



University
of Glasgow

Taylor, Alasdair (2014) *Aspects of optical meterology systems for space-borne gravitational wave detectors*. PhD thesis.

<http://theses.gla.ac.uk/5741/>

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

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

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

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

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

Aspects of Optical Metrology Systems for Space-Borne Gravitational Wave Detectors

Alasdair Taylor

School of Physics and Astronomy,
University of Glasgow

Presented as a thesis for the degree of Ph.D.
in the University of Glasgow, University Avenue,
Glasgow G12 8QQ

© Alasdair Taylor, 2013

November 2, 2014

Contents

Acknowledgements	xxiv
Acronyms	xxvi
Preface	xxix
Summary	1
1 Introduction to Gravitational Waves Research	1
1.1 Theory of gravity	1
1.2 Gravitational waves	2
1.2.1 Prediction of gravitational waves	2
1.2.2 Polarisation states	4
1.2.3 Production of gravitational waves	4
1.2.4 Indirect evidence of gravitational waves	5
2 Gravitational Wave Astronomy	7
2.1 A new window	7
2.2 Gravitational wave detection	8
2.2.1 Principles	8
2.2.2 Ground based detectors	8
2.2.3 Space based detection	9
2.3 Sources of gravitational waves	11
2.4 Testing G.R. in the strong field regime	14
2.4.1 Supermassive blackhole binaries	14
2.4.2 Extreme mass ratio inspiral	15
3 Reflectivity of a Hydroxide-Catalysed Silicate Bond	16

3.1	Introduction to bonding	16
3.1.1	Summary of bonding theory	18
3.1.2	Goal of this experiment	19
3.2	Experimental setup	21
3.2.1	Optical layout	25
3.2.2	Sample alignment	27
3.2.3	Probe beam wavelengths	28
3.3	Initial experimental results	29
3.3.1	Single disk testing	29
3.3.2	Angular tilt caused by rail	30
3.3.3	Measured bond reflectivity	31
3.4	Optical modeling of the bond layer	33
3.4.1	Model concept	33
3.4.2	Mathematical description	34
3.4.3	Model variables and outputs	36
3.4.4	Model results	38
3.5	Final results	42
3.6	Conclusion	44
4	Optical Modeling of Second Generation Fibre Injectors	45
4.1	Introduction	45
4.2	Design concepts	48
4.2.1	eLISA FIOS requirements	48
4.2.2	Design options	48
4.3	FIOS components	53
4.3.1	FMA and fibre guard	53
4.3.2	Silica spacer	53
4.3.3	Spacer to lens interface	54
4.3.4	Lens options	55
4.3.5	FIOS focal length	57
4.4	Optical performance	58
4.4.1	Modelled beam parameters	58
4.4.2	Experimental results	60
4.5	Conclusion	63

5	Optical Testing of Short Focal Length Components	64
5.1	Introduction	64
5.2	Method	66
5.2.1	Measurement principles	66
5.2.2	Methods of absolute measurement considered	67
5.2.3	Relative measurements	71
5.3	Experimental method and setup	72
5.3.1	Optical setup	72
5.3.2	Lens placement and orientation	73
5.3.3	Wavefront sensor alignment	74
5.4	Results	75
5.4.1	Measured focal positions	75
5.4.2	Expected uncertainty	76
5.5	Conclusion	77
6	Performance testing of photodiodes	78
6.1	Introduction	78
6.1.1	OSI InGaAs photodiodes	80
6.1.2	Lewicki/Silicon Sensor photodiodes	81
6.2	Photon detection by InGaAs and Silicon	83
6.2.1	Silicon photodiodes in the visible spectrum	83
6.2.2	InGaAs photodiodes in near infrared spectrum	83
6.2.3	Silicon photodiodes in near infrared spectrum	85
6.2.4	Enhancing Silicon responsivity to near infrared light	85
6.3	Photodiode characterisations	87
6.3.1	Spatial uniformity testing	87
6.3.2	Frequency response testing	88
6.4	Spatial intensity response	89
6.4.1	Introduction	89
6.4.2	Birmingham University scans	92
6.4.3	Glasgow University scans	94
6.4.4	Results	97
6.4.5	Conclusion	101
6.5	Frequency response	104

6.5.1	Introduction	104
6.5.2	Low frequency transfer functions	105
6.5.3	High frequency transfer functions	108
6.5.4	Conclusions	113
6.6	Pulse response times of photodiodes	114
6.6.1	Introduction	114
6.6.2	Theory	115
6.6.3	Experimental setup	115
6.6.4	Prompt response	117
6.6.5	Slow response	119
6.6.6	Conclusion	120
6.7	Summary	121
7	Induced Longitudinal Path Length Noise	122
7.1	Introduction	122
7.2	Model of QP22-E characteristics	124
7.2.1	Intensity response of QP22-E	124
7.2.2	Phase lag of QP22-E	125
7.2.3	Simulation results	126
7.3	Experimental investigation	129
7.4	Experimental results	130
7.5	Conclusion	132
8	Beam Dump Optical Design and Testing	133
8.1	Introduction	133
8.1.1	Stray light from unwanted beams	134
8.2	Light suppression goal	136
8.2.1	Interferometer requirement	136
8.2.2	Stray light in a LISA-type optical bench	139
8.3	Possible beam dump designs	141
8.3.1	Commercial beam dump design	141
8.3.2	Conic cutout design	142
8.3.3	Light pipe	144
8.3.4	Spirals	145

8.4	Geometric analysis of spiral designs	146
8.4.1	Method	146
8.4.2	Initial designs	146
8.5	Theoretical performance	151
8.5.1	Method	151
8.5.2	Model results	154
8.5.3	Simulation summary	156
8.6	Experimental testing of a prototype beam dump	159
8.6.1	Method	159
8.6.2	Optical setup	159
8.6.3	RF cross-talk	161
8.6.4	Results	163
8.7	Analysis	164
8.8	Conclusion	165
9	Outlook	166
	Bibliography	167

List of Figures

1.1	Illustration of the effect of a passing gravitational wave on a ring of test particles. The direction of propagation is through the page along the ‘ z ’ axis. The upper series of rings show the effect of ‘+’ polarisation, while the lower rings show ‘ \times ’ polarisation.	3
2.1	<i>Illustration of the range in frequency of gravitational wave sources. Credit: NASA [26]</i>	11
3.1	<i>Illustration of the measurement concept. Incident beam with amplitude Ψ_i produces three main beams from the front, middle, and back of the sample. These beams have amplitudes Ψ_a, Ψ_b, and Ψ_c. Note that beam Ψ_b is the combination of multiple beams from the bond layer.</i>	21

- 3.2 *Diagram of the optical layout showing the key optics. The quadrant photodiode (QPD) monitors the alignment of the overlapped beams, local oscillator and selected probe beam. The measurement PD is a single element photodiode. Both Lasers are shown, as are the acoustic-optic modulators (AOM), which modulated the frequency of both the reference and probe beams, to produce the desired heterodyne frequency. Also shown is the motorised mounting of the sample and a number of aperture stops which were needed to block unwanted beams from both the AOMs and the sample. The alignment mirrors were used to optimise the overlap from both lasers after they had been through the AOMs. The interference contrast was optimised at the final beam splitter (BS-2). The three beams from the sample, in both green and red, can be seen. 24*
- 3.3 *This graphic shows how the probe beams, in red, are selected through the motion of the sample. The reference beam, in blue, interferes with a selected probe beam reflection to form the purple beam. The purple beams are detected by the photodiodes. Due to the size of the photodiode and the prospect of stray beams there is a need for apertures. The aperture before the beam combiner, blocks all but a selected probe beam reflection from the sample. By moving the sample the reflected beam passing through the aperture is changed. Even though the beam combiner had an anti-reflection coating, a reflected local beam combining with a stray beam from the sample could be stronger than the bond layer signal. 26*

3.4	<i>This graphic shows the readout signal chain. The photodiodes were Thorlabs PDA36A-EC, single element photodiodes with a built in amplifier where the gain was adjustable. These were connected to a second ac coupled amplifier and finally the output of this amplifier was fed into the National Instruments (N.I.) cDAQ-9178, containing a NI 9215 module, which is a 16-bit adc system.</i>	27
3.5	<i>Diagram of the single fused silica disk and the beams produced through reflections at the air (n_1) to silica (n_2) boundaries. The input light Ψ_i was 633 nm of S-polarisation.</i>	29
3.6	<i>This collection of figures illustrates the ideas that the model is based on. The top left figure shows the way in which beams can be produced from multiple internal reflections inside the bond layer. The right figure shows the layout of main beams of the model: beams A, B, and C; being the front, middle and back beams. The boundaries are also labeled in terms of the reflection and transmission, e.g. R_{34} is the reflected beam amplitude between the boundaries of refractive index n_3 and n_4. The small circular image is the place where the phase of the beams is combined to calculate the interference.</i>	34
3.7	<i>This graph illustrates the different modeled amplitude ratios of the reflected beams for a fixed refractive index change of 0.005 from silica to bond at each wavelength. The natural look of this graph on a linear y-axis scale is a sin wave; however it is more convenient to have the y-axis in dB when drawing comparisons between the different parameters.</i>	38

3.8	<i>This graph illustrates how the amplitude changes with refractive index for a fixed thickness of bond. The values of refractive index at each wavelength are -0.005 to -0.5 of the value of fused silica at that wavelength. On a linear scale these plots are a parabolic shape centred on the refractive index of silica at each wavelength. Due to this, there is a strong correlation between amplitude size and the bond refractive index change for all wavelengths.</i>	39
3.9	<i>Amplitude of reflected beams from the bond layer compared to beam ‘A’. The left graph is 532 nm while the right is 633 nm. The X-axis in both is the refractive index of the bond layer, with a range of values which are -0.005 to -0.5 of the value of fused silica at that wavelength. The Y-axis is the bond layer thickness.</i>	40
3.10	<i>This graphic represents the probability of any pair of parameters shown, given measured values of bond layer reflectivity at both 633 nm and 532 nm. The measured values were given a Gaussian uncertainty distribution, the sigma of which come from the spread of the measured results, and was then compared to similar results at a different wavelength. The refractive index is with respect to fused silica. The sum of the probabilities was normalised to one, for all the possible results in this parameter space.</i>	43

