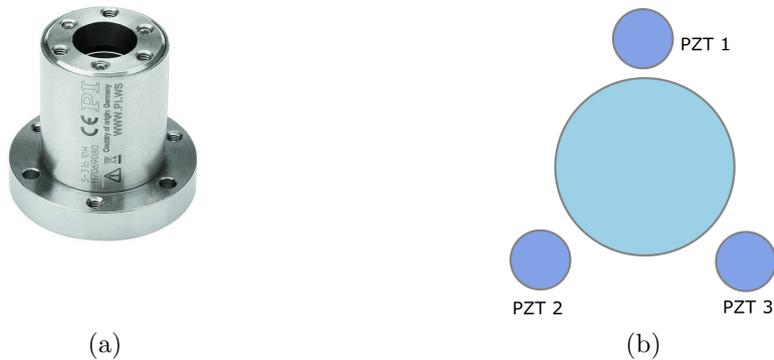


Figure 3.21: **a)** picture of the PI PI-S-315 steering mirror. **b)** the arrangement of the PZTs to control the tip/tilt of the steering mirror.

Due to the different position of the PZTs, a slightly different matrix is required for the PC side:

$$\begin{vmatrix} PZT_1 \\ PZT_2 \\ PZT_3 \end{vmatrix} = \begin{vmatrix} -Y \\ X + Y \\ -X + Y \end{vmatrix} = \begin{vmatrix} 0 & -1 \\ 1 & 1 \\ -1 & 1 \end{vmatrix} \begin{vmatrix} X \\ Y \end{vmatrix}. \quad (3.18)$$

3.3 250 m Cavity Results

I travelled to the ALPS site in Hamburg for three months to install the auto-alignment system and work on general commissioning of the detector (see picture 5.22). During this time a 250 m cavity was in place rather than the dual 120 m cavity setup described in section 3.2. This setup allowed troubleshooting of the experiment while other commissioning work was being carried out in parallel in the central hall.

Several tests were carried out on the auto-alignment system using the 250 m cavity. The most successful in terms of noise suppression was two minutes in duration and is described in section 3.3.1. The longest lock that was carried out was 2 hours long and is described in section 3.3.2.

3.3.1 Two minute Lock

As mentioned in section 3.2.4 there are 4 error signals in the Auto-Alignment (AA) scheme used to control 2 degrees of freedom (x and y) of 2 separate steering mirrors. In this section I will refer to the error signals produced by each QPD as QPD 1/2 X/Y. The control voltage refers to the error signal after it has passed through the filter stage but before it has gone to the high voltage amplifier i.e. it is a voltage between $0 \rightarrow 10V$. The control voltage will be referred as Steering Mirror (SM) 1/2 X/Y to identify which steering mirror and which degree of freedom the signal is controlling.

A comparison of the time series data of the error signal QPD 1 X for the case of the AA loop open and closed over a two minute time period is shown below. The time series for QPD 1 Y and QPD 2 X/Y are shown in appendix 5.12.

During this lock a single integrator was used in the filter stage and the integral gain applied to each error signal was increased until the auto-alignment system was on the verge of ringing (see figure 5.14 in appendix 5.10). Coefficients for the linear matrix combination were measured (method explained in section 3.2.4) but could not be implemented during this test due to limitations in the digital interface. Instead a unity matrix was applied.

It was found that the optimum alignment of the input beam to the 250 m cavity (based on the maximum light coupled into the 250 m cavity) corresponded to a non-zero setpoint for all 4 error signals. It is assumed that this is due to voltage offsets within the QPD readouts that are responsible for photocurrent to voltage conversion

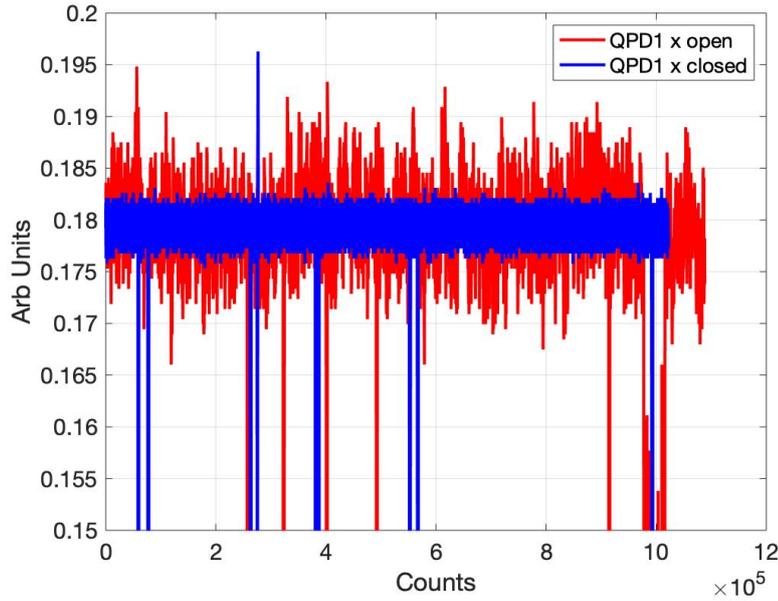


Figure 3.22: Error signal for QPD 1 X for open and closed AA loop over a 2 minute duration. A noticeable improvement in the error signal noise is seen when the loop is closed. The Amplitude Spectral Density (ASD) of this time series is shown in figure 3.23. The set point of the error signal was 0.18 in arbitrary units. The x-axis is displayed in counts rather than in seconds. This is due to uncertainty in the frequency that the cRIO system processed and saved data (see appendix 5.11). The spikes in the error signal do not necessarily indicate loss of cavity lock and are actually mostly glitches in the data saving (see section 3.4.7).

and demodulation. There is a significant DC offset of about 100 mV in quadrant B of QPD1 that is suspected to have come from damage to the mixer when a large voltage was mistakenly sent to the QPD electronics. This offset is then of course coupled into the X/Y error signals for QPD1. There are offsets that are an order of magnitude smaller seen on all other quadrants, the source of which is unknown but these offsets are not large enough to be concerning. The damaged mixer will be replaced in future commissioning work but it was possible to compensate for all the offsets and close the AA loop by adjusting the set point of the error signals accordingly.

Although the processing frequency of the cRIO time series is not reliable (see appendix 5.11) it is still possible to make a comparison between the open and closed loop cases assuming there are no changes to the rate that data is being processed. The ASD of all 4 error signals for the AA loop open and closed are shown below using a sampling frequency in the FFT of 16 kHz. Only stable sections of the time series data were used in producing the ASD.

Figures 3.23, 3.24 and 3.25 show that 3 degrees of freedom corresponding to the error signals from QPD1 X/Y and QPD2 X are all suppressed by about an order of

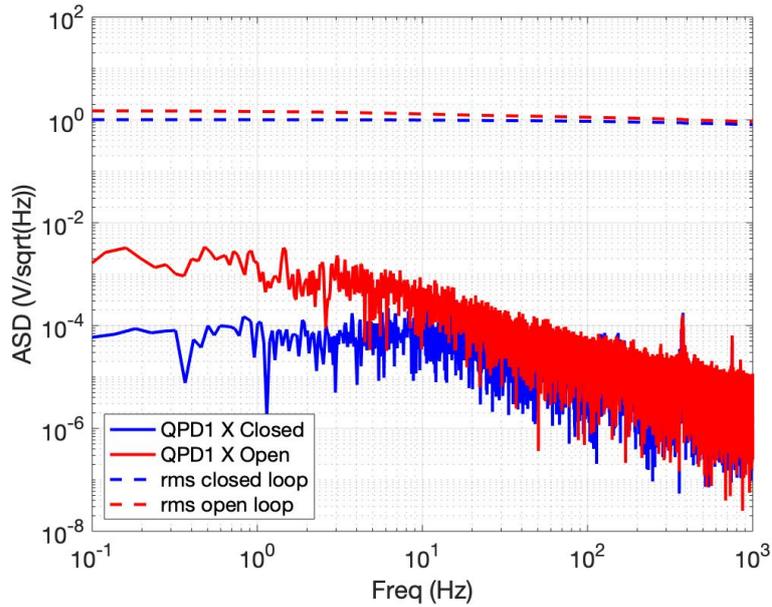


Figure 3.23: ASD for the error signals for QPD1 X in the case of both open and closed AA loop. The x axis is limited from 0.1 Hz to 1 kHz.

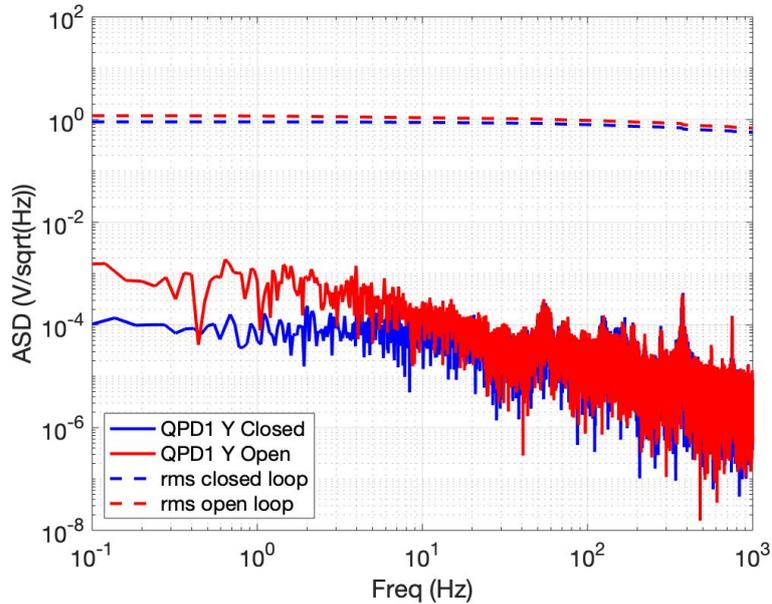


Figure 3.24: ASD for the error signals for QPD1 Y in the case of both open and closed AA loop. The x axis is limited from 0.1 Hz to 1 kHz.

magnitude at low frequencies; however figure 3.26 shows QPD2 Y is only suppressed by about a factor of 2. The integral gain coefficient on the QPD2 Y error signal could only be increased to 1/3 the magnitude of the other error signals gain coefficients before it began ringing. The reason for this is unknown, however, this asymmetry in the filter coefficients probably explains the reduced suppression of the error signal

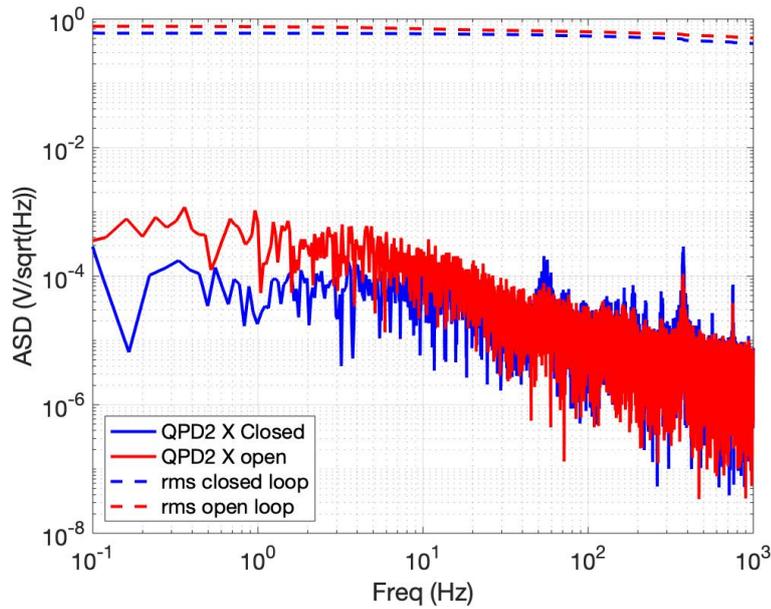


Figure 3.25: ASD for the error signals for QPD2 X in the case of both open and closed AA loop. The x axis is limited from 0.1 Hz to 1 kHz.

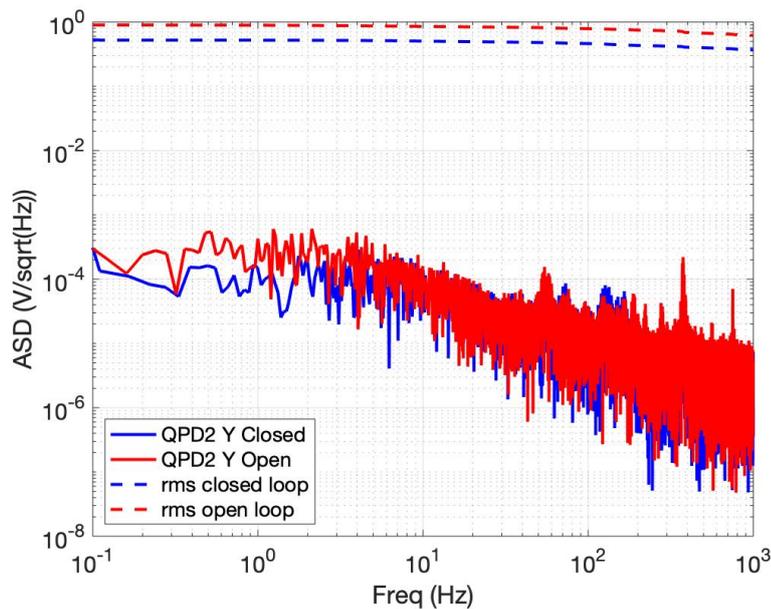


Figure 3.26: ASD for the error signals for QPD2 Y in the case of both open and closed AA loop. The x axis is limited from 0.1 Hz to 1 kHz. In this degree of freedom there is less suppression of the pointing noise at $f < 10$ Hz.

in QPD 2 Y.

The power incident on a photodiode in transmission of the optical cavity was also recorded for this 2 minute measurement and is shown in figure 3.27. The data was not recorded with the (unreliable) cRIO but instead with a MokuLab, which has a reliable data sampling frequency of 10 kHz for this measurement.

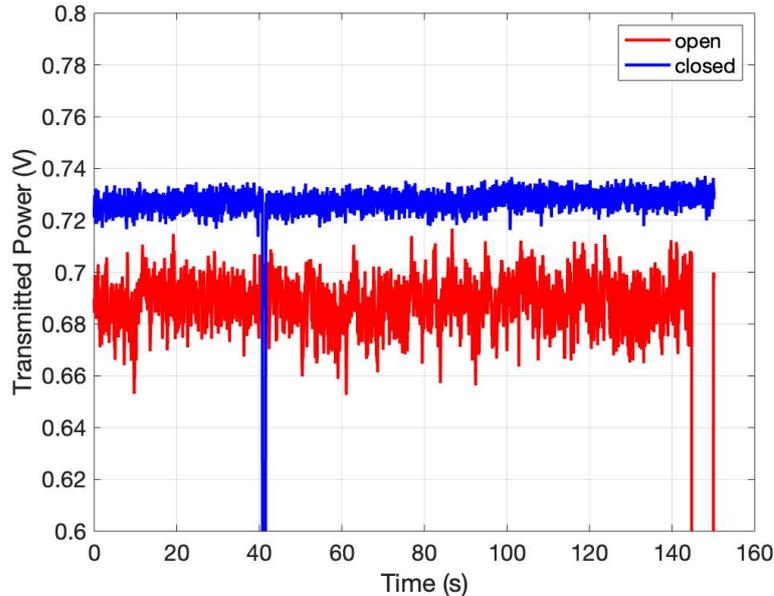


Figure 3.27: Comparing the power transmitted through the 250 m cavity for the AA loop open and closed during a 2 minute test of the NR AA setup. The power was recorded on a Moku in units of volts. Sharp drops in the transmitted power correspond to points where the cavity lost lock and were not used in the ASD analysis in figure 3.28

It was possible to adjust the set points of the error signals such that the power in transmission was increased as illustrated by the voltage level increasing by approximately 5%. It is also possible to observe an improvement in the rms when examining the time series. The amplitude spectral density of the cavity transmission time series is shown in figure 3.28.