4.1	<i>A picture of four LPF FIOS. Each FIOS shown has: a fibre holding block, a mounted lens, and a polarising cube, all bonded to a common baseplate. The fibre emerging from the back of the glass block is protected by a metal guard and a rubber guard (not shown). The end of the fibre is flush to the surface of the block facing the lens. The air gap between this face and the lens allows the beam to expand before being collimated by the lens. The polarising cube then cleans the polarisation of the output beam.</i>	46
4.2	<i>A selection of Zemax models showing the various designs considered. Design ‘A’, is similar to the LPF FIOS in design. Design ‘B’ is all fused silica, including the lens and cover slip. Design ‘C’ uses a fused silica spacer between fibre end and lens. Each design is modelled from the virtual waist of a fibre end (not shown).</i>	49
4.3	<i>Illustration of thermally driven effects on the LPF FIOS and on Design ‘A’. The blue centre line of the design, coaxial with the fibre and the lens at construction temperature. After a drop in temperature, an exaggerated effect on the dimensions is shown by the red outlines. The change in the fibre end height is denoted ‘δ_c’. The shorter physical distance ‘γ_f’ leads to a defocused fibre end and so less collimated output. The change in lens position ‘γ_c’ is not equal to ‘δ_c’, which results in the output vector being angled with respect to the centre line and is due to the CTE of the lens being different to that of fused silica.</i>	50

4.4	<i>Presented is a selection of collimated beam profiles from nominally identical fused silica aspheres. Each beam originated from a cleaved fibre, which had been shown to give acceptable results with other lens including the LPF FIOS lens.</i>	51
4.5	<i>Illustration of thermally driven effects on design ‘C’. The blue centre line of the design, coaxial with the fibre and with the lens at construction temperature. After a drop in temperature, an exaggerated effect on the dimensions is shown by the red outlines. The change in the fibre end height is denoted ‘δ_c’. The shorter physical distance ‘β_f’ leads to a defocused fibre end and so less collimated output. The change in lens position is ‘β_c’ is equal to ‘δ_c’, so the output beam vector is parallel to the centre line, resulting in no angular deviation.</i>	52
4.6	<i>Illustration of the quasi-monolithic FIOS design. Labeled are all the key parts, including the new item, the fused silica spacer.</i>	53
4.7	<i>Finite element analysis of the FIOS design by M. Perreur-Lloyd.</i>	57
4.8	<i>Graphic from Zemax showing the profile of the intensity profile with the phasefront of the beam at the lens output face. Marked onto the graph are two lines denoting $1/e^2$ or one beam radii of 0.79 mm. The scale of the phasefront plot is 10^{th} of a wave per division, while the x-axis is in millimeters.</i>	59
4.9	<i>Picture of the first fully bonded eLISA FIOS</i>	61
4.10	<i>Measured beam profiles of the first two eLISA FIOS. Shown are the measured vaules (\times), the best fit Gaussian beam profiles ($—$) and their ω_o (\otimes).</i>	62

5.1	<i>This illustration shows the basic variables of the measurement system. The separation of the fibre end and the lens rear surface ‘S’, is the key variable, as the thickness ‘T’ of each of the lenses was already known.</i>	67
5.2	<i>This illustration shows the variables and the reference points of the chosen measurement system. Changes of the separation between fibre end and lens moves the output beam waist position, which was measured by a wavefront sensor at a fixed distance from the fibre end.</i>	72
5.3	<i>Graphic to illustrate the details of the optical setup. The fibre was held by two clamps, which pushed it into a groove on the metal plate. The end of the fibre protruded a distance of 25 mm to allow it to be held close to the lens. The lens holder is a U-groove, and had a ball bearing glued at one side. The wavefront camera was atop a lab-jack which was adjusted in order to centre the beam on the wavefront camera.</i>	74
6.1	<i>Picture of the Silicon Sensor QP22-E photodiode.</i>	82

6.2	<p><i>This diagram shows the some of the specialised features of the QP22-E type photodiodes. Illustrated are the variable absorption depths for different wavelengths of light. Though not to scale it is clear to see that short wavelength light (blue) is absorbed in a shallow depth of silicon near to the p-n junction and inside the depletion region. Shown is the connector array on the backside of the photodiode. These are thought to be the cause of the array of spots seen on all of the 2D scans of the photodiodes. Note also the back scattered light. This is done to raise the Amps/Watt rating of the photodiode by reflecting unabsorbed light back through the device.</i></p> <p><i>note: this does not represent the actual QP22-E photodiode configuration and is for illustrative purposes only</i></p>	84
6.3	<p><i>This is the circuit diagram of the readout chain for each quadrant of the photodiodes tested for spatial intensity response. In this example, the photodiode has no reverse bias applied; however later scans did employ reverse bias. The front end electronics of this setup were made to be identical to the flight electronics of the LPF mission, and the signals are feed into a N.I. 9215 16-bit adc module, in a cDAQ-9178.</i></p>	89

6.4	<i>Optical layout used to produce 1D scans at the Univeristy of Glasgow. The light was delivered by a fibre coupler and then aligned onto the middle on the photodiode through the beam splitter. Some of input light, reflected from the beam splitter, was directed to a power meter. Light passing through the beam splitter was focused down onto the surface of the photodiode. The light reflected from the surface of the photodiode, passed back through the lens; was re-collimated, and partially reflected to the wavefront sensor. The radius of curvature of the light reaching the wavefront sensor was therefore a measure of the spacing between the lens and the photodiode surface. For a collimated beam from the fibre coupler the correct spacing of lens-to-photodiode would yield a focused spot on the photodiode surface and a well collimated beam at the wavefront sensor. . . .</i>	90
6.5	<i>Shown is the University of Birmingham scans of a QP22-E photodiode showing large scale non-uniformity of the response. . .</i>	92
6.6	<i>The graph shown here is the plots of both the X and Y axis of the beam used to probe the photodiode. A straight line fitting of these points gave an ω_o of $19.5 \pm 2 \mu\text{m}$.</i>	94
6.7	<i>The figure shows the region over which a typical 1-D scan of a QP22-E photodiode was conducted. The scan runs parallel to the vertical gap of the photodiode, but at a distance from it to ensure there is no significant signal from quadrants 1 and 3. The diagram is to scale: the black lines are the thickness of the gap, the quadrants are the correct area, and the thickness of the red line in the orange strip, is the width of the beam used to probe the photodiode.</i>	95

6.8	<i>Graph of repeatability of the measurements. Taken by rescanning the same path of a QP22-E photodiode with a 5–10 minute break between scans. Each repeated data point, shows variation at the $\sim 3\%$ level which is about the same as the error in the measurements due to the correction for the power variation. The scan starts each time in the centre of the gap between quadrants one and two. And shows an arbitrary, repeatable shape to the responsivity of the photodiode. The irregularities of the photodiode, are much larger than the $\pm 3\%$ uniformity of response requirement which is needed over the active area of each quadrant.</i>	96
6.9	<i>Graph showing the spacial response of a OSI InGaAs photodiode. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the sharp transition between the quadrants. The data shows variation at the $\sim 3\%$ level which is about the same as the error in the measurements due to the correction for the power variation.</i>	98
6.10	<i>Graph showing the spacial response of a Lewicki QP22-E Photodiode without reverse bias. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.</i>	99
6.11	<i>Graph showing the spacial response of a Lewicki QP22-E Photodiode with a reverse bias of 9 V. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.</i>	100

6.12	<i>Graph showing the spacial response of a Lewicki QP22-E Photodiode with a reverse bias of 100 V. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.</i>	101
6.13	<i>Graph produced by Astrium GmbH Friedrichshafen [52], measuring the spatial phase response of the QP22-E photodiodes. The measurements were done using 1064 nm light on an unbiased photodiode.</i>	102
6.14	<i>The circuit diagram of the readout chain for a single photodiode quadrant or photodiode, during the tested for frequency response. The non-recorded quadrants, were connected to ground through $47\ \Omega$. The minicircuits TB-409-1+ contains a GALI 1+, DC to 8 GHz amplifier. The output of which was fed into one of two Spectrum Analysers, depending on the experiment. .</i>	104
6.15	<i>The figure shows the optical setup used to test the QP22-E photodiodes over the frequency range of 1 – 100 kHz. The light from a laser was split into two AOMs driven at 80 MHz. The difference in heterodyne beat note was controlled by changing $\nu 2$.</i>	106
6.16	<i>The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at three different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the three different reverse bias voltages. .</i>	107

- 6.17 *The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at three different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the three different reverse bias voltages. 108*
- 6.18 *The figure shows the optical setup used to test the QP22-E photodiodes over the frequency range of ~ 1 kHz to ~ 100 MHz. The light from the lasers was passed through faraday isolators before being directed into a scanning cavity and a high bandwidth InGaAs photodiode. This photodiode was used to detect the beatnote between the two lasers. The frequency difference between the lasers was set to $\nu 2$, and was maintained via a servo which used feed back from the photodiode to drive the slave laser to follow the the master at the correct offset. The heterodyne beatnote at the test photodiodes was $\nu 2$ minus $\nu 1$ where $\nu 1$ is 80 MHz. The heterodyne frequency was changed by adjustment of $\nu 2$. The reference photodiode was again smaller in area than the QP22-E photodiode, so all results are again relative and not absolute in there magnitude. The measured phase difference was unaffected by this. 109*
- 6.19 *The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at four different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the four different reverse bias voltages. 111*

6.20	<i>The graph shows the phase lag of the QP22-E compared to an InGaAs photodiode. The QP22-E was measured at four different reverse bias voltages. The y-axis phase of the readouts from the two photodiodes.</i>	112
6.21	<i>Light step generation setup. The RF switch inputs were: ω_1 an 80 MHz signal amplitude modulated at 1 kHz, and ω_2 a square wave modulation signal. The output of which drove an AOM which produced a beam of light, amplitude modulated at 1 kHz, and periodically switching on and off with frequency ω_2.</i>	116
6.22	<i>Shown is the readout chain for the photodiodes used in the pulse response experiment. The amplifier was rated as DC to 8 GHz while the oscilloscope had a 350 MHz bandwidth.</i>	116
6.23	<i>Shown is the fast and slow response of both an InGaAs photodiode and a QP22-E photodiode. The InGaAs response is limited by the time taken for the acoustic wave to travel across the beam in the crystal. The prompt response of both is evident, with the QP22-E photodiode showing a slow rise over a $10\mu\text{s}$.</i>	118
7.1	<i>Response of a single quadrant, modeled in Matlab. The shape of the response is based on 1-D scans performed at the University of Glasgow. The graph is cropped to show a 500×500 section over which the photodiode would ideally have a flat response.</i>	125
7.2	<i>Phase response of a single quadrant, modeled in Matlab. The shape of the phase response is based on the Astrium scans [52].</i>	126
7.3	<i>Results showing a coupling between, beam motion across the simulated photodiode and the resulting apparent longitudinal signal.</i>	127

7.4	<i>Picture of the experimental setup. Shown in the picture is an initial test, involving a InGaAs photodiode. The LPF FIOS was mounted on a hexapod and used to deliver a geometrically stable beam of AM light which was centred on the photodiode. The photodiode was attach to a linear stage, which was actuated by a piezo, to produce the desired relative motion of beam to photodiode.</i>	130
7.5	<i>Experimental results showing a coupling between, beam motion across the photodiode and the resulting apparent longitudinal signal. The gradient of the best fit line, is equal to that of the simulation.</i>	131
8.1	<i>Shown is the area where the unwanted beam is produced. Labeled is the local oscillator light in red, and the received beam from the distant spacecraft in green. The coatings of the component ‘N’ are also labeled ‘C₁’ and ‘C₂’. Component ‘M’ does not affect the stray light. The light which is scattered or escapes the beam dump can couple back into the science interferometer along the optical path of the received beam.</i>	139
8.2	<i>Diagram of a commercial beam dump design. The light entering the beam dump is scattered off the conical point in the centre. The thick black line is the black powder coating which is a strongly absorptive material. The design makes it so that after 3 reflections, light can exit the beam dump along the same beam path it came in on. A typical commercial beam dump of this design has light amplitude suppression of 10⁻⁵ or less of the input light exiting the beam dump.</i>	141

8.3	<i>Sectioned views of the conic cutouts. Shown for each design is a single ray entering the beam dump in the top half of the beam (red-pink area), before exiting the beam dump in the lower half of the beam, after multiple reflections in the conic cut out. The number of reflections for a ray traveling parallel to the axis of the conic shape is determined by the angle ‘θ’. The ray in the left beam dump design (purple ray) has four reflections, while in the right beam dump design (blue ray), has eight.</i>	142
8.4	<i>Shown are three different beam dump designs. Designs ‘A’ and ‘B’ are both light traps, where light is funneled into a cylinder to cycle. Design ‘C’ is a finite spiral, with an extra large opening diameter. The diagonal lines are artifacts of the Zemax output.</i>	147
8.5	<i>Design ‘D’ also called the ‘final design’. Shown is the location of the several objects: the origin of the beam, and two of the three detector planes (mentioned in more detail in the following section). Also shown are two exiting rays, each having made 12 internal reflections. Although it looks like ~ 10 reflections, there are reflections (on both rays) which cannot be seen in this projection.</i>	150
8.6	<i>Input beam profile at the origin labeled in figure 8.5. The red circle is the number of rays incident on the detector plane (10^7), the Green arrow is the ‘peak irradiance’ and the units of the colour bar (W/cm^2), and the pink arrow is the total power incident on the detector (1 W).</i>	152
8.7	<i>Plot showing the statistical spread of rays exiting from the beam dump after 11 – 13 reflections. The area of this detector is 80 mm square and it is positioned 30 mm from the beam dump exit. Shown on the data for the plot is the number of rays which hit the detector, these number 3.2×10^4 of the 10^7 that entered.</i>	153

8.8	<i>Plot of the beam dump designs, showing the variation in performance. Each design shows a different spread of rays exiting after a set number of reflections. Note that it is \log_{10} of the number of rays exiting that is plotted. The number of input rays was 10^5.</i>	157
8.9	<i>A graphical render of the final beam dump design. Shown are the two halves of the beam dump, which allow the spiral to be machined out and the absorbing black coating applied before assembly. A small footprint on the optical bench area is achieved by the spirals, which make use of the available height above the optical bench. CAD model and render produced by Michael Perreur-Lloyd.</i>	158
8.10	<i>The optical setup for the beam dump testing. Not shown are the fibre AOMs which fed the fibre couplers or a waveplate and polariser pairing which cleaned the output polarisation of each coupler. The red beam is the local oscillator beam 'ω_1', the blue beam is the probe beam 'ω_2', and the purple beam is the result of their combination into a heterodyne beat note. The unused port of each beam splitter 'BS-1' & 'BS-2', was dumped into a commercial beam dump. The optic behind the test object was a mirror, used to produce a reference heterodyne signal with which to compare the effectiveness of the beam dump.</i>	160
8.11	<i>Shown is the RF cross-talk between the probe beam signal and the local oscillator signal.</i>	161

List of Tables

3.1	Single non-bonded disk test results	30
3.2	Angular deviation of a reflected beam for a given rail position	31
3.3	Amplitude reduction ratio between ‘ Ψ_1 ’ and ‘ Ψ_2 ’ for each wavelength of light and the coefficient of reflection for the bond layer.	32
4.1	<i>Geltech 352220 lens specification</i>	56
4.2	<i>modelled beam profile from eLISA FIOS</i>	58
4.3	<i>Measured beam profile of a collimated output beam from two nominally identical Geltech lenses C & D, as compared to the model predictions where the waist is 0 m or 1 m from the lens.</i>	60
5.1	<i>Relative EFL of selected Geltech 352220 aspheres</i>	75
6.1	Specifications of InGaAs and QP22-E photodiodes	80
8.1	<i>Conic cutout beam dump results: angle, number of reflections, length and the minimum area of the beam dump foot print assuming a metal thickness of at less 5 mm.</i>	143
8.2	<i>Total light power escaping from the ‘Final design’, assuming a total input power of 1 W (each ray being 10^{-7} W), no scatter on reflection and a 98% absorptive coating.</i>	154
8.3	<i>Number of rays escaped from a design given 10^5 input</i>	155
8.4	<i>Amplitude suppression values for: prototype beam dumps, and flat sections with coating and finish consistent with the prototype beam dumps.</i>	163

Acknowledgements

Where do you start when a list is so long as to include almost countless names, that one is definitely assured to forget some of them? I feel unbelievably fortunate to have stumbled into such an exciting project, which has taken me around the world and shaped all aspects of my personality for the better.

To Harry: your ability to snooze through half a tele-o-con only to arise suddenly and deliver a telling question, is unlike anything I have ever seen. This thesis would be a mere figment of my imagination were it not for the time and effort you have invested and for this and so much more I am truly thankful.

To Dave: your kindness and patience are seemingly endless. Your methods of teaching have always been relaxing and informative. Even if it felt like most of the time was spent with you waiting for me to catch up.

To Christian and Ewan, I have to say a massive thank you. Your informal teachings and advice ensured that my experiments flowed smoothly whether they be experimental or computational. There was also the odd pint in the “Judges” after work and a Christmas night out which ended in the casino. The details of why, are still a mystery to me.

To Michael, thank you for all the technical assistance and prompt delivery of all those little pieces which are the really important bits of an experiment.

Also thank you, Russel and Steven for ensuring no work day was ever free of a little work place sporting banter.

To all those who spent many an evening unwinding after a long week with a few refreshing beverages, a big thank you.

I would like to acknowledge the support of STFC, who made this Ph.D a possibility.

To my friends and family, who have been there for me during all the ups and downs of life during this Ph.D, a huge thank you to you all.

And last but by no means least, I would like to thank the students and staff of Physics and Astronomy as a whole. Together you have made this Ph.D experience truly special and I can't think of anywhere I would rather have been than right here. Testament to this is the prestige that Glasgow University has in both appearance and in its research, the very reasons I wanted to come here in the first place; however it is the people that made it so special, thank you.

Acronyms

Acronym	Definition
AM	Amplitude Modulation
AOM	Acousto-Optic Modulator
BS	Beamsplitter
CAD	Computer Aided Design
CCD	Charge-Coupled Device
CMM	Coordinate Measurement Machine
DWS	Differential Wavefront Sencing
eLISA	Evolved Laser Interferometer Space Antenna
EMRI	Extreme Mass Ratio Inspiral
FIOS	Fibre Injector Optical Subassembly
LIGO	Laser Interferometer Gravitational wave Observatory
LISA	Laser Interferometer Space Antenna
LPF	LISA Pathfinder
MBH	Massive Blackhole
OB	Optical Bench
QPD	Quadrant Photodiode
RF	Radio Frequency
TM	Test Mass

Preface

This thesis is an account of the research carried out by the author in developing a range of necessary components for future gravitational wave detectors, between October 2009 and September 2013 in the Institute for Gravitational Research at the University of Glasgow.

The aims of this thesis were to design, develop and test technologies which enable future space-borne interferometric gravitational wave detectors. The research was as part of a large international collaboration focused on developments for a space-borne gravitational wave detector by the name of LISA and all of her variants. At the university of Glasgow this collaboration included H.Ward, D.I. Robertson, C. J. Killow, E. D. Fitzsimons, M. Perreux-Lloyd and the author.

In Chapter 1 the context of the research is introduced as a contribution to opening the field of gravitational wave astronomy through direct detection of gravitational waves taken from published literature.

In Chapter 2 the detection methods and sources of gravitational waves are introduced, with a particular focus towards important sources which could be detected and studied with a LISA-type mission. This work is derived from published literature.

In Chapter 3 a non-destructive method of testing the hydroxide-catalysis bond layer between two optics is described. The measurement of the reflection coefficients of the bond layer is used to calculate the bond layer refractive index and thickness. An optical model of the bond layer was written by the author in Matlab. The experimental work was carried out by the author with assistance from H. Ward, C. J. Killow and M. Kennedy.