The ASD of the transmitted power shows an improvement in low frequency noise (< 1 Hz) of about an order of magnitude. The unity gain frequency is approximately 7 Hz.

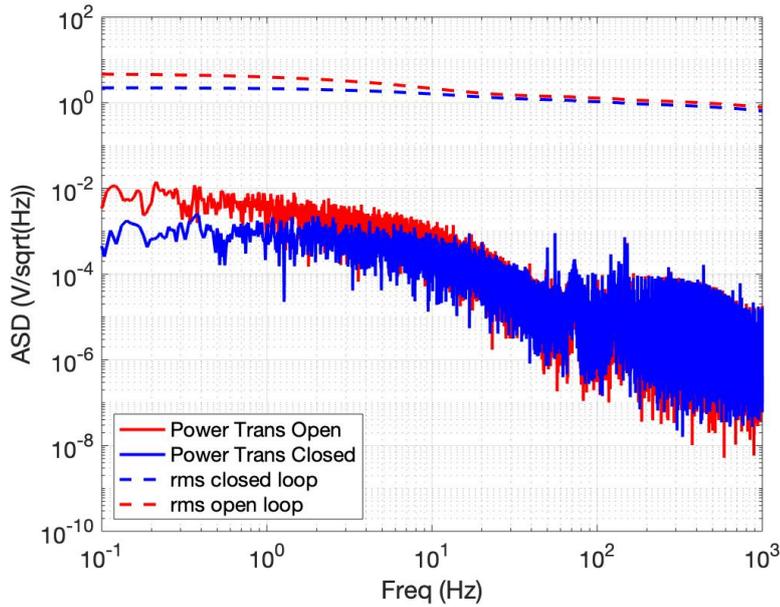


Figure 3.28: ASD of the 2 minute lock shown in figure 3.27. The noise features beyond 50Hz are not fully understood but are suspected to be a feature of the Moku:lab that was used to monitor the transmitted power.

3.3.2 Two Hour Lock

The auto-alignment system was monitored over a 2 hour period in order to observe the suppression of long term drift in the 250 m cavity eigenmode. As in section 3.3.1, there was a single integrator stage used in the closing of the loop and each integrator gain was increased until the system was on the verge of ringing (see figure 5.14).

The ASD of the data in figure 3.29 is shown in figure 3.30.

Figure 3.30 shows roughly a factor of 2 suppression in the low frequency cavity noise, about 5 times less suppression that seen during the 2 minute lock in figure 3.28. It is suspected that this is due to improper spot positioning on QPD2 at the start of the test. Unfortunately there was a limited window for this measurement and no time to optimise it further.

The interesting result from this lock however, is that there is a significant drift in the DC transmitted power level in figure 3.29 when the AA loop is not engaged, with the voltage level dropping by roughly 13% over the 90 minute test. From equation 3.7 the coupling sensitivity is proportional to $\left(\frac{1}{P_{PC}\beta_{RC}}\right)^{\left(\frac{1}{4}\right)}$. Assuming an identical drift in the PC and RC cavities, this decrease in cavity power due to eigenmode drift is equivalent to the minimum coupling sensitivity increasing from $g_{\alpha\gamma\gamma} = 2 \times 10^{-12}$ to $g_{\alpha\gamma\gamma} = 2.14 \times 10^{-12}$ over the 90 minute period. If the cavities continued to drift at the same rate the minimum sensitivity would continue to worsen until eventually the cavities would become completely misaligned and come unlocked.

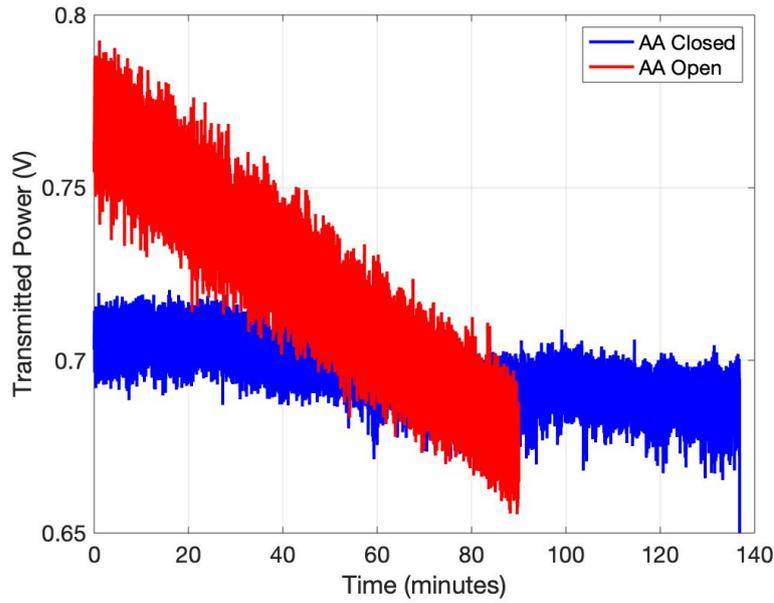


Figure 3.29: Comparing the power in transmission during a long measurement for the case of auto-alignment open and closed. The power was recorded on a Moku in units of volts. The closed loop data lasted a little over two hours whereas the open loop data lasted about 90 minutes and was taken by a colleague from the week earlier. There is less noise suppression seen in the rms compared to the 2 minute lock in figure 3.27. When the loop is open there is a steady decrease in the DC transmitted power. When the loop is closed the DC level is fairly stable.

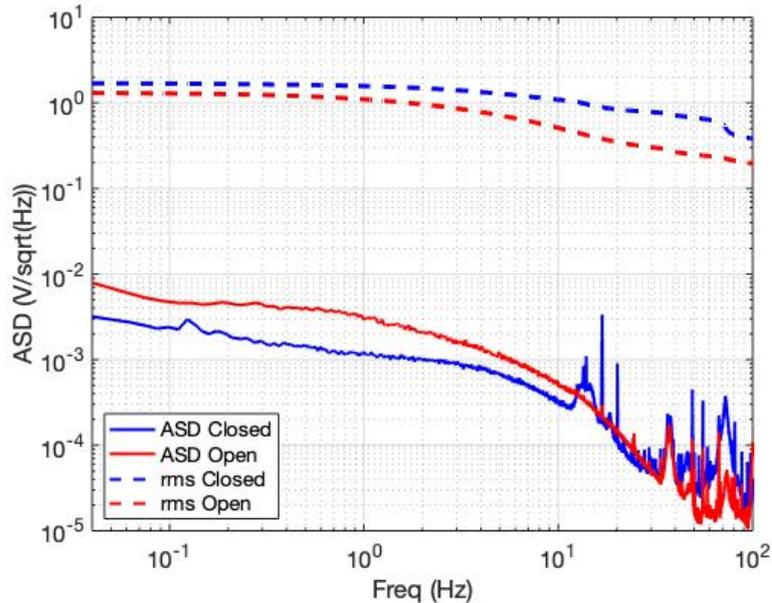


Figure 3.30: ASD of the power in transmission for the auto-alignment loop open (1.5 hour duration) and closed (2 hour duration). The noise features beyond 30 Hz are not fully understood and it is suspected that they are features of the Moku:lab used to record this data. The x-axis is limited between 40 mHz to 100 Hz.

Conversely, when the AA loop is engaged a change of roughly 4% is observed. Since the control voltages for the AA steering mirrors were not near saturation during this test it is suspected that the 4% change is not due to cavity eigenmode drift but rather some other uncontrolled system. There are several potential suspects that can be investigated for this including laser power drift, thermal lensing effects and drift in the mode matching lenses. This result indicates that auto-alignment will be necessary to achieve a successful science run at ALPS IIc.

By measuring the gradient of the control signal drift during the test it is possible to estimate the time until the auto-alignment control output is saturated.

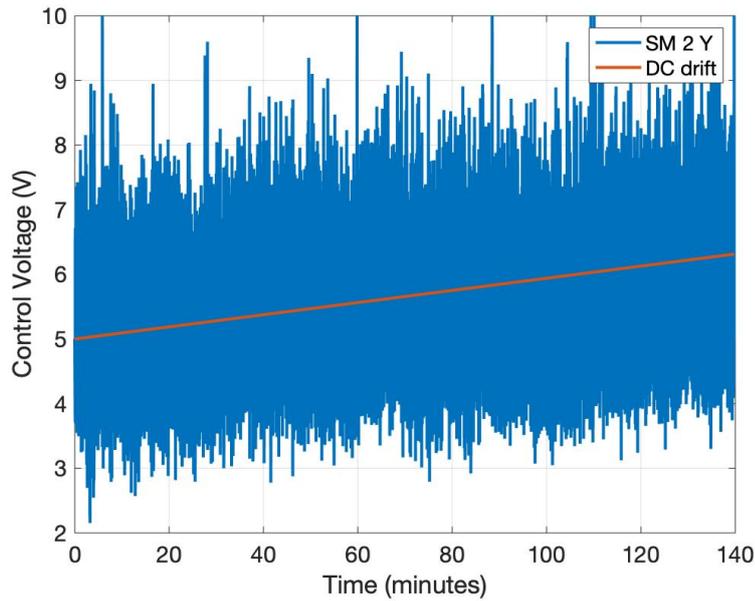


Figure 3.31: **Blue**: control voltage (before amplification) for SM 2 Y during the ~ 2 hour lock of the AA system. **Orange**: fit for the DC drift of the control voltage. The sample frequency for this data is 10 Hz. This control signal exhibited the largest amount of DC drift out of the 4 control voltages and is approximated by the orange fit. A plot of all 4 control voltages during this test is shown in appendix 5.13

The most significant DC drift in the control voltages is seen in SM 2 Y control. With reference to figure 3.31 it is estimated that 1 V of DC drift occurs over a time period of 2 hours. If it is assumed this drift continues at the same rate then the output would saturate after 10 hours. This is also a slight underestimate of the time taken to saturate the output as the rms fluctuates by approximately 2 V around the DC value so it is more likely that we would start to see reduced suppression of the input alignment noise after only 6 hours due to clipping. During a science run it is a target to take data for 20 days which corresponds to 480 hours.

There are however, options to compensate for prolonged drift in the experiment. There are 2 pico-motor steering mirrors that can be manually adjusted to control

the same 4 degrees of freedom as the 2 auto-alignment mirrors. If the control voltage starts to drift towards their output limit then the pico motors can be adjusted to bring the control voltages back to the mid point of their range. This has not been fully investigated but a preliminary test is shown in figure 3.32.

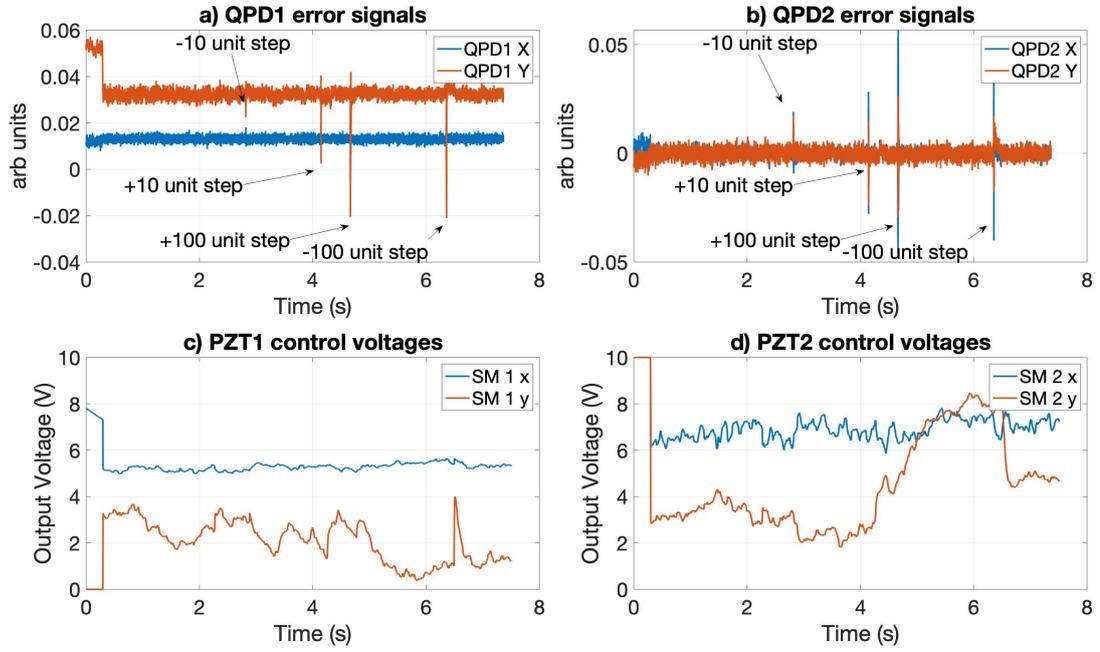


Figure 3.32: **a)** error signals for QPD 1 X/Y, **b)** error signals for QPD 2 X/Y, **c)** control voltages (before amplification) for SM 1 X/Y, **d)** control voltages (before amplification) for SM 2 X/Y. Non-zero setpoints were used for all error signals. Spikes in the error signal correspond to times when a pico-motor steering mirror took some step in the y-direction. 10 unit steps were taken at 2.8s and 4.1s. 100 unit steps were taken at 4.6s and 6.4s. The cavity remained locked throughout the test.

The spikes in the error signal in figure 3.32 correspond to times when the pico-motor steering mirrors were tilted along the y-axis. The high frequency return to the setpoint is not coming from the auto-alignment steering mirrors (since there is no high frequency response in the control voltage), but is suspected to be the pico motor mirrors settle time after each adjustment. It's also worth noting that the pico motors do not take uniform steps with each adjustment, hence why a step of +100 steps followed by -100 steps results in a slightly different control voltage for the SM 2 control. The same pico-motor mirror was used to make all tilt adjustments. The corresponding Gouy phases of the pico-motor steering mirrors are not known but it appears that the pico motor mirror under adjustment is coupled more strongly to SM 2. At ~ 2.8 s into the measurement, a step of -10 'units' was taken and the control voltage in SM 2 decreases by roughly 1V to compensate. At ~ 4.1 s a step of +10 units was taken and the control voltage increases by ~ 2 V. At ~ 4.6 s the

pico-motor mirror was tilted +100 units and the control voltage in SM 2 changes by ~ 6 V. Finally at ~ 6.4 s the pico-motor tilted -100 units, causing the control voltage to decrease by ~ 4 V.