In Chapter 4 the optical modeling of the fibre injectors for eLISA are described. The design of the eLISA fibre coupler was based on the previous LISA Pathfinder FIOS also constructed at Glasgow University. The design was the product of optical modeling carried out by the author and H. Ward, with assistance from D.I. Robertson, C. J. Killow, E. D. Fitzsimons and M. Perreur-Lloyd. The optical design analysis was carried out by the author.

In Chapter 5 an optical test of the focal length of the lens used in the eLISA FIOS. The focal length variation of a selection of lenses was measured. The testing method was developed by the author and the measurements were completed with the the assistance of E. D. Fitzsimons.

In Chapter 6 a series of photo-electric tests are performed to investigate the suitability of a Silicon quadrant photodiode with enhanced response at 1064 nm. The tests were to determine the photodiode's suitability for LISA Pathfinder and eLISA. The tests were an international collaboration between researchers at: Let labs Astrium GmbH, AEI Hanover, the University of Birmingham and the University of Glasgow. Low frequency testing was carried out at Glasgow University was by H.Ward, D.I. Robertson, C. J. Killow, E. D. Fitzsimons, F. Guzman and the author. Other testing carried out at Glasgow University was by H.Ward, D.I. Robertson, E. D. Fitzsimons and the author.

In Chapter 7 an investigation of a potential coupling between beam motion

across the surface of a photodiode and an apparent change of phase in the received signal. An optical model of the photodiode response was developed by the author in Matlab. Subsequent experimental measurements were carried out by the author with assistance from H. Ward.

In Chapter 8 a sequence of design and testing of a proto-type beam dump design for eLISA is described. The initial models were created by the author and M. Perreur-Lloyd, before the author carried out optical analysis with the optical modeling program Zemax. Proto-type beam dumps were procured by AEI Hannover, and tested by the author with assistance from H. Ward

Summary

Gravitational waves, predicted by General Relativity, are quadrupolar in nature and produce fluctuating strains in space-time. The strains cause measurable changes in the proper distance between test masses, detectable with interferometric techniques.

The predicted amplitude of gravitational waves is small and requires displacement sensitivity of 1 part in 10^{21} m/ $\sqrt{\text{Hz}}$. The field of gravitational wave astronomy is awaiting its inaugural detection; however there is strong indirect evidence to support their existence. One example of indirect evidence comes from tightly orbiting binary systems of two compact objects – one a pulsar – are detectable from Earth. Through pulsar timing, detailed analysis of the evolution of the binary orbit parameters can be acquired. So far all observed systems show orbit evolution which is in agreement with the prediction that the systems are losing orbital energy through the emission of gravitational waves.

Gravitational wave astronomy would open a new window on the universe and allow for complementary methods of detection to further our understanding of complex systems. The weakly interacting gravitational waves would pass undisturbed through matter which would block traditional EM observations. When thinking about a blackhole it is easy to understand the benefits of grav-

itational wave astronomy. These objects are only ‘observable’ through indirect methods using light, while using gravitational waves they are directly measurable.

Testing General Relativity in the strong gravitational environments of compact objects is a key goal for gravitational wave astronomy. Two of the sources predicted to generate gravitational waves which would yield information for the study of such strong gravitational fields are: massive blackhole (MBH) mergers, and extreme mass ratio inspirals (EMRI).

One of the ways to reduce thermally generated noise in gravitational wave detectors is to use glass to glass connections. These help reduce thermal noise in detectors for both ground and any planned space mission. Hydroxide-catalysis bonding is a chemical process which joins two conformal glass surfaces together. The resultant bond is of similar strength to the materials which are joined and has a similar chemical structure. The current method of testing the parameters of the bond layer are typically destructive. Described and demonstrated is a new method of bond parameter measurement which is non-destructive. This method uses two different wavelengths of light to measure the bond layer reflectivity and then through comparison to a model, the bond layer thickness and refractive index can be estimated.

For low frequency detectors such as a LISA-type mission, eLISA or LISA Pathfinder, the need for long term stability is essential. The use of very low expansion glass for the optical bench and the interferometer optics joined with hydroxide-catalysis bonding provides a stable base in which to build the interferometers. Another area where stability is required is in the beam delivery onto the optical bench. A beam delivery system called a FIOS – fibre injector optical subassembly – is designed which provides isolation from thermally driven beam pointing variation. The design also used hydroxide-catalysis

bonding to provide a fibre end to lens output optical path which is entirely in a glass medium. Their design and preliminary testing of a several proto-type FIOS are described. An investigation into the output of the first proto-type FIOS leads to an experimental measurement of the chosen lens focal length. The strict tolerance on the beam output parameters of the FIOS leads to a need to measure the focal length variation of the lens to $\pm 10 \mu\text{m}$.

The measurement technique used in a LISA-type mission is a heterodyne interferometer. A phase sensitive measurement is made using quadrant photodiodes and as such the fidelity of these photodiodes are critical. Following an observed phase lag between a silicon replacement photodiode and an InGaAs photodiode at a low heterodyne frequency of 1 kHz a large series of tests were conducted by an international collaboration. Reported is the analysis of the QP22-E Silicon Sensor quadrant photodiode response to both signal magnitude and phase over a range of frequencies of 1 kHz to 50 MHz.

The final topic covered in this thesis is an investigation into designing and testing a beam dump. The goal of this was to design a beam dump which could be space qualified and used on a LISA-type mission. The challenges were the limited space and the need for 10^{-6} suppression of the stray light amplitude. Using Zemax and SolidWorks several initial designs were analysed before a spiral design was chosen. Then following the procurement of a proto-type beam dump with non-representative surface roughness, an experimental investigation was performed to test their suppression. The results of which implied that the fully representative beam dumps would be suitable for use in eLISA.

Chapter 1

Introduction to Gravitational Waves Research

1.1 Theory of gravity

We experience gravity in many different forms. It keeps our feet on the Earth, the Moon in the sky, and the Sun shining. Visible on large scales in the universe is gravity's dominance over the others forces, shaping matter into many of the objects observable in the night sky. This force is commonly known to all of us; however due to its weak interaction with matter, it is a challenge to study the subtleties of its nature. Indeed we must rely on nature to provide suitable test samples, if we are to further our understanding of gravitation.

Albert Einstein's *General theory of Relativity* was published almost one hundred years ago, and in it he laid out one of the greatest modern theories about the universe [1]. It gave a mathematical framework for understanding gravity as a manifestation of curved spacetime, the curvature of which was dependent

on its matter and energy content. The essence of the theory was that, ‘*matter bends spacetime*’, and ‘*spacetime guides the motion of matter*’. The theory is based upon the description of spacetime with a metric. For weak gravitational fields – like the Earth’s gravity – these metrics show spacetime to be almost flat. Indeed these weak fields can generally be described by Newtonian dynamics with sufficient accuracy, or, if needed, Special Relativity. However, when looking closely at the orbital precession of Mercury, the positions of the stars behind the sun, or binary systems with solar-mass-sized compact objects, this description of gravity begins to break down. It is under these situations that General Relativity is needed.

1.2 Gravitational waves

1.2.1 Prediction of gravitational waves

A solution to Einstein’s equations, takes the form of equation (1.1). Where the energy momentum tensor ‘ \mathbf{T} ’, describing the matter energy content of the spacetime, is equated to the Einstein tensor ‘ \mathbf{G} ’, which describes the curvature of the spacetime.

$$\mathbf{T} = \frac{c^4}{8\pi G} \mathbf{G} \quad (1.1)$$

This is an important result as it predicts the stiffness, of the medium, in which gravitational waves propagate. The stiffness or ‘spring constant’ of spacetime is $c^4/8\pi G \sim 10^{42} \text{ kg m s}^{-2}$, implying that spacetime is a very stiff medium [2].

The deviation between curved spacetime and flat spacetime is measured by the strain ‘ h ’. Through suitable select of the form of ‘ h ’ equation (1.2) can be derived for free space propagation.

$$\left(-\frac{\partial^2}{\partial t^2} + c^2 \nabla^2\right) \bar{h}_{\mu\nu} = 0 \quad (1.2)$$

Where $h_{\mu\nu}$ represents the tensor components of the linearised perturbation of the gravitational field in Minkowski spacetime. This equation is a second important result as it has the mathematical form of the wave equation. This predicts oscillatory perturbations of spacetime, propagating at the speed of light. These perturbations are the predicted gravitational waves of General Relativity [2].

Because these waves propagate in the ‘stiff’ medium that is spacetime, all detectable sources of gravitational waves, for the foreseeable future, are astrophysical. This is due to the enormous levels of energy required to produce perturbations of a level detectable by modern detection methods.

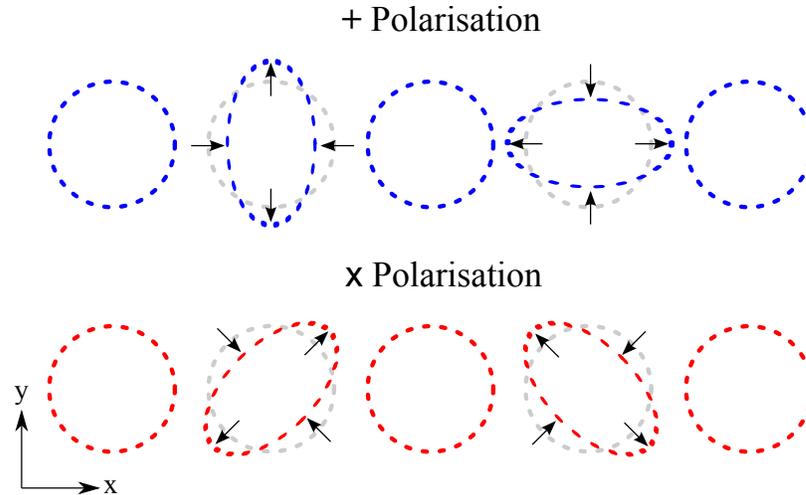


Figure 1.1: Illustration of the effect of a passing gravitational wave on a ring of test particles. The direction of propagation is through the page along the ‘z’ axis. The upper series of rings show the effect of ‘+’ polarisation, while the lower rings show ‘x’ polarisation.

1.2.2 Polarisation states

Gravitational waves come in two polarisation states, shown in figure 1.1, and a general case can be described by a linear combination of ‘+’ and ‘×’ polarisation states. The effect of a passing gravitational wave is to squeeze and stretch in orthogonal directions, perpendicular to the direction of propagation, where the angle of these orthogonal directions depends on the polarisation. The change in positions of the test particles in figure 1.1 is detectable as a change in their proper distance ‘ L ’, the diameter of the ring of test particles in figure 1.1. The maximum proper distance change between a pair of free test particles, is described by equation 1.3. The displacement of a test particle, when there is a perturbation present is ΔL .

$$\frac{h}{2} = \frac{\Delta L}{L} \quad (1.3)$$

This is for a ‘best case’ scenario, where the direction of propagation is perpendicular to the detector, in this case a ring of test particles. Were the wave not to propagate along the ‘ z ’ axis in figure 1.1, then the measured displacement would be lower. A wave traveling along the ‘ y ’ axis in this figure would only displace the test particles in the x axis.

1.2.3 Production of gravitational waves

Any system in which there is an asymmetric time varying quadrupolar moment will be a source of gravitational waves. A typical example is a binary system in which two massive compact objects – neutron stars – are in orbit. The gravitational waves are a means to releasing gravitational energy from the system. An example of the magnitude of ‘ h ’ for a typical system can be

calculated by equation 1.4. Where ‘ R ’ is the orbital radius, ‘ f ’ the orbital frequency, and ‘ r ’ is the distance to the source, ‘ M ’ is the mass of both stars, ‘ c ’ and ‘ G ’ are the speed of light and gravitational constant respectively.

$$h = \frac{32\pi^2 G M R^2 f^2}{c^4 r} \quad (1.4)$$

Assuming an $M \sim 1.4 M_{\odot}$, $R = 20 \text{ km}$, $f = 1000 \text{ Hz}$, and $r = 15 \text{ Mpc}$ (the distance to the Virgo galaxy cluster), then strain h is $\sim 6 * 10^{-21}$. This strain demonstrates the scale of the problem when it comes to detection of gravitational waves.

1.2.4 Indirect evidence of gravitational waves

Although gravitational waves have not yet been directly detected, there is strong evidence to support their existence. One of the measurable effects that gravitational wave emission has on a system is to decay the orbits.

In 1974, Hulse and Taylor using the Arecibo radio antenna, discovered a radio pulsar – a neutron star – which showed a variation in the timing of the received radio pulses. The reason was that it was in a short period – 7.75 hours – orbit around another compact object, later determined to be another neutron star. This unique system provided the first measurements of the orbital decay through gravitational wave emission. The orbital decay has, since discovery, agreed with the predicted rate of decay due to emission of gravitational waves to 0.2% [3, 4]. The discovery and following results earned Hulse and Taylor the 1993 Nobel Prize in Physics. In 2004 Taylor and Weisberg analysed the data for the 30 years of observation and concluded that the error in the agreement was limited by the poorly known galactic constants, the Sun’s distance from the galaxy centre, the proper motion of the pulsar, and its distance from Earth.

In 2013 the existence of a massive $2M_{\odot}$ pulsar, in a compact binary system, with a white dwarf star was confirmed [5]. The systems high mass and short period of 2 hours and 27 minutes make it an ideal test subject for General Relativity, due to the curvature of the spacetime the binary experience.

Chapter 2

Gravitational Wave Astronomy

2.1 A new window

The field of astronomy has been dominated by a single observation method since its conception, the EM spectrum. This has revealed truly beautiful images of the universe; however given that the universe is 4.9% ‘normal’ matter, 26.8% ‘dark’ matter [6] with the rest being dark energy, the EM spectrum only shows us a small fraction of the universe. Gravitational waves are not part of the EM spectrum and are a complementary method of observation. Furthermore, dark matter does interact gravitationally, thus leading to a possibility of detection through gravitational waves. Due to the weak interaction with matter, gravitational waves pass through matter thus revealing previously unseen parts of the universe. Through gravitational waves astronomers could ‘see’ binary systems with a common envelope, a blackhole binary, and even past the surface of last scattering. Gravitational wave astronomy would provide a new window onto the universe.

2.2 Gravitational wave detection

2.2.1 Principles

Gravitational waves as mentioned and shown in figure 1.1, change the proper distance between ‘test masses’, one way to measure this is through laser metrology between these test masses. A method of detecting gravitational waves is through the use of interferometry. A laser beam is typically split into two beams which travel in two different directions, preferably perpendicular to each other. These beams are reflected back after a certain distance by the ‘test masses’. To ensure the sensitivity of the detector, the optics in the interferometer, specifically the ‘test masses’, are isolated from vibrational and other external disturbances [7].

2.2.2 Ground based detectors

There are several ground based gravitational wave detectors currently operating. These include LIGO [8], VIRGO [9], GEO600 [10], TAMA300 [11]. The largest of these systems is LIGO, comprised of three interferometric detectors, which are spread over two sites. At both sites there are 4 km arm length interferometers, the largest in the world, which are designed to detect changes in the separation of test masses of $\sim 10^{-18}$ m.

These detectors typically operate over a range of frequencies from ~ 50 Hz to 5000 Hz. Ground based detectors all suffer from the same low frequency noise limit (~ 10 Hz) from seismic and gravity gradient noise. This blocks attempts made from Earth to observe many of the predicted gravitational wave sources available below these frequencies [12]. This limit is why space based detectors

are planned to operate at lower frequencies, which would be undetectable from the ground.

These current generation ground based detectors had reached their design sensitivity and are undergoing upgrades. The upgraded detectors will be called advLIGO [13], advVIRGO [9] and LCGT [14]. GEO-HF [15] has transitioned to a higher frequency test bed for more novel interferometric techniques. A general overview of the gravitational wave detectors can be found in the above references and [16, 7, 17].

2.2.3 Space based detection

The Laser Interferometric Space Antenna (LISA) [18, 19, 20, 21] is the umbrella under which a global collaboration is working towards having a gravitational wave detector in space. The project has existed in many forms over its 30 year life and is in a mature form. Due to the unique challenges faced in moving precision interferometer to the space, LISA has a prototype mission to mitigate technical risk. The highly successful LISA Technology Package (LTP), the project which developed and built the required flight components for the prototype spacecraft named LISA PathFinder (LPF).

The LISA mission class has at its core, the measurement of free falling test masses through optical metrology over a much longer base line than would be possible on earth. Indeed what was called ‘the LISA mission’ had a triangular configuration of interferometer arms, each 5×10^9 m long. Large enough to encircle the Sun, the mission location was to follow 20° behind the Earth.

LPF, due for launch in 2015 and will mitigate most of the technical risk of a LISA-type mission. The spacecraft will be placed at the first Lagrangian

point (L1) and is a miniature version of a single LISA arm, containing two test masses, an optical bench, data readout electronics, test mass housing and control, and micro-newton thrusters. A more detailed description of the LPF mission and status can be found here [22, 23, 24].

Evolved LISA (eLISA), is a European only mission with reduced arm length of 1×10^9 m and only two arms. These and other changes were made to reduce the cost to an available budget, in order to be selectable as a candidate mission by European Space Agency (ESA). The potential of eLISA vs LISA is a slight reduction of the directional sensitivity and overall strain sensitivity; however all of the primary science goals are maintained. An overview of the science goals can be found here and the references there in [25].

2.3 Sources of gravitational waves

Over the years since gravitational waves were predicted a wealth of potential sources have been studied. These sources can be classified into four main types: transient events, binary systems, non-axisymmetric rotating objects, and a stochastic background. The characteristic frequencies and amplitudes of these sources cover an entire spectrum, with a general rule being that the more massive objects tend to have the lower frequencies, while the smaller more objects have higher frequencies and more rapidly evolving signals.

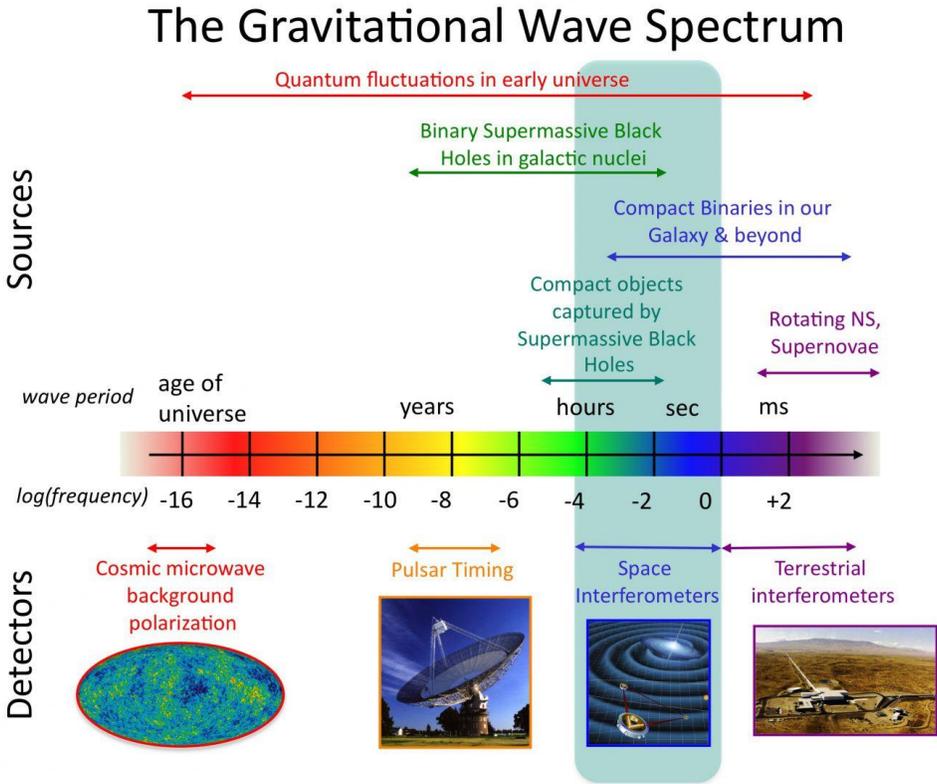


Figure 2.1: Illustration of the range in frequency of gravitational wave sources. Credit: NASA [26]

Transient or burst events are typically sources of gravitational waves over a

finite timescale. A typical source is expected to be a supernovae, where a star, unable to support its own weight, implodes to form a remnant core – a compact object – while the rest of the matter is ‘bounced off’ this core into the surrounding space. The ground based detectors have an appropriate bandwidth to detect these sources.