Although the step size of the pico-motor is inconsistent, the manufacturer states that each step should correspond to approximately $0.7 \mu\text{rad}$. 10 steps on the pico-motor is therefore $\sim 7 \mu\text{rad}$ tilt which required a -1 V correction signal on steering mirror 2. This 1 V corresponds to $1/10^{\text{th}}$ of the total angular range of the steering mirror: $73 \mu\text{rad} \times 2 \times 1/10 = 14.6 \mu\text{rad}$. The correction tilt values for all pico-motor steps are shown in table 3.2.

Time (s)	Pico-Motor Step	Steering Mirror Correction
2.8	$-7 \mu\text{rad}$	$-14.6 \mu\text{rad}$
4.1	$+7 \mu\text{rad}$	$+29.2 \mu\text{rad}$
4.6	$-70 \mu\text{rad}$	$+87.6 \mu\text{rad}$
6.4	$+70 \mu\text{rad}$	$-58.4 \mu\text{rad}$

Table 3.2: Auto-alignment correction tilts for Pico-motor adjustment

If it is assumed that the alignment noise is coming predominantly from the pico-motor mirrors then it would also be possible to use these numbers to roughly characterise the amount of noise coming from the pico-motor mirrors.

This preliminary result in figure 3.32 shows that it is possible to utilise the larger range of the pico-motor steering mirrors to artificially increase the range of the auto-alignment system. If the two pico-motors were adjusted simultaneously one would be able to shift the DC value of the control signals and prevent saturation during a long run without losing lock.

More work is needed to investigate how feasible this could be during a data run however. The power in transmission of the cavity was not monitored during this test and it is likely that the pico-motor adjustments couple noise into the cavity power, in a similar way to how the error signal spikes whenever the pico-motors are adjusted. It may be possible to pause data taking whilst the pico-motors are adjusted but this is not ideal in an experiment that wants to run continuously for at least 1.7×10^6 s. Another problem is that, according to the manufacturers, the pico-motors cannot be used as part of an active feedback system due to the gradual degradation of their *slip/stick* mechanism and the noise their operation introduces. This would mean that the control voltages for auto-alignment would require supervision during a run with manual adjustment of the pico-motor mirrors in case saturation of the control signal occurs i.e. a Human In The Loop (HITL). 24 hour supervision is not unusual in experiments such as this; however it does introduce a potential source of human error into the auto-alignment setup which is not desirable.

A fairly straightforward solution would be to replace the $\pm 73 \mu\text{rad}$ auto-alignment

mirrors with higher range mirrors, for instance the KC1-P actuators which have a $\pm 275 \mu\text{rad}$ range and are compatible with the existing AA setup for the RC side. This would of course require the replacement of optics in the input path to the cavity which would cause the input beam to become misaligned, however realignment is fairly straightforward and it has taken as little as 1 hour to successfully realign the cavity after mirrors have been changed in the past.

3.4 Future Work

An improvement of a factor 10 in low frequency cavity noise has been observed in the 250 m cavity when the auto-alignment system is engaged; however there are several areas that should be the focus of future commissioning efforts in order to improve the system further. This section outlines some of the key areas of improvement for the ALPS auto-alignment system going forward.

3.4.1 Higher Range Mirror Actuators

As mentioned in section 3.2.6, the mirror actuators in the RC optics layout have about 1/8 the range of the mirrors in the PC optical layout and it is predicted that the output voltage of the auto-alignment system will saturate after less than 10 hours of operation. Switching the Thorlabs KC1-PZ(/M) actuators ($\pm 73 \mu\text{rad}$ range) with KC1-P actuators ($\pm 275 \mu\text{rad}$ range) would be the simplest solution here since the same connectors can be used to the high voltage amplifier and the PZT layout on the mirror is the same so the same PZT matrix coefficients can be used. This switch may not completely solve any range limitations of the auto-alignment system but it would be a significant improvement and would only be a minor disruption to the experiment.

3.4.2 Linear Combination Matrix

Measurements for the linear combination matrix were taken by injecting a line to each SM in X and Y and measuring the amplitude of the signal on each QPD. A significant asymmetry was observed in the X/Y degrees of freedom:

$$A_X = \begin{vmatrix} 1 & 0.34 \\ 1.27 & 1 \end{vmatrix}, \quad (3.19)$$

$$A_Y = \begin{vmatrix} 1 & 0.16 \\ 0.30 & 1 \end{vmatrix}. \quad (3.20)$$

Unfortunately the digital interface was built for a symmetric response in X and Y and there was no way to use different coefficients for the 2 degrees of freedom.

Instead a unity matrix was used in the tests. This has since been changed and it will now be possible to use the correct combination coefficients.

This should reduce the amount of cross-coupling seen between the two SMs and result in less noise when the loop is closed.

3.4.3 Rotation into PZT Matrix

Only two out of three PZTs in each steering mirror were used when the auto-alignment system was closed. The steering mirrors in the North Right optics setup have the PZTs arranged in an L-shape (see figure 3.20) so one can actuate on the two PZTs on the outermost edge to correct for X and Y degrees of freedom. At the time this was done in order to reduce the complexity of the system.

The PZT rotation matrix is now included in the digital system such that the corner PZT can also actuate on the steering mirror as explained in section 3.2.6. The inclusion of this matrix should double the actuation range of the steering mirrors compared to the range seen in the prior tests described in section 3.3.1 and 3.3.2.

3.4.4 Filter Stages

A single integrator stage was used in all measurements of the auto-alignment loop in section 3.3.1 and 3.3.2. This provides a basic $1/f$ slope on the error signal filtering stage and a maximum unity gain frequency of about 7 Hz. The integrator gain coefficient could not be increased to arbitrarily large values because the system would begin ringing. Preliminary measurements show resonances in the steering mirrors at around 1 kHz which are the likely source of this ringing.

The ringing could be avoided by more aggressive low pass filtering. If a second integrator stage were introduced with a pole frequency at 100 Hz there would be an additional 20 dB of suppression on this resonance at 1 kHz and the overall gain of the filtering could be increased (with a proportional gain stage) thereby increasing the unity gain frequency.

There is no clear target for the unity gain frequency so this is a fairly open ended question. How high the unity gain frequency can be pushed without causing the system to ring will be explored by other commissioners.

3.4.5 Spot Position Control

One of the sub-goals of the commissioning of the ALPS IIc experiment is to be able to remotely control as many key systems as possible. This is not just due to the convenience of controlling the experiment from your home office but also to aid in low noise testing during commissioning of the experiment. There are significant thermal drifts introduced when someone works in the cleanroom, which have been

seen to misalign the beam into the cavity. There are also misalignments and loss of cavity lock introduced by opening the door to the cleanroom. For this reason, physical presence in the cleanrooms should be kept to a minimum when low noise testing is desired.

In the case of auto-alignment the beam needs to be centered on the QPD sensors in order to generate error signals. This is achieved by steering mirrors LTMT24 and LTM28 (see figure 3.8). However these steering mirrors do not have pico-motors or PZTs attached and cannot be operated remotely. Someone therefore needs to physically be in the clean room in order to centre the beam spot on the QPD cameras which introduces the aforementioned thermal drift. Realistically this improvement would only be useful in the commissioning effort. Once a good input beam position to the cavity is found (based on mirror losses/scattering) the spot position on the cameras can be adjusted once and should not change.

3.4.6 Digital to Analogue Converter and High Voltage Amplifiers

On the RC side, the digital to Analogue Converter (DAC) is designed to output voltages between $-10\text{ V} \rightarrow 10\text{ V}$; however the High Voltage Amplifier (HVA) can only accept $0 \rightarrow 10\text{ V}$. Appropriate limits and offset are used within the cRIO signal processing to prevent saturation and clipping on the HVA; however there is effectively twice as much noise on the DAC than we should have based on the $0 \rightarrow 10\text{ V}$ requirement.

The impact of this DAC noise on the auto-alignment system has not been investigated thus far.

The simplest and cheapest solution would most likely be to change the range of the DAC which would require some new custom electronics.

3.4.7 General Debugging

There are various bugs and glitches in the labview code that is used to digitally process the data in the RC auto-alignment system.

Sampling/Processing Frequency: The sampling/processing frequency of the cRIO seems to vary between different channels which makes quantitative analysis difficult (see figure 5.15 in appendix 5.11). This is why many time series in sections 3.3.1 and 3.3.2 are displayed in counts rather in units of time. An investigation into this is currently ongoing.

Data saving: In the final design of the ALPS experiment, auto-alignment data will be saved and monitored using a custom built monitoring system called DOOCS. However, there is currently no way to send data to the DOOCS system so commissioners at ALPS are currently building a connection to pick off the error and control

signals at the input and output of the cRIO to send to the DOOCS monitoring system.

During commissioning the data was being saved internally on the cRIO; however it seems that there is a glitch in the data processing where values are randomly saved into the incorrect data column.

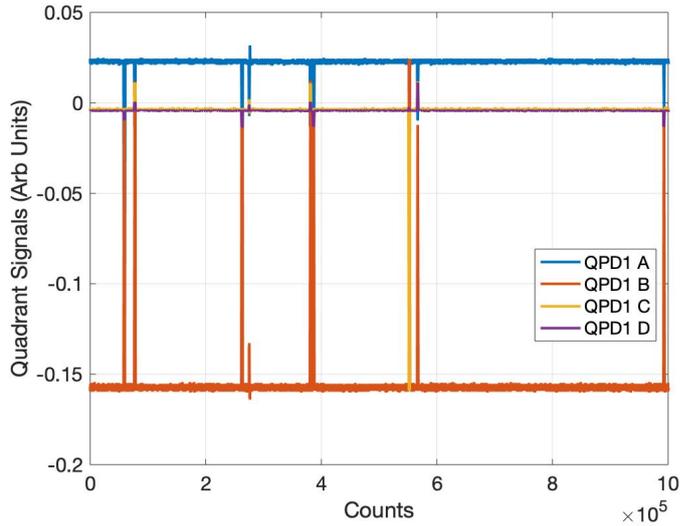


Figure 3.33: Demodulated signals from the 4 quadrants of QPD 1 during the 2 minute test described in section 3.3.1. The sampling frequency of the cRIO is not certain but is expected to be roughly 16 kHz. Glitches can be seen where the error signals instantaneously change value and occasionally swap values with each other. The cavity was locked throughout the test. As shown in figure 5.15 there is no corresponding change in the control signal that can be seen.

The glitches occur instantaneously and in all 4 quadrants which results in a sudden change in the error signals at the same frequency as the cRIO sampling frequency. There is no change that can be seen in the control signals which may indicate that this could be a glitch in the data *saving* rather than in the data *processing*, however this has not yet been verified and should be the subject of further investigation.

If these glitches are only seen in the data saving then they can be avoided entirely once the DOOCS data readout is in place for the auto-alignment signals.

3.4.8 NL Setup

Now that the auto-alignment system in the RC side of the experiment has been shown to work, commissioning efforts should shift to include the setup of the twin system in the PC side. As mentioned in section 3.2.4, 3.2.5 and 3.2.6, this system achieves the same result as the RC system but uses different techniques to achieve it. Notably the signal processing in the PC side is all performed with analogue electronics as opposed to digital controllers. Some time was spent electronically

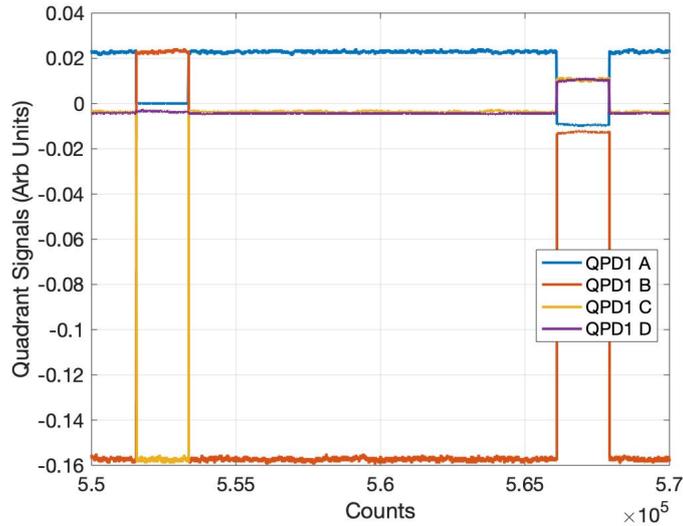


Figure 3.34: Zoomed in view of the signals from figure 3.33. Glitches can be seen at 5.52×10^5 counts and 5.67×10^5 counts and lasts about 0.03×10^5 counts before reverting back to the original value. Assuming a 16 kHz sample frequency each glitch is about 0.2s long.

characterising the analogue servos for the PC auto-alignment but there have been no tests performed with them at the ALPS site as of writing this thesis.

The optics for the auto-alignment system (Gouy phase telescopes and sensing QPDs) are being installed as of writing, however the PC mirrors will not be installed at the ALPS site until mid-to-late 2022 so no formal testing of the PC auto-alignment system will be able to take place until then.

3.5 ALPS Concluding Remarks

An auto-alignment system based on the well known Pound-Drever Hall and differential wavefront sensing techniques has been designed and installed at the ALPS IIc site. The auto-alignment system has been tested on a 250m test cavity, whilst a twin system that uses analogue electronics is in the process of being commissioned. The optical layout for the NL and NR optical tables has also been designed and commissioned and the NR optics were used to successfully lock the 250m optical cavity.

In section 3.3.1 it was shown that the auto-alignment system was able to suppress low frequency (< 10 Hz) cavity noise by up to an order of magnitude compared to the open loop scenario. It was also shown in section 3.3.2 that there is significant loss of the circulating cavity power over the course of 90 minutes when the auto-alignment loop is not engaged. The cause of this drift may not be entirely due to cavity alignment but figure 3.29 showed that this drift was greatly reduced during

a comparable measurement time once the auto-alignment loop was engaged. Figure 3.31 also indicates that the auto-alignment would be able to maintain lock over approximately a 10 hour measurement, although this has not been verified.

As shown in section 3.4 there are still significant improvements to be made to the auto-alignment system and efforts are underway to improve the suppression of higher frequency alignment noise and to commission the twin system that will actuate on the PC side of the experiment. Now that the auto-alignment system has been shown to work, I hope that my colleagues can build off this foundation of work to more rapidly improve the system beyond what it has already achieved.

Chapter 4

Conclusion

This thesis has outlined advanced systems that were designed and implemented in two high precision interferometry experiments: the LIGO gravitational wave detector (chapter 2) and the ALPS dark matter search (chapter 3).

The ALPS dark matter search is still in its commissioning phase and is currently expected to begin the first science run at the end of 2022. As part of this commissioning effort, an auto-alignment system has been designed and tested and was shown to reduce low frequency cavity noise by up to an order of magnitude. It has also been shown to maintain DC alignment of the laser beam to the optical cavity eigenmode over a period of 2 hours, an essential feature for the long measurement times of the ALPS experiment. I hope that in the upcoming science run the auto-alignment system is used to set new limits on axion-like particles and perhaps if we are very fortunate, ALPS IIc will be the experiment that solves the decades old question of dark matter in our universe.

In the LIGO gravitational wave search, the experiment has already been successfully detecting gravitational waves for over 7 years. The goal in this project was to further improve the detector's sensitivity as part of the *A+* upgrade, in order to increase the likelihood of detecting astronomical binary mergers (specifically binary neutron stars). This goal was ultimately successful and new photodetector electronics have been designed and installed that exhibit dark noise levels 6.5 times lower than the previous generation readouts at 1 kHz. Based on noise floor projections, the new readouts give a predicted improvement in the DARM sensitivity of 10% between 600 Hz and 3 kHz in the upcoming observing run in 2023. I look forward to seeing the results from the next observing run and for the first detection that will be made using these new readouts.

Chapter 5

Appendix

5.1 Gaussian Beams

The simplest model for laser propagation is to model the beam as a series of non-diverging plane wave fronts propagating in the same direction. However this approximation breaks down quickly when considering laser propagation over larger distances or for small spot sizes. The propagation of phase fronts described by the paraxial Gaussian equations is a more accurate representation of the laser as it travels.

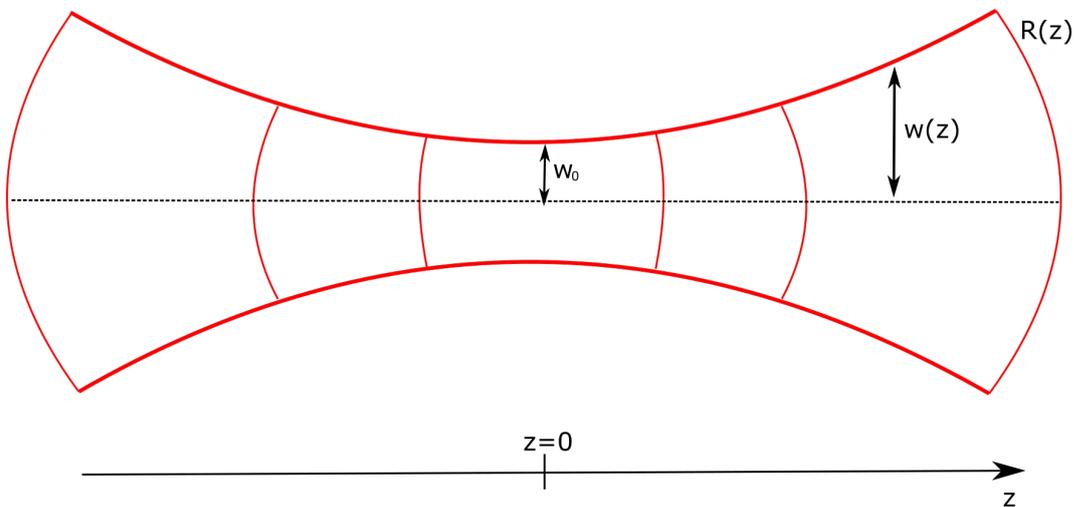


Figure 5.1: Diagram showing phase front propagation described by the paraxial Gaussian beam formula [51]. The wave is travelling along the z-direction and has its narrowest point at $z=0$ with radius, w_0 . The radius of curvature at this point is $R(z) = \infty$. z_R is the Rayleigh range and corresponds to the point at which the beam area has doubled, i.e. $w(z) = \sqrt{2}w_0$

For a beam with its waist, w_0 , located at the origin $z = 0$, the beam width after a distance z may be described by equation 5.1

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (5.1)$$

where z_R is the Rayleigh range and is defined by $z_R = \frac{\pi W_0^2}{\lambda}$.

The radius of curvature of the beam after a distance z may be described by equation 5.2

$$R(z) = z \left[1 + \left(\frac{z_R}{z}\right)^2 \right]. \quad (5.2)$$

Equations 5.1 and 5.2 may be described together by the q-value of a Gaussian beam at a point z [51]:

$$\frac{1}{q(z)} = \frac{1}{R(z)} - j \frac{\lambda}{\pi w^2(z)}. \quad (5.3)$$

The q-value will be transformed whenever it passes through an optical component (or free space) as described in appendix 5.2.

5.2 Mode Matching a Fabry-Perot Cavity

An optical cavity is an arrangement of mirrors that enables a beam to continuously circulate its optical path.

There are a few conditions that must be satisfied in order to achieve optical resonance. The first and simplest is the longitudinal resonance condition. This states that the cavity length must be equal to a half integer of the wavelength of the light

$$L = \frac{n\lambda}{2}, \quad (5.4)$$

or equivalently when the cavity roundtrip length is equal to an integer number of laser wavelengths. This is achieved using the well-established Pound-Drever-Hall locking scheme [53] to actuate on either cavity length or laser frequency.

As well as locking the length of the cavity it is important to match the beam profile to the shape of the cavity mirrors such that the transverse resonance condition is also satisfied. This process, known as *mode matching*, takes into account the factors involved in laser beam propagation in the Gaussian approximation (see section 5.1)

Incorrect mode matching will result in higher order mode coupling [153] and an overall reduction to the total power circulating in the fundamental mode within the

cavity. The beam parameter of the input laser is manipulated using either lenses or curved mirrors such that the input beam matches the parameters of the cavity.

ABCD matrices are used to calculate the final complex beam parameter q' from its initial parameter q shown in equation 5.3.

An arbitrary optical component will have elements within its corresponding transformation matrix, M :

$$M = \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \quad (5.5)$$

that will transform the complex beam parameter q according to:

$$q' = \frac{Aq + B}{Cq + D}. \quad (5.6)$$

A few common matrices are shown below:

$$M_1 = \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix} \textit{Free Space} \quad (5.7a)$$

$$M_2 = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix} \textit{Lens} \quad (5.7b)$$

$$M_3 = \begin{bmatrix} 1 & 0 \\ -\frac{2}{R_m} & 1 \end{bmatrix} \textit{Curved Mirror} \quad (5.7c)$$

where M_1 describes free space propagation along a distance, d , M_2 describes a lens with focal length, f and M_3 describes reflection from a mirror with radius of curvature, R_m . To model propagation through multiple components, the matrices can simply be multiplied. For example the ABCD matrix for propagating through a free space d and then a lens, f is given by:

$$M_2 M_1 = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix} \begin{bmatrix} 1 & d \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & d \\ -\frac{1}{f} & 1 \end{bmatrix}. \quad (5.8)$$

As long as the mirrors that make up the Fabry-Perot cavity contain a stable eigenmode it is possible to construct a transformation matrix using a series of optical components and spaces to match the beam's waist size and position to that of the cavity eigenmode.

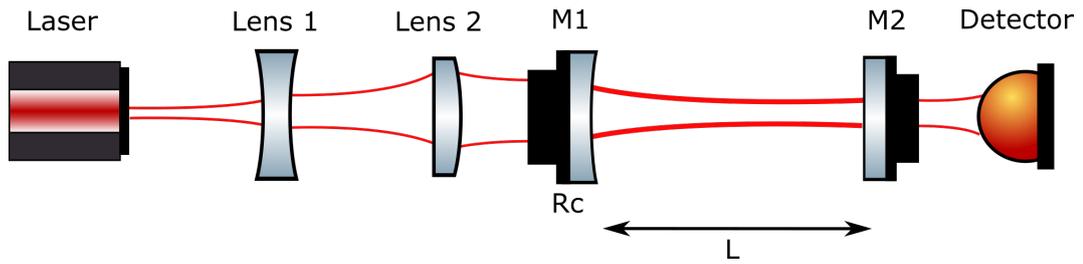


Figure 5.2: Two lenses are used to form a mode matching telescope that matches the beam parameter of the Fabry-Perot cavity. The cavity eigenmode is determined by the radius of curvature of the two cavity mirrors and the separation between the two mirrors.

5.3 DCPD pin layout

For reference the pin layout of the circuit and the voltage level used are shown in figure 5.3, table 5.1 and table 5.2.

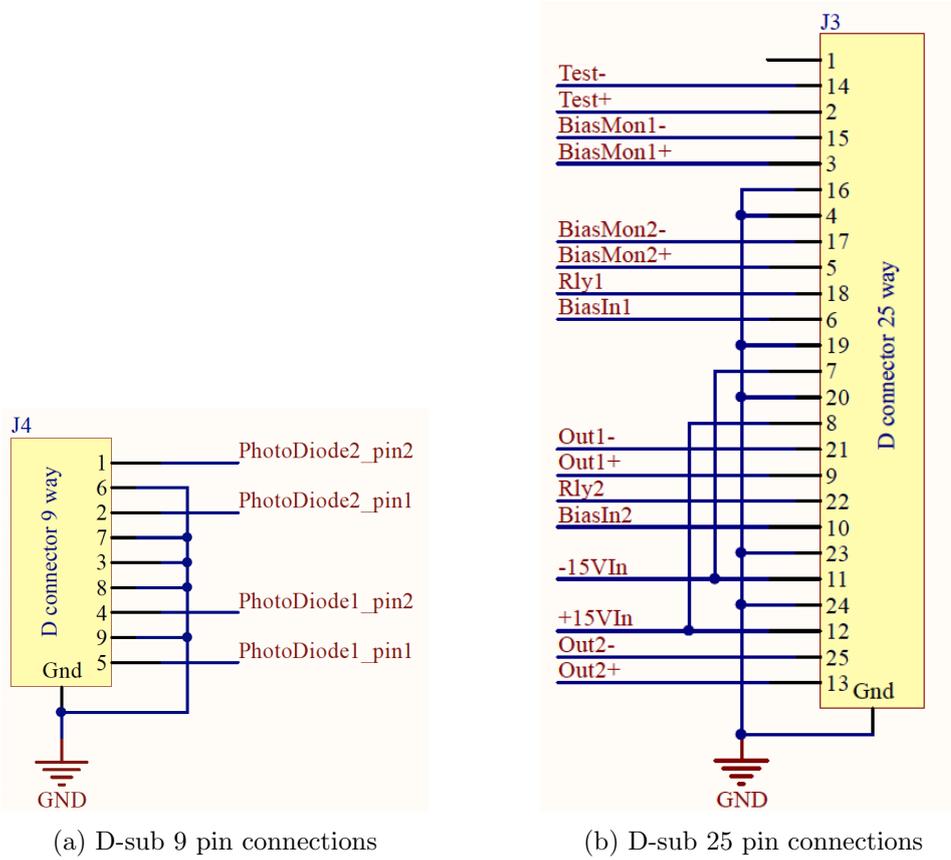


Figure 5.3: Pin layout for the D-sub 9 and D-sub 25 connectors

For the D-sub 25 connector the following input voltages were used for testing.

Pin Number	Pin Name	Input Voltage Used
2	Test+	1 V _{rms}
4, 16, 19, 20, 23, 24	Ground	GND
14	Test-	GND
18	Rly1	5 V _{DC}
6	BiasIn1	100 mV _{rms}
22	Rly2	5 V _{DC}
10	BiasIn2	100 mV _{rms}
11	-15VIn	-15 V _{DC}
12	+15VIn	+15 V _{DC}

Table 5.1: Input pin voltages used for D-sub25

For the D-sub 9 connector the following inputs were used for testing.

Pin Number	Pin Name	Input Voltage Used
1	Photodiode 2 pin 2	GND
2	Photodiode 2 pin 1	-10 dbm
4	Photodiode 2 pin 2	GND
5	Photodiode 2 pin 1	-10 dbm

Table 5.2: Input pin voltages for the D-sub 9

5.4 PCB Photos

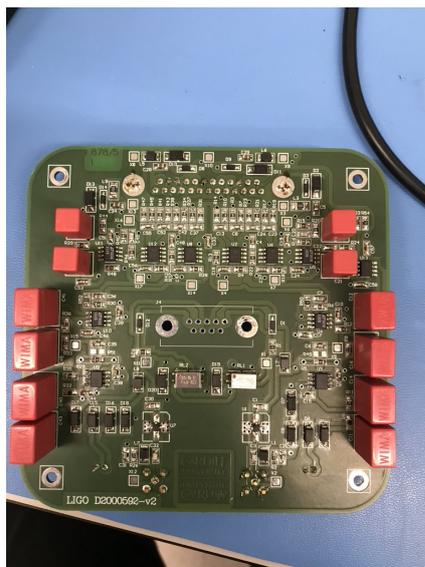


Figure 5.4: Top view of the DCPD PCB. The board has two independent channels and is symmetric between the left and right hand sides.

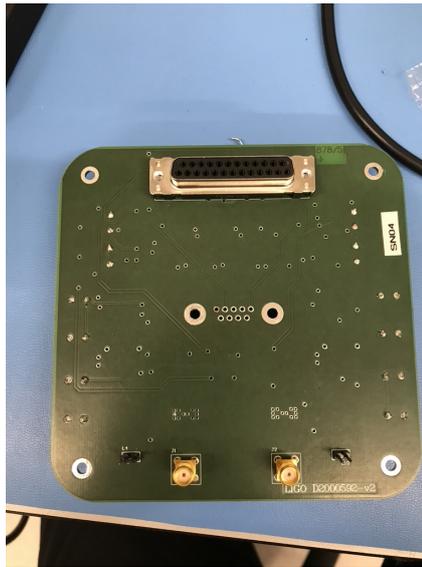


Figure 5.5: Bottom view of the DCPD PCB. The SMA connectors provide the RF output. The D-sub 9 connector in the middle provides the photodiode input and the D-sub 25 provides all other connections.

5.5 VIRGO-Style Inductor

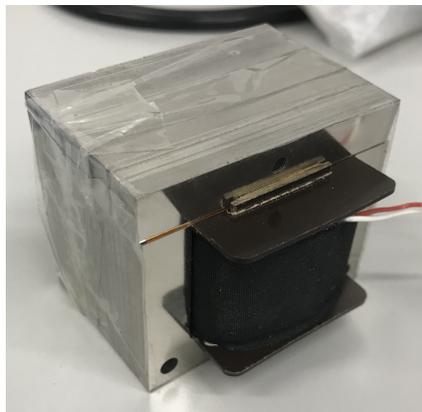


Figure 5.6: Photo of the inductor used in the VIRGO-style readout. This inductor has the same core as the inductor used in the LIGO readout but has double the wire windings which increases the inductance from ~ 2.6 H to ~ 10.6 H

5.6 Variations of the transfer function with DC Current

The effect that only an inductance change would have on the overall transfer function was explored. As the DC current through the inductor is increased saturation of the inductor core occurs and the inductance decreases. The following plots use the 50 mA case as a baseline for comparison.