A system of two gravitationally bound compact objects with a combined mass of $\sim > 1 M_{\odot}$, are prime sources of gravitational waves for current and future ground based detectors. The binary system can be made of any combination of compact object, white dwarfs (WD), neutron stars (NS) and pulsars, or blackholes (BH).

Each type of binary system can emit gravitational waves during three separate phases: inspiral, merger, and ring-down. The inspiral stage is the longest and would be a continuous source of gravitational waves. The frequency of which would be steady. The merger is more like a burst or transient event, where the two objects merge together into a single object. During this time the gravitational wave output peaks, with strong emission over a chirp in frequency. The ring-down post merger is the phase where the single object releases gravitational energy in order to relax and dampen excited vibrational states. Depending on the mass of the system and their orbital period, the frequency of these sources fall over a broad range between mHz and kHz. This makes them both space and ground based targets.

The best example of non-axisymmetric source is a neutron star with a bump on the surface. Such a bump could be the result of deformation of the outer crust during the formation. It is known that pulsars, have glitches – changes or breaks in their continuous frequency of emission – which might be the result of structural changes on the neutron star outer crust. Current generation ground based detectors have set upper limits on the size of such bumps of the crab

nebula and vela nebula, pulsars [27, 28].

A stochastic background of gravitational waves from the big bang, loop quantum gravity or string theory are possibilities. Detection of a gravitational waves from the big bang would be direct experimental evidence and as such would be a very important result. These are targets for radio telescopes which use pulsar timing and cosmic microwave background detectors.

For further reading beyond the scope of this thesis see [29, 25, 30, 7] and the references within.

2.4 Testing G.R. in the strong field regime

Two important sources for eLISA which can test General Relativity in the strong gravitational field regime are binary systems where at least one of the objects is a $10^4 - 10^8 M_{\odot}$ blackhole. These sources are important due to the strong curvature of the spacetime around a Massive Blackhole (MBH). Another important fact is that the frequency of gravitational waves from such sources is typically < 1 Hz, making them ideal sources for space-borne gravitational wave detectors.

2.4.1 Supermassive blackhole binaries

One of the strongest sources of gravitational waves are binary systems where both objects are MBH. Here the systems are expected to form as part of galaxy merger and evolution in the early to late universe. The merger of two MBH is predicted to have a gravitational wave luminosity which is equivalent to $\sim 10^{23} L_{\odot}$, where L_{\odot} is the solar luminosity in the EM spectrum. This type of event would output more power - via gravitational waves - during its peak, than all the stars in the observable universe.

The detection of equal mass MBH binaries $10^4 M_{\odot}$ is expected to be visible out to redshift $z = 20$ with eLISA [25]. This detection range is sufficient to study galaxy formation and how MBH - like that at the centre of the Milky Way - are formed. The ‘intermediate mass’ blackholes ($\sim 10^4 M_{\odot}$) are the expected unobserved missing link between the stellar mass sized blackholes and the $\sim 10^8 M_{\odot}$ blackholes at the centre of galaxies [25].

2.4.2 Extreme mass ratio inspiral

Where the binary system is composed of a massive blackhole $> 10^6 M_{\odot}$ and a compact object $\sim 10 M_{\odot}$ this is called an Extreme Mass Ratio Inspiral (EMRI). EMRIs yield finely resolved details of the nature of the gravitational potential around MBH and are expected to be a very clean astrophysical source of gravitational waves. A typical system could have 10^5 cycles of the MBH that should reveal the a precision map of the spacetime matrix in a way which is almost better than could be hoped for. EMRIs are expected to be the best sources to test the strong field regime predictions of G.R.

Over the eLISA timescales EMRI wavefronts are predicted to display extreme forms of periastron and orbital plane precession due the dragging of the interial frames by the MBH's spin [31]. The mass and spin of the central MBH should be determined to $10^{-3} - 10^{-4}$ fractional accuracy for an SNR of 20 [32]. The methods used to determine the gravitational field mapping of the MBH are similar to those used in the mapping of the Earth's gravitational field by satellites. Further information can be found in the following [25, 33, 34, 35] .

Chapter 3

Reflectivity of a Hydroxide-Catalysed Silicate Bond

3.1 Introduction to bonding

A commonly used method for the joining of optics in the field of gravitational wave astronomy is, hydroxide-catalysis bonding [36]. The process enables oxide producing substrates to be joined together with a thin bond, which is a chemical composite of the two adjoined substrates. This can be done by applying a solution, with hydroxide ions, between two flat surfaces. Typically, for a bond to be successful the surfaces of the two objects must be conformal due to a bond layer being only ~ 100 nm thick [37, 38].

The bonding process has a long heritage at the University of Glasgow, both in space applications and on the ground based detectors. The University of Glas-

gow developed monolithic glass suspensions for Advanced LIGO. The reason bonding is used here is the high mechanical strength and potential to lower thermal noise by avoiding glass to metal connections [39, 37, 40]. The lower thermal noise comes from the use of similar materials in both the mirrors, their holders and the very thin bond layer. For space, the reasons bonding is used are primarily strength and stability, with the additional feature of adjustability during bonding. The strength ensures that it survives the stress of launch, while the stability ensures that the optics will remain aligned in the long term which is very important for a LISA-like mission where the signal time scales are 0.1 Hz to 0.1 mHz. This long term stability ensures the path length is stable on the optical bench.

A typical example of a bond carried out for the LISA pathfinder optical bench at Glasgow University, would be between a mirror-coated glass optic and an optical bench also made of a glass-like material. The optics's placement can be assured through the measurement of a beam, reflected off the optic during bonding [37]. As the bonding fluid initially acts as a lubricant between the two, the optics generally move freely upon placement. This almost frictionless motion stiffens over time until the two become optically contacted and then bond. Adjustability of components – for a short period after the application of the bonding fluid – helps ensure they are aligned correctly. During this time, an optic will typically move, as it floats on the bonding fluid. Precise control over the position of the optic is of great importance when constructing an optical bench using bonding. Once the bond has taken it is permanent. Removal of a bond is possible only for the first few minutes after the bond starts to set, the bonding process is not reversible without damaging the surfaces. This window of adjustability suffices to meet the tolerances on position and angle of beams which are present on many optical benches. The following papers, detailing LPF construction, describe precision bonding as was used, in the building of

both flight benches for LPF [37, 24, 41].

3.1.1 Summary of bonding theory

Hydroxide-catalysis bonding is a chemical process, where substrates with oxide layers can be joined through hydration/dehydration [38]. The bonding process can be used to bond silicon compounds, although it is not limited to them alone. With two silicon compounds with oxide layers; a strong layer of similar chemical composition to that of the oxide layer is formed; provided the surfaces to be joined are in close proximity. The sequence of chemical reactions in the bond layer formation are: hydration and chemical etching, followed by polymerisation, and dehydration. In the case described here, the material was fused silica.

The chemical processes for such a bond is as follows:

- The hydroxide ions etch the oxide layer at the surfaces, releasing silicate ions which mix with the hydroxide ions.



- The pH of the solution is lowered, due to the hydration of silicate ions by hydroxide ions; when the pH is < 11 , the silicate ions dissociate.



- The dissociated silicate ions polymerise to form siloxane chains by dehydration; this leads to water molecules being produced.



The bond is defined as a join between two surfaces with a silicate-like network. The siloxane chains that form the bond layer, are similar to silicon dioxide in structure, but are less dense due to the number of locally terminated groups e.g. Silanol groups (Si–OH).

The bond layer is thin enough that the silicate-like network can be considered to be uniform through the bond thickness; however it is possible that it is not uniform over the full area of the bonding surface.

The thin bond layer ~ 100 nm, means that the two surfaces must be globally flat in order to be in close proximity. This leads to a general requirement for the global surface flatness – of the surfaces to be bonded – of ~ 50 nm to provide a good bond over the whole surface. This is usually provided to optical shops as a $\frac{\lambda}{10}$ peak to valley ($\lambda = 633$ nm) specification on the surfaces to be bonded, typically resulting in a surface roughness, better than $\frac{\lambda}{30}$ RMS.

In order to guarantee a good bond, the surfaces should be clean at the molecular level, over the the macroscopic area of the bonding surfaces, with no particles larger than ten nanometers on the surface. Well filtered, high purity solutions in a clean environment is necessary to consistently achieve high quality bonds.

3.1.2 Goal of this experiment

Currently the optimisation of the processes to ensure a good bond and the properties of that bond are the focus of research. The thickness and refractive index of a bond, and their relation to the bonding parameters, are relatively unknown. Being able to measure these numbers in a non-destructive way would aid in the understanding of bonding techniques. This would also provide addi-

tional parameters for models, specifically ground based thermal noise research. The Young's modulus of a bond has been measured [40] and a refractive index could yield a better description of a bond layer. The ability to measure the bond reflectivity coefficients is, itself providing a measure of the quality of a bond. This experiment was designed to optically measure these important bond characteristics, and could be extended to determine their dependence on the bonding parameters such as: bonding fluid concentration/volume, substrate material, etc.

3.2 Experimental setup

A bond layer cannot be measured in isolation as it is a rigid join being two substrates. The concept of the method used to probe the bond layer can be seen in figure 3.1. In order to measure the absolute size of the reflected beam Ψ_b , comparisons were drawn to the expected reflected amplitudes Ψ_a and Ψ_c .