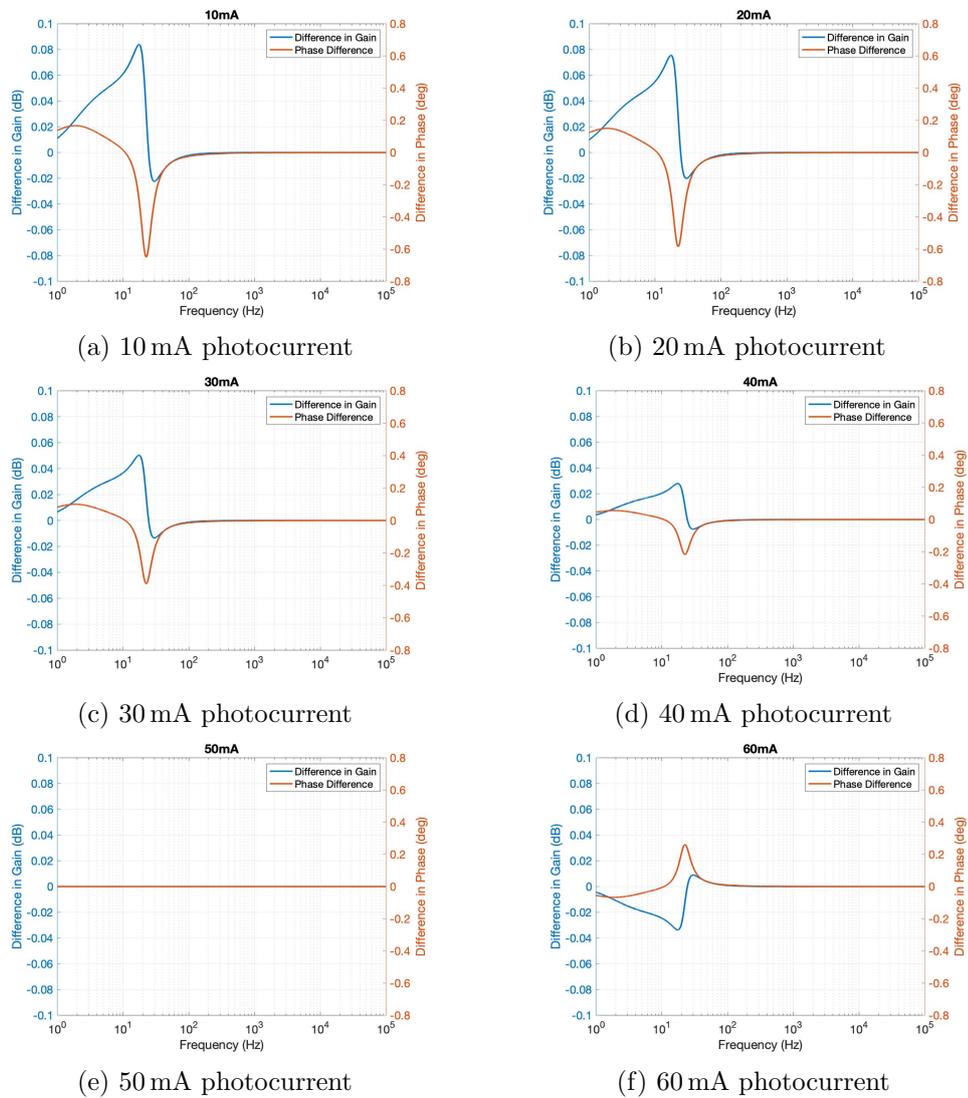


Figure 5.7: LISO [73] simulation for the change in transfer function for various photocurrent (and thus inductance) levels. The transfer function for a 50 mA DC photocurrent is used as the reference.

5.7 Full Circuit Schematic

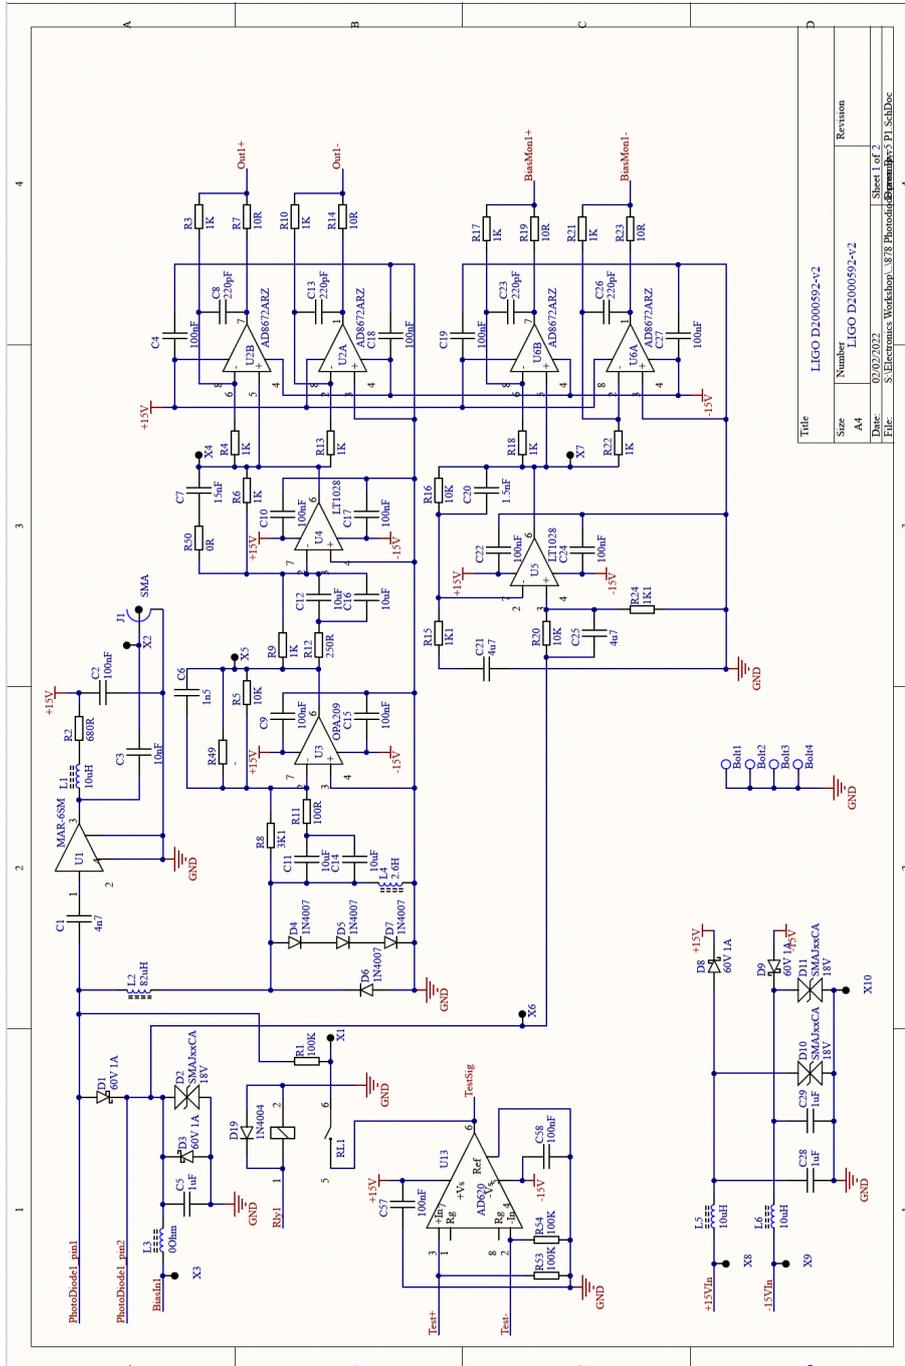


Figure 5.8: First page of the DCPD electronics schematic. **Far Left:** the test relay input. **Upper Left:** photodiode input. **Upper middle:** the RF readout. **Centre left:** $L4$, the inductor that directs photocurrent preferably to ground (DC) and to the transimpedance stage (U3) and **Upper Right:** additional whitening (U4) and the differential outputs drives. **Lower Right:** bias voltage monitor with differential output drive.

5.8 Inductor Mounting

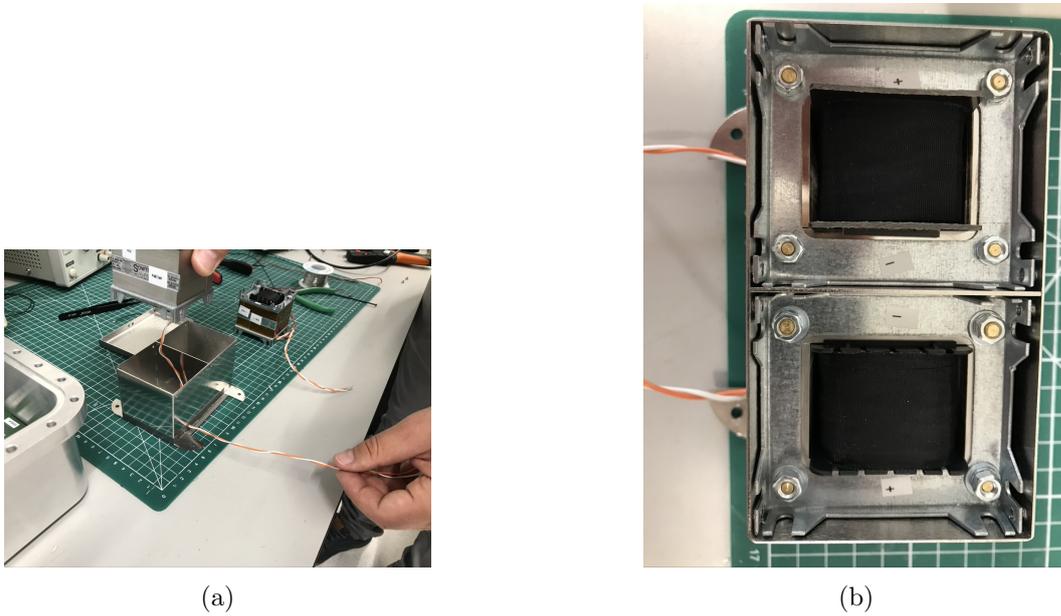


Figure 5.10: **a)** Feeding the inductor wires through the mu-metal shield. **b)** Orientation of the inductors. Labelled on them are the polarity of the induced signal from magnetic pickup.

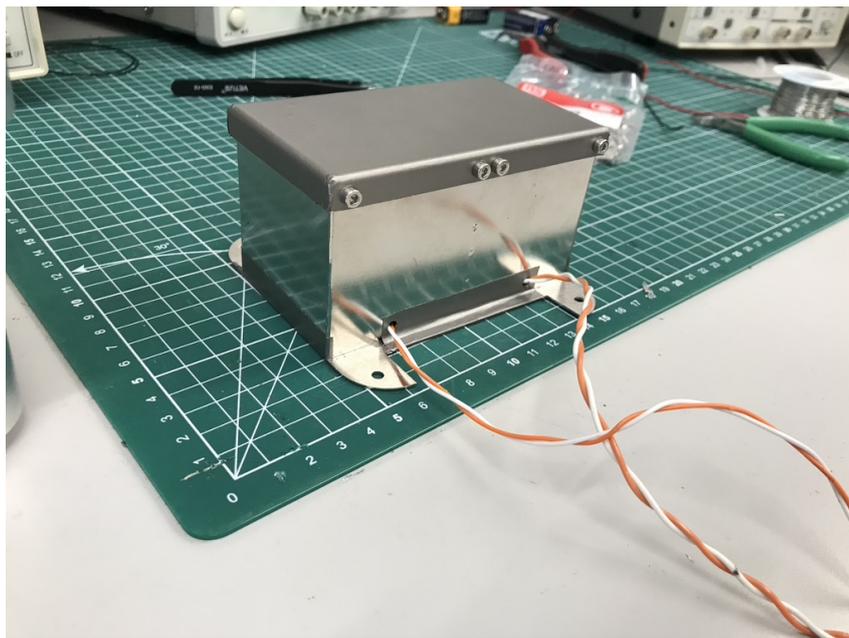


Figure 5.11: Inductor shield with inductors mounted

5.9 Complex zero/pole Formulation

In figure 2.23 there are components L_4 and C_{11} that have imaginary impedances. To simplify the DCPD model we can use the complex impedances Z_{eff} and Z_{llel} .

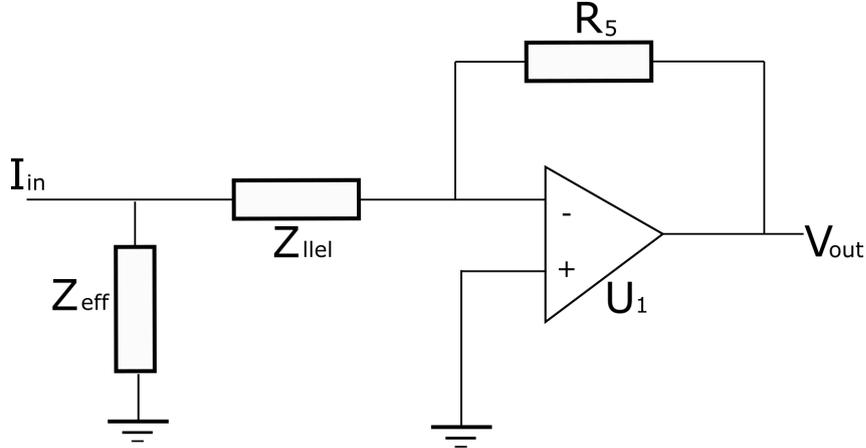


Figure 5.12: Same as the DCPD schematic in figure 2.23 expect some components are now represented by complex impedances Z_{eff} and Z_{llel} .

Z_{eff} is the model of the inductor due to the ohmic resistance and parasitic capacitance that is shown in figure 2.44:

$$Z_{eff} = \frac{Z_{cl} (\mathbf{j}2\pi fL + R_L)}{Z_{cl} + R_L + \mathbf{j}2\pi fL}, \quad (5.9)$$

where $z_{cl} = \frac{1}{\mathbf{j}2\pi fC_L}$ is the complex impedance of the parasitic capacitance of the inductor with capacitance C_L , R_L is the ohmic resistance of the copper wire in the inductor, f is the frequency, L is the inductance of the ideal inductor.

Z_{llel} is given by:

$$Z_{llel} = \frac{R_8 (Z_{C11} + R_{11})}{z_{c11} + R_{11} + R_8}, \quad (5.10)$$

where $Z_{c11} = \frac{1}{\mathbf{j}2\pi fC_{11}}$ is the complex impedance of C_{11} .

The current passing through the transimpedance stage can be found using:

$$I_{llel} = \frac{Z_T}{Z_{llel}}, \quad (5.11)$$

where Z_T is the total impedance of Z_{llel} and Z_{eff} in parallel with each other.

The voltage across the transimpedance resistor is then simply given by:

$$|V_{out}| = |I_{lcl}| R_5, \quad (5.12)$$

where R_5 is the transimpedance resistor and we are now taking the absolute value of the current which is complex.

This output voltage is shown below:

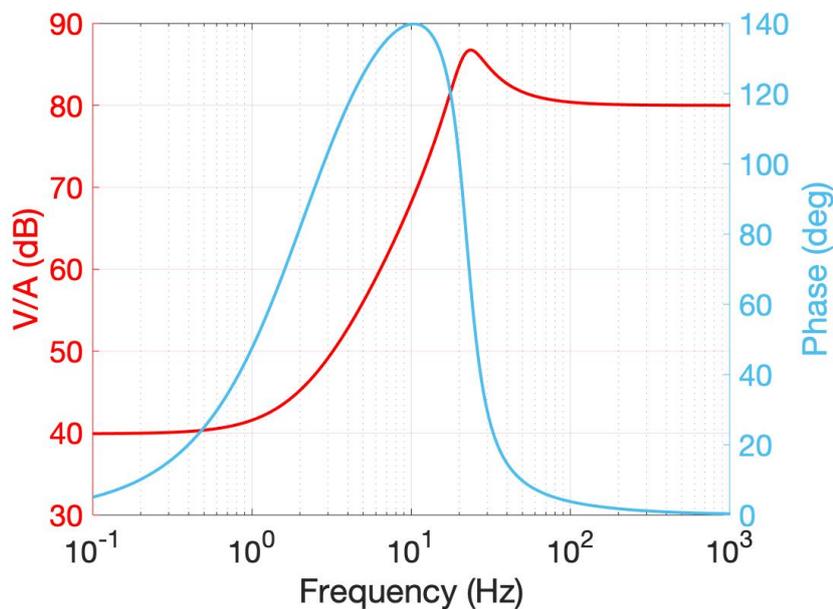


Figure 5.13: Model of the transfer function from the photodiode input pin to the output of the transimpedance stage. This model is only used to illustrate the complex zero/pole in the transimpedance stage and does not include the low pass filter with a pole at 10 kHz or the contributions of the second amplifier stage that is present in the final design.

5.10 Auto-alignment Ringing Error Signals

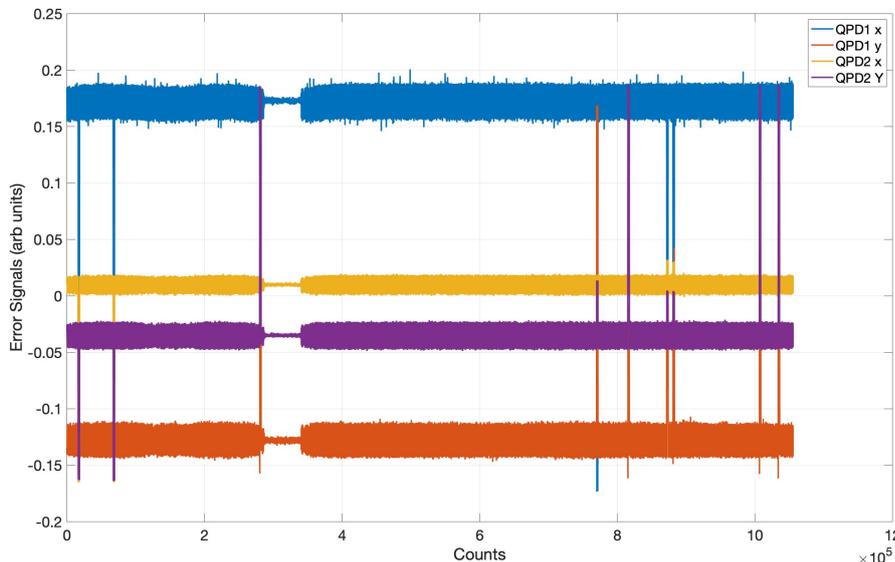


Figure 5.14: Ringing error signals from QPD1 and 2 in X and Y during a 2 minute test of the AA loop. The gain coefficients of all integral gain stages were increased until the system was on the edge of ringing. It can be seen that the system briefly stops ringing at 3×10^8 counts and then quickly resumes ringing once more. During closed loop operation the gain integrals would be set slightly lower than the value used for this test such that the system was on the edge of ringing.

5.11 Auto-Alignment Sampling Frequency

The error signals and control voltage over the course of the 2 minute window are shown below:

Despite both the control and error signals being saved internally on the cRIO (and being saved simultaneously) they have different numbers of total counts (10×10^5 and 13×10^5). Even more confusingly, a stopwatch was used to check the length of the actual time series and it was 120 ± 1 s; however we can calculate the apparent duration of the time series based on the counts and the expected sample rate (16 kHz) and see that the error signal time series is apparently only ~ 62 s long and the control time series is ~ 81 s long. This issue is still not understood and it is hoped that it is a feature of the internal data saving capabilities of the cRIO which will not be used during science runs.

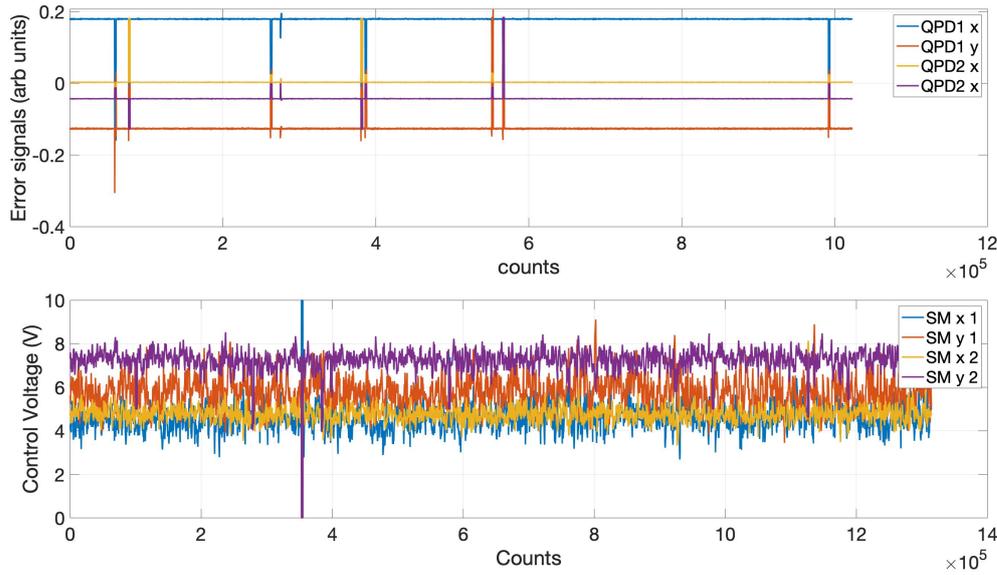


Figure 5.15: **Top:** Time series of error signals QPD 1 X/Y and QPD2 X/Y data was saved internally on the cRIO. **Bottom:** Control voltages for mirror 1 X/Y and mirror 2 X/Y before amplification, data was saved internally on the cRIO.

5.12 Auto-Alignment Two minute time-series

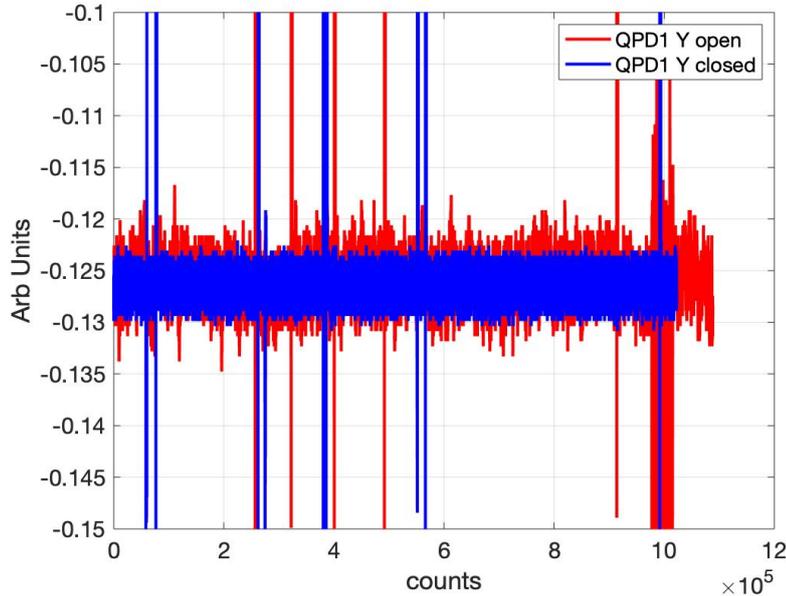


Figure 5.16: Error signal for QPD 1 Y for open and closed AA loop over a 2 minute duration. A noticeable improvement in the error signal noise is seen when the loop is closed. The set point of the error signal was -0.0126. The x-axis is displayed in counts rather than in seconds. This is due to uncertainty in the frequency that the cRIO system processed and saved data (see appendix 5.11 for more information).

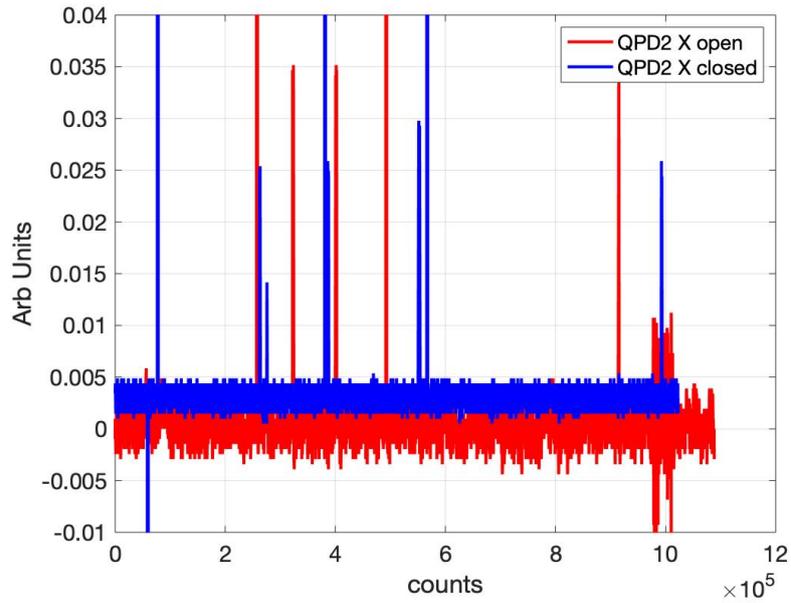


Figure 5.17: Error signal for QPD2 X for open and closed AA loop over a 2 minute duration. A noticeable improvement in the error signal noise is seen when the loop is closed. The set point of the error signal was 0.004. The x-axis is displayed in counts rather than in seconds. This is due to uncertainty in the frequency that the cRIO system processed and saved data (see appendix 5.11 for more information).

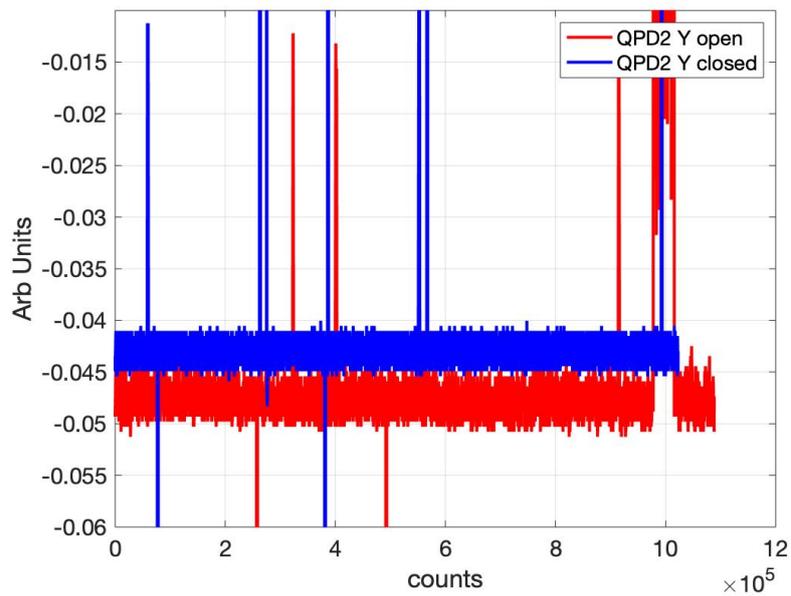


Figure 5.18: Error signal for QPD2 Y for open and closed AA loop over a 2 minute duration. The set point of the error signal was -0.043. A shift in the DC level is seen when the AA system is closed. There is a smaller improvement in the noise seen in this degree of freedom. The x-axis is displayed in counts rather than in seconds. This is due to uncertainty in the frequency that the cRIO system processed and saved data (see appendix 5.11 for more information).

5.13 Auto-alignment Two hour control signals

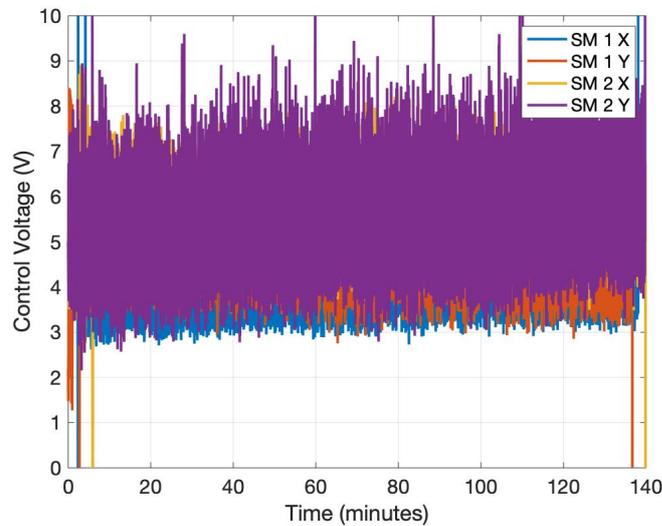


Figure 5.19: All 4 control voltages during the 2 hour lock. The controls for SM 1 X/Y and SM 2 X all exhibited minimal drift during the test whereas SM 2 Y showed a drift in alignment around 1.5 V

5.14 ALPS Photos

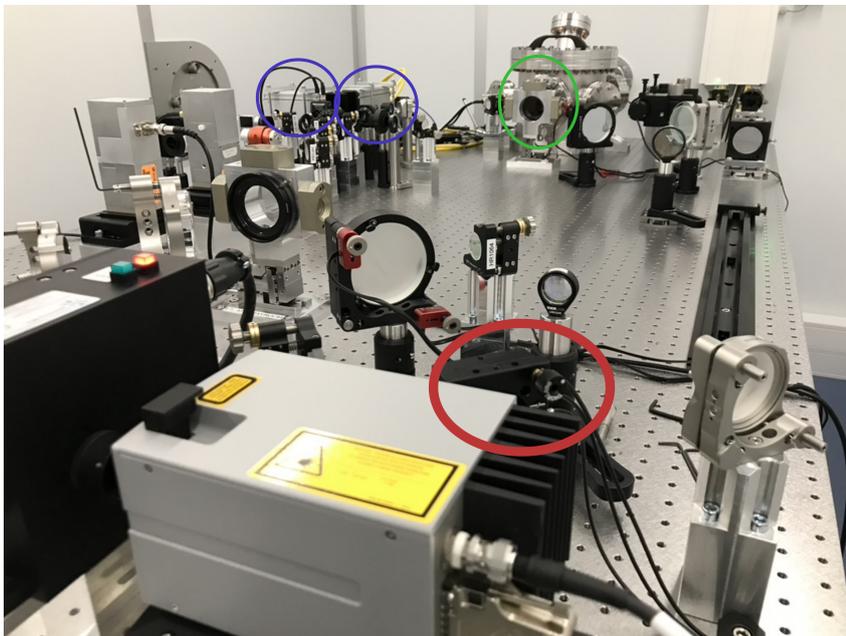


Figure 5.20: The optics table used in the RC side of the experiment. **Red:** One of the auto-alignment steering mirrors, the other steering mirror is on the left just out of frame. **Blue:** the two QPD sensors used in auto-alignment. **Green:** the input mirror to the 250 m cavity. In the final configuration this will be the input mirror to the regeneration cavity.