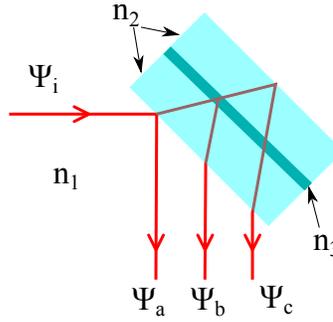


Figure 3.1: *Illustration of the measurement concept. Incident beam with amplitude Ψ_i produces three main beams from the front, middle, and back of the sample. These beams have amplitudes Ψ_a , Ψ_b , and Ψ_c . Note that beam Ψ_b is the combination of multiple beams from the bond layer.*

Both the bond's thickness and refractive index affect the amplitude of light reflected from a bond. In a pre-experiment using 1064 nm light, the reflected light from the bond layer was $\sim 10^{-4}\%$; therefore an amplitude sensitive measurement – a heterodyne interferometer – was chosen due to its greatly improved sensitivity. A power sensing detector, measures the intensity of a signal which is the square of its amplitude. A heterodyne interferometer senses the signal as proportional to the amplitudes of two interfering beams as follows:

$$I_{het} \propto (\Psi_{het})^2 = (\Psi_{LO} + \Psi_x)^2 \quad (3.4)$$

where I_{het} is the heterodyne intensity, Ψ_{het} is the heterodyne amplitude, Ψ_{LO} is the local oscillator beam amplitude, and Ψ_x is the reflected beam amplitudes

from the sample for $x = a, b, c$. It follows that for:

$$\Psi_x = E_x \cos(\omega_x t + \phi) \quad \& \quad \Psi_{LO} = E_{LO} \cos(\omega_{LO} t) \quad (3.5)$$

where ϕ is the phase between the two beams, ω is the frequency, t is time, and E is the amplitude of the electric field.

$$\Psi_x = [E_x \cos(\omega_x t + \phi) + E_{LO} \cos(\omega_{LO} t)]^2 \quad (3.6)$$

$$\begin{aligned} \Psi_x &= E_x^2 \cos^2(\omega_x t + \phi) + E_{LO}^2 \cos^2(\omega_{LO} t) \\ &\quad + 2E_x E_{LO} \cos(\omega_x t + \phi) \cos(\omega_{LO} t) \end{aligned} \quad (3.7)$$

Given that:

$$\cos^2(\theta) = \frac{1 + \cos(2\theta)}{2} \quad (3.8)$$

it follows that:

$$\begin{aligned} \Psi_x &= E_x^2 \left[\frac{1 + \cos(2\omega_x t + 2\phi)}{2} \right] + E_{LO}^2 \left[\frac{1 + \cos(2\omega_{LO} t)}{2} \right] \\ &\quad + 2E_x E_{LO} \cos(\omega_x t + \phi) \cos(\omega_{LO} t) \end{aligned} \quad (3.9)$$

and by dropping the terms without ω and the terms with 2ω and then expanded up using:

$$\cos(\theta) = \frac{e^{i\theta} + e^{-i\theta}}{2} \quad (3.10)$$

this gives:

$$\begin{aligned} \Psi_x &= \frac{E_x E_{LO}}{2} (e^{i[(\omega_x + \omega_{LO})t + \phi]} + e^{-i[(\omega_x + \omega_{LO})t + \phi]} \\ &\quad + e^{i[(\omega_x - \omega_{LO})t + \phi]} + e^{-i[(\omega_x - \omega_{LO})t + \phi]}). \end{aligned} \quad (3.11)$$

Using equation 3.10 again gives:

$$\Psi_x = E_x E_{LO} [\cos((\omega_x + \omega_{LO})t + \phi) + \cos((\omega_x - \omega_{LO})t + \phi)] \quad (3.12)$$

and finally, the equation for the lower frequency beat note is:

$$\Psi_x = E_x E_{LO} \cos((\omega_x - \omega_{LO})t + \phi). \quad (3.13)$$

As can be seen in equation 3.13, the signal received by the photodiode proportional to the amplitude of the signal beam (small), multiplied by the amplitude of the local beam (large). This is why a heterodyne is such a good method for detection of weak signals.

To ensure reliable comparison of the three amplitudes to be measured, each beam must be consistently and equally overlapped with the local oscillator. In a system where the beams are fixed, this is done once; however for this experiment the three reflected beams Ψ_a , Ψ_b & Ψ_c must each be interfered with the same reference beam. The most practical way of doing this was to place the sample on a rail, such that it could be moved, to select which of the three reflected beams was to be interfered with the reference (see figure 3.3). To ensure the beams were overlapped at all the relevant times, the sample must be polished and ideally have plane parallel sides. Without the polished sides the light is scattered and is unmeasurable; without the parallel sides the reflected beams are not parallel, which greatly increases the difficulty in maintaining the alignment of the interferometer. The bonded sample used during the testing was two fused silica etalon disks from ICOS (Imperial College Optical Systems). These disks were parallel sided and polished to better than $\frac{\lambda}{10}$. The bond between them was visually flawless and mature. This was important as bonds can take several months before becoming fully cured and stable.

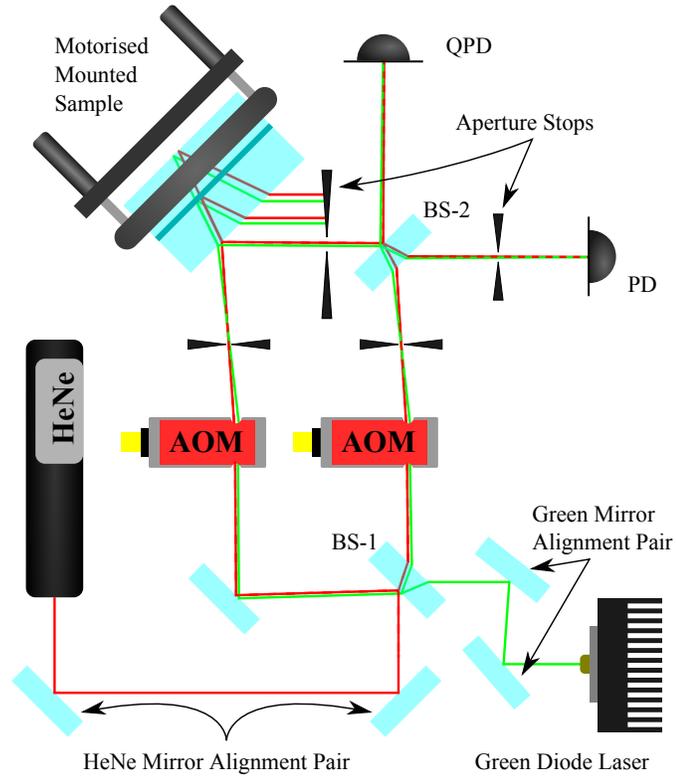


Figure 3.2: Diagram of the optical layout showing the key optics. The quadrant photodiode (QPD) monitors the alignment of the overlapped beams, local oscillator and selected probe beam. The measurement PD is a single element photodiode. Both Lasers are shown, as are the acoustic-optic modulators (AOM), which modulated the frequency of both the reference and probe beams, to produce the desired heterodyne frequency. Also shown is the motorised mounting of the sample and a number of aperture stops which were needed to block unwanted beams from both the AOMs and the sample. The alignment mirrors were used to optimise the overlap from both lasers after they had been through the AOMs. The interference contrast was optimised at the final beam splitter (BS-2). The three beams from the sample, in both green and red, can be seen.

3.2.1 Optical layout

Shown in figure 3.2 is the optical layout of the final experiment. The layout has two heterodyne interferometers, one at 633 nm and another at 532 nm. Both laser sources are ‘dog-legged’ into the beam splitter (BS-1) before the AOMs. This means that they reflect off at least two mirrors and are therefore fully adjustable in their beam vector without needing to move the source. This was needed to ensure that both beams could be adequately maneuvered independently. One key factor was ensuring that both beams passed through the same aperture stops, which required careful alignment of their entry to the AOMs. As an AOM works by refraction and because it is dependent on the wavelength of the light, both beams exit the AOM with a different angle. Using the following equation 3.14 where v is the speed of sound in the AOM crystal (typically 4000 ms^{-1}), f is the frequency of the sound waves, λ is the wavelength of the light, and Θ_n is the angle between the zero order, and the n^{th} order refraction.

$$\Theta_n = \frac{n\lambda f}{v} \quad (3.14)$$

The angle between them was 2 milli-radians, which meant that without corrections by the time they reached the sample they would be $\sim 800 \mu\text{m}$ apart. It was later considered that by using different frequencies of modulation, 95 MHz and 80 MHz, the output angles would have been equal. However that would have required precision alignment into both AOMs and was determined to be unnecessary. This was one of the largest factors which required independent adjustment. By angling and offsetting their entry position to the AOMs, both beams passed through the same aperture stops and were aligned to less than $100 \mu\text{m}$ on the photodiodes after the beam recombiner (BS-2).

In Figure 3.3 are shown the details of the beam selection method. The motion

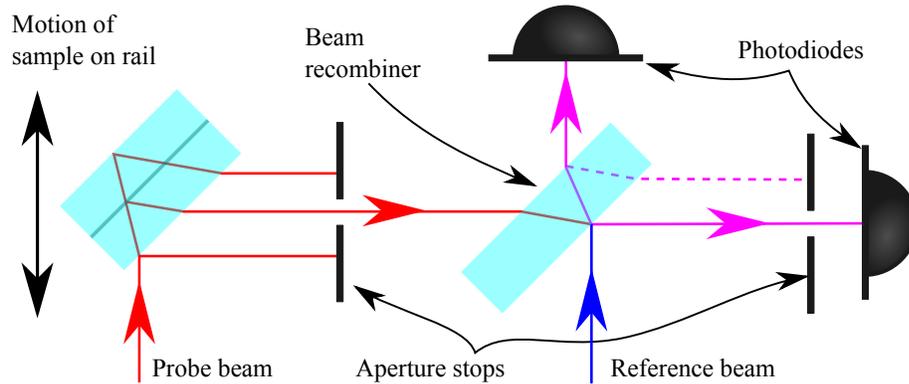


Figure 3.3: *This graphic shows how the probe beams, in red, are selected through the motion of the sample. The reference beam, in blue, interferes with a selected probe beam reflection to form the purple beam. The purple beams are detected by the photodiodes. Due to the size of the photodiode and the prospect of stray beams there is a need for apertures. The aperture before the beam combiner, blocks all but a selected probe beam reflection from the sample. By moving the sample the reflected beam passing through the aperture is changed. Even though the beam combiner had an anti-reflection coating, a reflected local beam combining with a stray beam from the sample could be stronger than the bond layer signal.*

provides the same effective measurements as that which would be achieved by scanning the probe beam across the sample to probe different parts of the sample, without it moving. This was considered more complicated when compared to a single rail. The quadrant photodiode detected the phasefront tilt of the interference between the probe and the reference beams. The phasefront tilt is caused by the misalignment of the two interfering beams. These measurements were used to direct a servo system, which was used to tilt the sample in real time. This was needed to correct a bow in the rail which caused the reflected beams to become misaligned. Other optical components used in the setup were a pair of polarisers.