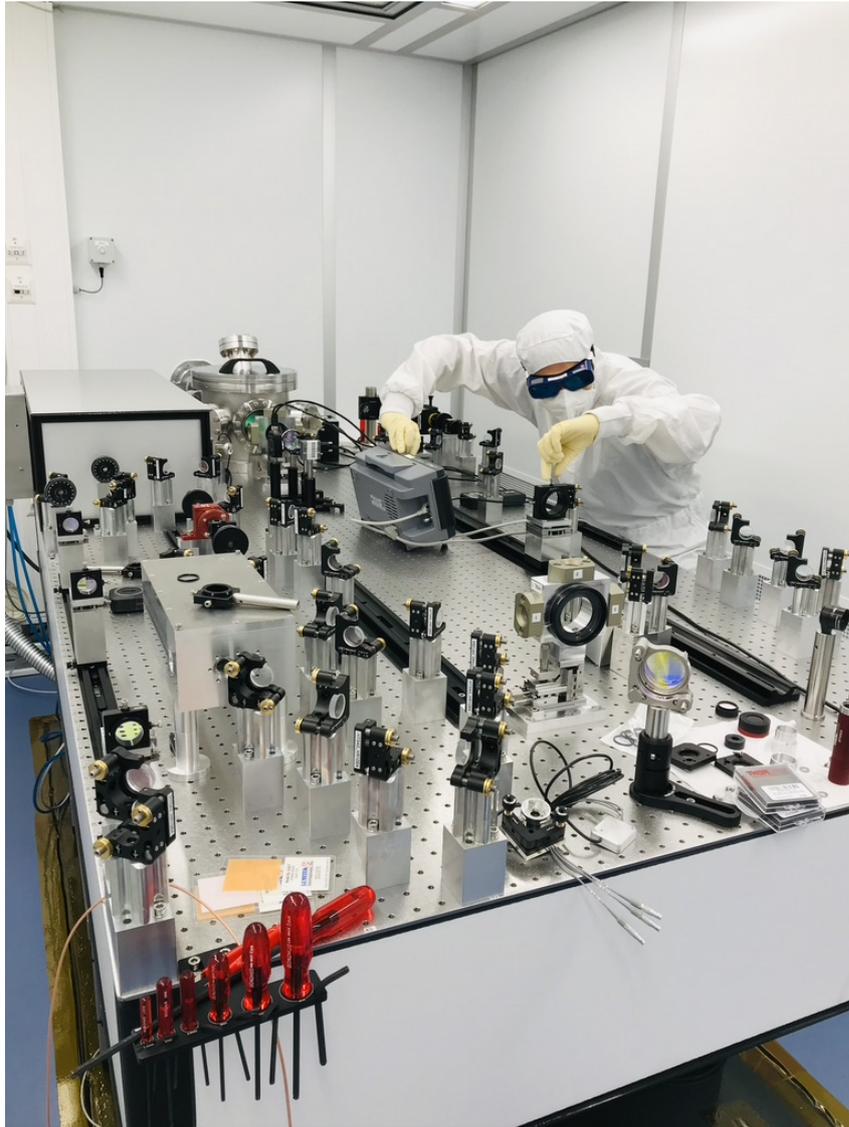


Figure 5.21: Hard at work on the PC optical table.

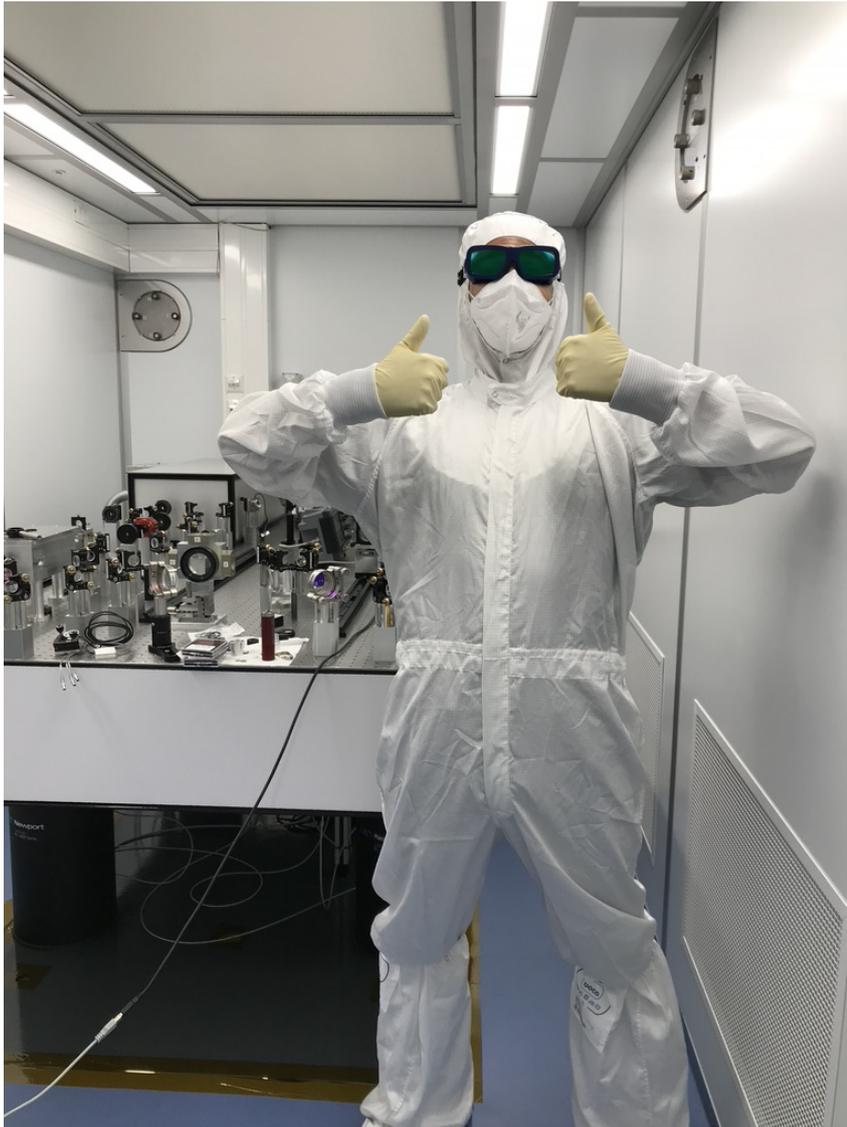


Figure 5.22: A successful day! :)

5.15 LIGO Filtering Stages

The CDS system has an analogue to digital converter (ADC) that saves data as ‘counts’ rather than as Volts. There is also an additional whitening board in between the DCPD and the ADC that contains a whitening filter. Both of these effects need to be taken into account in order to compare the data saved on the CDS to the data taken directly from the DCPD. The counts to volts factor is 8192 whilst the other inverse filters are shown below.

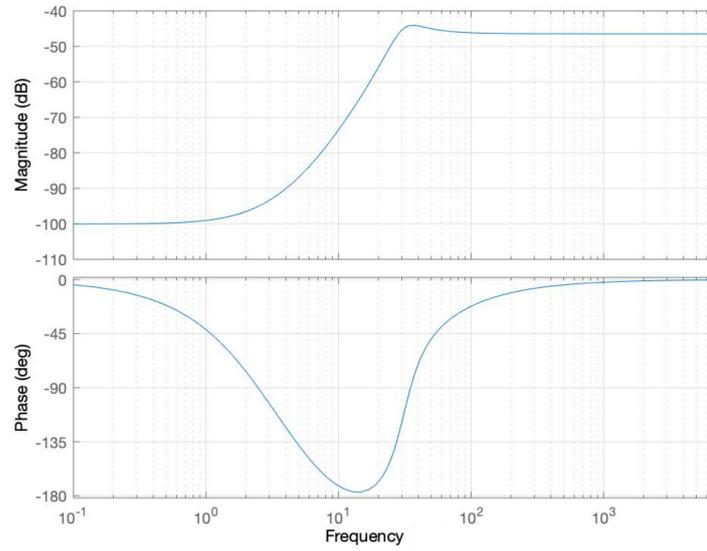


Figure 5.23: The inverse DCPD filter

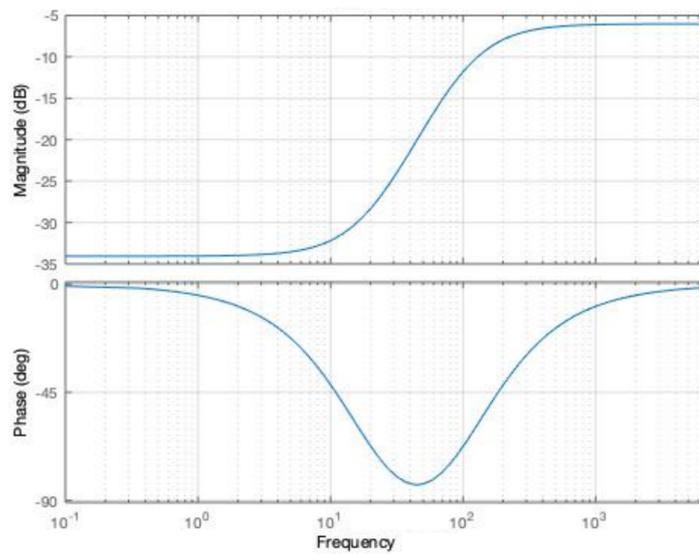


Figure 5.24: The inverse whitening filter

Bibliography

- [1] Thomas Young. The Bakerian Lecture. Experiments and calculations relative to physical optics. <https://doi.org/10.1098/rstl.1804.0001>, 1804.
- [2] Claus Jönsson. Electron Diffraction at Multiple Slits. *AJP* 42, 4, 1974.
- [3] A. Michelson. Experimental determination of the velocity of light. 1880.
- [4] Michelson A A and Morley E W. On the relative motion of the Earth and the luminiferous ether. *Am. J. Sci. (3rd series)* 34 333–45, 1887.
- [5] Harvey R Brown. Michelson, FitzGerald and Lorentz: the origins of relativity revisited. 2003.
- [6] Perot A. and Fabry Charles. On the Application of Interference Phenomena to the Solution of Various Problems of Spectroscopy and Metrology. *apj*, 9: 87, feb 1899. doi: 10.1086/140557.
- [7] G Zavattini et al. The PVLAS experiment: detecting vacuum magnetic birefringence. *J. Phys.: Conf. Ser.* 442 012057, 2013.
- [8] F. Della Valle et al. Extremely long decay time optical cavity. *Vol. 22, Issue 10, pp. 11570-11577*, 2014.
- [9] Cahillane. C et al. Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four. *Galaxies*, 10, 36, 2022.
- [10] Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.* 116, 061102, 2016.
- [11] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916.
- [12] Diego Bersanetti et al. Advanced Virgo: Status of the Detector, Latest Results and Future Prospects. *Universe* 2021, 7, 322, 2021.
- [13] K L Dooley and the LIGO Scientific Collaboration. Status of GEO 600. *J. Phys.: Conf. Ser.* 610 012015, 2015.

-
- [14] KAGRA collaboration. KAGRA: 2.5 generation interferometric gravitational wave detector. *Nature Astronomy volume 3, pages 35–40*, 2019.
- [15] Michele Maggiore et al. Science case for the Einstein telescope. *JCAP03(2020)050*.
- [16] David Reitze. Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO. *arXiv:1907.04833 [astro-ph.IM]*, 2019.
- [17] M. Armano et al. Beyond the Required LISA Free-Fall Performance: New LISA Pathfinder Results down to 20 μ Hz. *Phys. Rev. Lett. 120, 061101*, 2018.
- [18] Charles W Misner, Kip S Thorne, and John Archibald Wheeler. *Gravitation. New York, NY : Freeman*, 2008.
- [19] P. Saulson. *Interferometric Gravitational Wave Detectors. World Scientific*, 2017.
- [20] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. DOI:10.1093/acprof:oso/9780198570745.001.0001.
- [21] Pisin Chen. Gravitational beamstrahlung. *Modern Phys. Lett. A Vol. 06, No. 12, pp. 1069-1075*, (1991).
- [22] D. Cuomo et. al. *Experimental Gravitational Physics. ed. P. F. Michelson and E.-K. Hu and G. Pizzella World Scientific, Singapore*, (1987).
- [23] E. Teller. talk to NATO ASI on Cosmology and Particle Physics. *Erice*, 1993.
- [24] Lars Bildsten. Gravitational radiation and rotation of accreting neutron stars. *The Astrophysical Journal, 501:L89–L93*, 1998.
- [25] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Astron. Soc. 331, 203–220*, 2002.
- [26] Nihan S. Pol et a. Astrophysics Milestones for Pulsar Timing Array Gravitational-wave Detection. *ApJL 911 L34*, 2021.
- [27] A. Burrows and D. Vartanyan. Core-collapse supernova explosion theory. *Nature volume 589, pages29–39*, 2021.
- [28] H.-Th. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller. Theory of core-collapse supernovae. *Physics Reports Volume 442, Issues 1–6, Pages 38-74*, April 2007.
- [29] Christian D. Ott, Adam Burrows, Eli Livne, and Rolf Walder. Gravitational waves from axisymmetric, rotating stellar core collapse. *The Astrophysical Journal, 600:834–864*, 2004.
-

- [30] LIGO collaboration. Search for core-collapse supernovae in initial LIGO and VIRGO data. *ligo.org*, accessed:2022.
- [31] Alison J. Farmer and E. S. Phinney. The Gravitational Wave Background from Cosmological Compact Binaries. *arXiv:astro-ph/0304393*, 2003.
- [32] Maggiore. Stochastic backgrounds of gravitational waves. *arXiv:gr-qc/0008027*, 2000.
- [33] P. de Bernardis et al. A Flat Universe from High-Resolution Maps of the Cosmic Microwave Background Radiation. *arXiv:astro-ph/0004404*, 2000.
- [34] K. Danzmann et al. Laser Interferometer Space Antenna. *arXiv:1702.00786 [astro-ph.IM]*, 2017.
- [35] S'ebastien Clesse, Juan Garc' ra Bellido, and Stefano Orani. Detecting the Stochastic Gravitational Wave Background from Primordial Black Hole Formation. *arXiv:1812.11011 [astro-ph.CO]*, 2018.
- [36] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. *arXiv:2111.03606 [gr-qc]*, 2021.
- [37] Abbott et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. *The Astrophysical Journal Letters*, 896:L44 (20pp), 2020.
- [38] M. Zevin et al. Exploring the Lower Mass Gap and Unequal Mass Regime in Compact Binary Evolution. *ApJL* 899 L1, 2020.
- [39] Luiz L. Lopes and Debora P. Menezes. On the nature of the mass-gap object in the GW190814 event. *arXiv:2111.02247 [hep-ph]*, 2021.
- [40] Dennis Overbye. A Black Hole's Lunch Provides a Treat for Astronomers. *The New York Times*, June 24, 2020.
- [41] LIGO and VIRGO collaboration. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.* 119, 161101, 2017.
- [42] LIGO and VIRGO collaboration. Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *The Astrophysical Journal Letters*, 848:L13 (27pp), 2017.
- [43] LIGO and VIRGO collaboration. Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run. *Phys. Rev. Lett.* 118, 121101, 2017.

-
- [44] LIGO and VIRGO collaboration. Search for post-merger gravitational waves from the remnant of the binary neutron star merger GW170817. *arXiv:1710.09320 [astro-ph.HE]*, 2017.
- [45] LIGO and VIRGO collaboration. A gravitational-wave standard siren measurement of the Hubble constant. *Nature volume 551 pages 85–88*, 2017.
- [46] Tim Dietrich et al. Matter imprints in waveform models for neutron star binaries: Tidal and self-spin effects. *Phys. Rev. D 99, 024029*, 2019.
- [47] J. Weber. Gravitational-Wave-Detector Events. *Phys. Rev. Lett. 20(1307)*., 1968.
- [48] J. Weber. Gravitational Radiation. *Phys. Rev. Lett. 18, 498*, 1967.
- [49] Taylor J. H. and Weisberg J. M. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. *The Astrophysical Journal, 253:908-920*, 1982.
- [50] G Hobbs et al. The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Class. Quantum Grav. 27 084013*, 2010.
- [51] Siegman. Lasers. *University Science Books*, 1990.
- [52] Tobin Thomas Fricke. Homodyne detection for laser-interferometric gravitational wave detectors. *PhD thesis for Louisiana State University and Agricultural and Mechanical College*, 2011.
- [53] Eric D Black. An Introduction to Pound-Drever Hall Laser Frequency Stabilization. *Am. J. Physics*, 69, 2001.
- [54] LIGO Collaboration. https://labcit.ligo.caltech.edu/ll_news/s5_news/s5article.html. 2007.
- [55] Abbott et al. Directional limits on persistent gravitational waves using LIGO S5 science data. *Phys. Rev. Lett. 107, 271102*, 2011.
- [56] S Hild et al. DC-readout of a signal-recycled gravitational wave detector. *Class. Quantum Grav. 26 055012*, 2009.
- [57] J Degallaix et al. Commissioning of the tuned DC readout at GEO 600. *J. Phys.: Conf. Ser. 228 012013*, 2010.
- [58] Robert L. Ward. Length Sensing and Control of a Prototype Advanced Interferometric Gravitational Wave Detector. *PhD thesis, California Institute of Technology, Pasadena, CA*, 2010.
-

- [59] Brian J. Meers. Recycling in laser-interferometric gravitational-wave detectors. *Phys. Rev. D* 38, 2317, 1988.
- [60] The LIGO Scientific Collaboration et al. Advanced LIGO. *Class. Quantum Grav.* 32 074001, 2008.
- [61] Peter Fritschel, Matthew Evans, and Valery Frolov. Balanced homodyne read-out for quantum limited gravitational wave detectors. *Optics Express* 4225, Vol. 22, No. 4, 2014.
- [62] LIGO DCC. Block Diagram for Squeezer 3.125Mhz Combiner for O4. 2022.
- [63] LIGO Collaboration. Instrument Science White Paper 2021. *LIGO DCC*.
- [64] Hans-A Bachor and Timothy C Ralph. A Guide to Experiments in Quantum Optics. *Wiley-VCH*, 2009.
- [65] Lisa Barsotti. LIGO Document T1800447-v7 Design Requirement Document of the A+ filter cavity (HSTS) and relay optics (HDS, HXDS) for frequency dependent squeezing. *LIGO DCC*, 2022.
- [66] R. Weiss. Quarterly Progress Report: Electromagnetically coupled broadband gravitational antenna. *LIGO Document P720002-v1*, 1972.
- [67] Rana Adhikari. Sensitivity and Noise Analysis of 4 km Laser Interferometric Gravitational Wave Antennae. *PhD thesis, Massachusetts Institute of Technology*, 2006.
- [68] The LIGO Scientific Collaboration. Sensitivity and performance of the Advanced LIGO detectors in the third observing run . *APS, Phys. Rev. D* 102, 062003, 2020.
- [69] H. Grote et al. High power and ultra-low-noise photodetector for squeezed-light enhanced gravitational wave detectors. *Optics Express*, Vol. 24, No. 18, 2016.
- [70] John Horgan. Is the Gravitational-Wave Claim True? And Was It Worth the Cost? *Scientific American*, 2016.
- [71] Abbott et al. Population Properties of Compact Objects from the Second LIGO–Virgo Gravitational-Wave Transient Catalog. *The Astrophysical Journal Letters*, Volume 913, Number 1, 2021.
- [72] R. Abbott. In-vacuum DCPD for OMC (O3 and before). *LIGO DCC D060572-v1*, 2009.
- [73] G. Heinzel. LISO, Program for Linear Simulation and Optimization of analog electronic circuits. *MPQ Garching*, 1998.

-
- [74] W. Korth. OMC in-vacuum DCPD preamp LISO model. *LIGO DCC T1500142-v1*, 2015.
- [75] Paul Horowitz and Winfield Hill. The Art of Electronics. 1980.
- [76] Gianfranco Durin and Stefano Zapperi. The Barkhausen effect. *arXiv:cond-mat/0404512 [cond-mat.mtrl-sci]*, 2006.
- [77] Col. Wm. T. Mclyman. Transformer and Inductor Design Handbook. *JPL CalTech*, 1988.
- [78] Streetman Ben. Solid state electronic devices. *Pearson Education*, 2015.
- [79] Vincent John David. Fundamentals of infrared and visible detector operation and testing . 2016.
- [80] R. Abbott. Reduction of Bias Voltage Noise in OMC DCPDs. *LIGO DCC - E1600061-v1*, 2016.
- [81] Aldo Ejlli. Case drawing for A+ DC-PD readout. *LIGO doc: D1900453-v2*, 2019.
- [82] Aldo Ejlli. A+, ISC, OMC GW DCPD and SQZ, 105 kHz QPD enclosure. *LIGO doc: D2000246-v9*, 2022.
- [83] Aldo Ejlli. Researcher at Cardiff University. *EjlliA@cardiff.ac.uk*.
- [84] LIGO collaboration. L1 x axis magnetic noise. https://ldas-jobs.ligo-la.caltech.edu/detchar/summary/day/20200511/pem/magnetic_x/, 2020.
- [85] LIGO Collaboration. LIGO, VIRGO and KAGRA observing run plans. <https://observing.docs.ligo.org/plan/>, 2021.
- [86] V. Frolov. LLO.
- [87] LIGO collaboration. Input referred photocurrent noise during O3. https://alog.ligo-la.caltech.edu/aLOG/uploads/50726_20200103123344_FSR_CLF.pdf.
- [88] GWINC. <https://git.ligo.org/gwinc/pygwinc>.
- [89] H. Harari. A New Quark Model for Hadrons. *Phys. Lett. B Volume 57, Issue 3, 7 July*, 1975.
- [90] The CDF Collaboration. Evidence for Top Quark Production in pp Collisions at $J_s = 1.8$ TeV. *Phys.Rev.Lett.* 73, 1994.
- [91] The CDF Collaboration. Observation of Top Quark Production in p^-p Collisions. *Phys.Rev.Lett.* 74, 1995.
-

- [92] P. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.* *13*, 508, 1964.
- [93] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.* *13*, 321, 1964.
- [94] T. W. Kibble. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.* *13*, 585, 1964.
- [95] The CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B Volume 716, Issue 1*, 2012.
- [96] The ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B Volume 716, Issue 1*, 2012.
- [97] B. Schutz. A First Course in General Relativity. *Cambridge University Press*, 2009.
- [98] C. Rovelli. Quantum Gravity. *Cambridge University Press*, 2004.
- [99] P. J. E. Peebles and Bharat Ratra. The cosmological constant and dark energy. *Rev. Mod. Phys.* *75*, 559, 2003.
- [100] Lord Kelvin. Lecture on the Velocity of Stars. *C.J. Clay and Sons. p. 274*, 1884.
- [101] J. C. Kapteyn. First Attempt at a Theory of the Arrangement and Motion of the Sidereal System. *Astrophysical Journal, vol. 55*, 1922.
- [102] F. Zwicky. The red shift of extragalactic nebulae. *Helvetica Physica Acta, Vol. 6*, 1933.
- [103] Richard Massey et al. The dark matter of gravitational lensing. *Rep. Prog. Phys.* *73 086901*, 2010.
- [104] E. Komatsu et al. Five-year wilkinson microwave anisotropy probe (WMAP) observations: Cosmological interpretation. *arXiv:0803.0547 [astro-ph]*, 2008.
- [105] Planck Collaboration. Planck 2018 results - V. CMB power spectra and likelihoods. *AA Volume 641*, 2020.
- [106] Simone Aiola et al. The Atacama Cosmology Telescope: DR4 maps and cosmological parameters. *JCAP12(2020)047*, 2020.
- [107] Silvia Gallia, Fabio Iococ, and Alessandro Melchiorri Gianfranco Bertone. CMB constraints on Dark Matter models with large annihilation cross-section. *Phys. Rev. D 80, 023505*, 2009.

-
- [108] Gianfranco Bertone and Tim M. P. Tait. A new era in the search for dark matter. *Phys. Rev. D* *100*, 2018.
- [109] Hui L., Ostriker J. P., Tremaine S., and E Witten. Ultralight scalars as cosmological dark matter. *Phys. Rev. D* *95*, 2017.
- [110] Sébastien Clessea and Juan García-Bellido. Detecting the gravitational wave background from primordial black hole dark matter. *Phys. Dark Universe* *18*, 105–114, 2017.
- [111] Leszek Roszkowski. WIMP dark matter candidates and searches—current status and future prospects. *Rep. Prog. Phys.* *81* 066201, 2018.
- [112] Sander M. Vermeulen. Direct limits for scalar field dark matter from a gravitational-wave detector. *Nature*, 2021.
- [113] Abhishek Banerjee, Hyungjin Kim, and Gilad Perez. Coherent relaxion dark matter. *Phys. Rev. D* *100*, 2019.
- [114] A. Ringwald. Axions and Axion-Like Particles. *arXiv:1407.0546 [hep-ph]*, 2014.
- [115] Mordehai Milgrom. On the use of galaxy rotation curves to test the modified dynamics. *The Astrophysical Journal*, *333:689-693*, 1987.
- [116] Constantinos Skordis and Tom Złóśnik. New Relativistic Theory for Modified Newtonian Dynamics. *Phy. Rev. Lett.* *127*, 161302, 2021.
- [117] Katia Moskvitch. Troubled Times for Alternatives to Einstein’s Theory of Gravity. *QuantaMagazine*, 2018.
- [118] Bigi I. I. and Sanda A. I. CP Violation. *Cambridge University Press*, 2009.
- [119] S. M. Barr. Solving the Strong CP Problem without the Peccei-Quinn Symmetry. *Phys. Rev. Lett.* *53*, 329, 1984.
- [120] R. D. Peccei and Helen R. Quinn. CP Conservation in the Presence of Pseudoparticles. *Phys. Rev. Lett.* *38*, 1440, 1977.
- [121] R. D. Peccei and Helen R. Quinn). Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D* *16*, 1791, 1977.
- [122] R.D. Peccei. The Strong CP Problem and Axions. *Lect. Notes Phys.* *741:3-17,2008*, 2008.
- [123] Leanne D. Duffy and Karl van Bibber. Axions as Dark Matter Particles. *New J.Phys.* *11:105008,2009*, 2009.
-

- [124] Paola Arias, Joerg Jaeckel, Javier Redondo, , and Andreas Ringwald. Optimizing light-shining-through-a-wall experiments for axion and other weakly interacting slim particle searches. *Phys. Rev. D* *82*, 2010.
- [125] I.G. Irastorza et al. Conceptual design of the International Axion Observatory (IAXO). *Journal of Instrumentation Vol 9*, 2014.
- [126] CAST Collaboration. New CAST limit on the axion–photon interaction. *Nature Physics vol. 13*, 2017.
- [127] K. Barth, D. Delikaris, G. Passardi, M. Pezzetti, O. Pirotte, L. Stewart, B. Vullierme, L. Walckiers, and K. Zioutas. Commissioning and First Operation of the Cryogenics for the CERN Axion Solar Telescope (CAST). *AIP Conference Proceedings 710, 168*, 2004.
- [128] B. Majorovits and for the MADMAX interest group. MADMAX: A new road to axion dark matter detection. *J. Phys.: Conf. Ser. 1342 012098*, 2020.
- [129] A. Wagner et al. Search for Hidden Sector Photons with the ADMX Detector. *Phys. Rev. Lett. 105, 171801*, 2010.
- [130] Jonathan L. Ouellet et al. First Results from ABRACADABRA-10 cm: A Search for Sub- μeV Axion Dark Matter. *Phys. Rev. Lett. 122, 121802*, 2019.
- [131] D.B. Tanner and N.S. Sullivan. The “Gen 2” Axion Dark Matter Experiment (ADMX). *Department of Physics, University of Florida Final report*, 2019.
- [132] Andreas Ringwald. Exploring the role of axions and other WISPs in the dark universe. *Physics of the Dark Universe*, vol. 1, 116-135, 2012.
- [133] OSQAR Collaboration. New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall. *Phy. Rev. D: Particles and fields 92(9)*, 2015.
- [134] ALPS collaboration. Design of the ALPS II optical system. *Physics of the Dark Universe Volume 35*, 2022.
- [135] Dieter Horns et al. Hardening of TeV gamma spectrum of active galactic nuclei in galaxy clusters by conversions of photons into axionlike particles. *Phys. Rev. D 86*, 2009.
- [136] Kim J.E. Weak-interaction singlet and strong CP invariance. *Phys. Rev. Lett.*, *43*, 1979.
- [137] Shifman M., Vainshtein A., and Zakharov V. Can confinement ensure natural CP invariance of strong interactions? *Nuclear Phys. B, 166 (3)*, 1980.

-
- [138] Dine M., Fischler W., and Srednicki M. A simple solution to the strong CP problem with a harmless axion. *Nuclear Phys. B*, 166 (3), 1981.
- [139] Zhitnitsky A. On Possible Suppression of the Axion Hadron Interactions. (In Russian). *Sov. J. Nucl. Phys.*, 31, 1980.
- [140] Cheng S.L., Geng C.Q., and Ni W.-T. Axion-photon couplings in invisible axion models. *Phys. Rev. D*, 52, 1995.
- [141] Meyer M., Horns D., and Raue M. First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations. *Phys. Rev. D*, 87, 2013.
- [142] Giannotti M., Redondo J. Irastorza I., and Ringwald A. Cool WISPs for stellar cooling excesses. *Cosmol. Astropart. Phys.*, 2016 (05), 2016.
- [143] Abbott L. and Sikivie P. A cosmological bound on the invisible axion. *Phys. Lett. B*, 120 (1), 1983.
- [144] Dine M. and Fischler W. The not so harmless axion. *Phys. Lett. B*, 120, 1983.
- [145] Preskill J., Wise M.B., and Wilczek F. Cosmology of the invisible axion. *Phys. Lett. B*, 120 (1), 1983.
- [146] Ipser J. and Sikivie P. Can galactic halos be made of axions? *Phys. Rev. Lett.*, 50, 1983.
- [147] A. Sokolov and A. Ringwald. Photophilic hadronic axion from heavy magnetic monopoles. *Journal of High Energy Physics* 123, 2021.
- [148] Ayman Hallal, Mauricio Diaz Ortiz Giuseppe Messineo, Joseph Gleason, Harold Hollis, D.B. Tanner, Guido Mueller, and Aaron Spector. The heterodyne sensing system for the ALPS II search for sub-eV weakly interacting particles. *Physics of the Dark Universe* 35, 2022.
- [149] R. Shah et al. TES Detector for ALPS II. *arXiv:2110.10654v1 [physics.ins-det]*, 2021.
- [150] Ayman Hallal, Giuseppe Messineo, Joseph Gleason, Harold Hollis, Mauricio Diaz, Ada Uminska, David Tanner, and Guido Mueller. REGH FDD. *ALPS internal document.*, 2020.
- [151] Hartmut Grote, Alasdair James, Jan Pödl, and Benno Willke. GEN FDD. *ALPS internal document*, 2021.
- [152] Clemens Albrecht. Straightening of superconducting HERA dipoles for the any-light-particle-search experiment ALPS II. *EPJ Techniques and Instrumentation volume 8, Article number: 5*, 2021.
-

- [153] D. Anderson. Alignment of Resonant Optical Cavities. *Applied Optics*, 23, 1984.
- [154] Daniel Brotherton. PhD student, University of Florida.
daniel.cai.brotherton@desy.de.
- [155] Andreas Freise. FINESSE. 2019.