Figure 3.4: This graphic shows the readout signal chain. The photodiodes were Thorlabs PDA36A-EC, single element photodiodes with a built in amplifier where the gain was adjustable. These were connected to a second ac coupled amplifier and finally the output of this amplifier was fed into the National Instruments (N.I.) cDAQ-9178, containing a NI 9215 module, which is a 16-bit adc system.

The cDAQ module (Figure 3.4) was controlled using LabView which also saved the data. Labview was also used to control the motorised mirror mount and rail. This mirror mount held the sample on the rail and so was able to correct any misalignment of the interferometer by actuation of the sample.

3.2.2 Sample alignment

It was important to maintain a constant angle of incidence at the sample. To ensure this, the displacement of the beam as the light passed through the sample, was used to check that the angle of incidence was correct. First a CCD camera was placed in the beam path behind where the sample would be located and the position of the beam recorded. The thickness of the sample was measured and the sample was then placed into the beam and adjusted in angle until the CCD camera measured the beam to have been displaced the correct amount. There was a slight difference in angle between the 532 nm and 633 nm due to the AOM. This angle was estimated to be less than 0.1° ;

however its effect was judged to be too small to be resolved by the experimental setup.

3.2.3 Probe beam wavelengths

An issue that was considered was the different reflection coefficients which the optics had for 633 nm and 532 nm. All of the optics were optimised for 633 nm and there was a noticeable difference when comparing the beam splitting coatings. As we were using ratios this was not a problem; however better optimised coatings, could have helped improve the signal to noise ratios of the 532 nm measurements. These coatings did inhibit the ability to use a 1064 nm laser in the same setup, as the mirror coatings only reflected $< 10\%$ of the incident 1064 nm light.

3.3 Initial experimental results

3.3.1 Single disk testing

The experimental setup was initially tested using a single fused silica disk. This disk was identical to the pair of disks bonded together in the test sample. This disk had no bond layer and was a test to see how accurately the experiment could measure the predictable amplitudes for the first three beams. These beams were labeled Ψ_1 , Ψ_2 and Ψ_3 in figure 3.5.

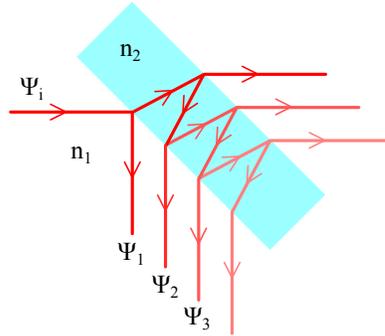


Figure 3.5: *Diagram of the single fused silica disk and the beams produced through reflections at the air (n_1) to silica (n_2) boundaries. The input light Ψ_i was 633 nm of S-polarisation.*

The goal of this experiment was to validate the accuracy and repeatability of the equipment. The initial results of the single disk test showed substantial variation and inaccuracy, see table 3.1. Run 3 in particular showed results which were clearly incorrect.

To investigate this a series of tests were performed to check that the beams were not being clipped or obstructed by dirt. The alignment of the component to the plane of the experiment was also checked. This was done for a single position of the rail. If the component were to be tilted, then the front

Table 3.1: Single non-bonded disk test results

Beams	Run 1		Run 2		Run 3	
	Amplitude	Ratio	Amplitude	Ratio	Amplitude	Ratio
Ψ_1	1.67	1	2.51	1	1.66	1
Ψ_2	1.44	0.86	2.14	0.85	2.21	1.33
Ψ_3	0.12	0.07	0.19	0.08	0.18	0.11

* Results taken by M. Kennedy

and back beams would differ in their height. The local beam height is fixed throughout and so this would lead to a change in the overlap of the beams, and a corresponding loss of signal. This would lead to erroneous measurements of the comparison between the reflected beams. Great care was taken to ensure that the beams were of equal height by measuring their heights after a propagation distance of 5 m which ensured the component was aligned to less than a milli-radian. With this alignment accuracy the potential for overlap change was $\sim 10 \mu\text{m}$ which would result in a negligible loss of contrast, given the beams were approximately 1 mm in diameter.

3.3.2 Angular tilt caused by rail

Having eliminated all other likely sources of error, the linearity of the motion of the rail was examined. The motion of the optic holder along the rail was found to introduce a repeatable angular change in the optic holder. This resulted in the component being tilted with respect to the incident beam, thus changing the beam vector of the reflected light and causing misalignment with the reference beam.

To measure this effect a quadrant photodiode was used to measure a beam's

position after being reflected from a mirror mounted on the rail. The photodiode was placed on top of a manual linear translation stage, which allowed the photodiode to be moved an equivalent distance for each rail position. Thus ideally keeping the reflected beam centred on the photodiode. The measurement was primarily focused on the changing beam height with rail position. Though the lateral position of the beam on the photodiode did show variation, it was less than 10% of the variation in height. The measured beam positions on the photodiode were then converted to angular deviations, the results of which are shown in table 3.2.

Table 3.2: Angular deviation of a reflected beam for a given rail position

Rail position (cm)	3.5	4	4.5	5	6	7	8	9	10	11
Deviation (milli-radians)	10	6	9	23	29	31	31	19	-14	-19

The magnitude of the variation between the probe beam at the different rail positions was ~ 50 milli-radians. This problem was solved using an actuated mirror mounting system which the author manually adjusted to maximise the heterodyne signal.

3.3.3 Measured bond reflectivity

The results of the bond reflectivity measurements are seen in table 3.3. The uncertainty in the values seen in table 3.3 are mainly due to the uncontrolled intensity of the 532 nm and 633 nm Lasers. Due to the nature of the experiment, the amplitudes of beams ‘ Ψ_1 ’ and ‘ Ψ_2 ’ were not measured simultaneously. Thus variation of the Laser power between these measurements was possible, resulting in an error in the ratio of these amplitudes. To mitigate this a series of five measurements at each wavelength was made. The uncertainty on the value was taken from the spread of these results.

Table 3.3: Amplitude reduction ratio between ‘ Ψ_1 ’ and ‘ Ψ_2 ’ for each wavelength of light and the coefficient of reflection for the bond layer.

Wavelength	Ratio between heterodyne amplitudes	Coefficient of reflection
532 nm	-53.5 ± 1 dB	0.37 ppm
633 nm	-52.5 ± 1 dB	0.48 ppm
1064 nm	-62 ± 3 dB	0.05 ppm

These results suggest that the bond layer is indeed similar in refractive index to that of the substrate material from which it is made. To further understand these results and evaluate the possible bond parameters which might produce such results, a model of the bond was produced.

3.4 Optical modeling of the bond layer

3.4.1 Model concept

Using the results presented in table 3.3 and some assumptions about the bond layer, a concept for a model was conceived. The bond layer was modeled as a thin plane of material between two large blocks of substrate. It was assumed that the refractive index change was a step function; not a gradient or smoothly changing refractive index, similar to a dielectric coating.

The siloxane molecules – which dominate the bond’s molecular structure – have refractive indexes which range from 1.4 – 1.6 at 600 nm [42]. Reflections from the bond in the sample – a pair of fused silica disks – were very weak in the visible and near infrared spectrum (see table 3.3). This suggest that the bond layer refractive index is not markedly different to that of the substrate material. This is perhaps not a surprising result given that the bond layer is made from the substrate material in its surroundings.

Following this it was taken as an assumption that the dispersion of the bond layer was similar to that of substrate material - fused silica - and the author didn’t find published literature which supported or disproved this assumption. Should it be proved that the dispersion of the bond layer, in the sample used for this experiment, is different over the wavelengths 633 nm and 532 nm, then this model will be incorrect in its analysis of the experimental results.

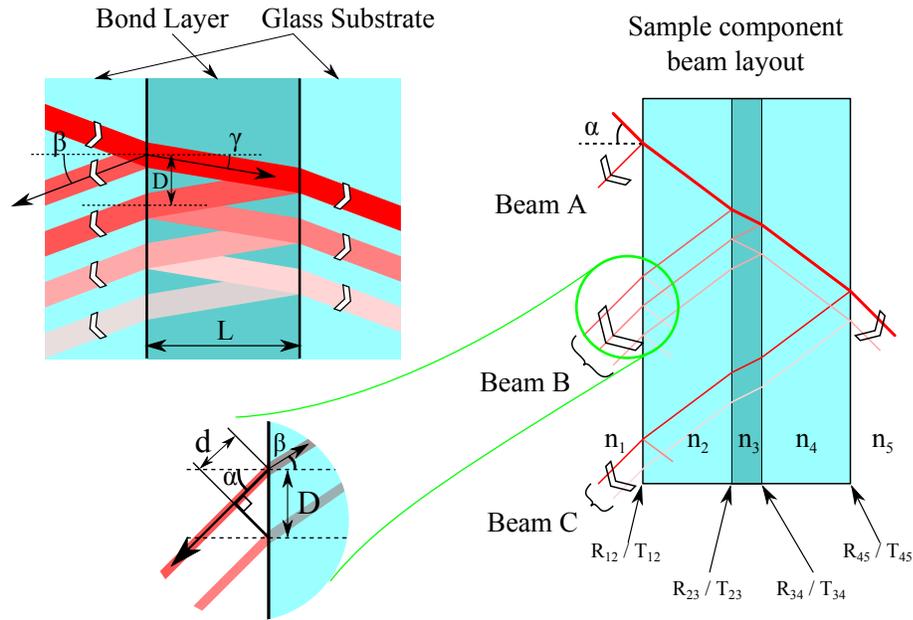


Figure 3.6: This collection of figures illustrates the ideas that the model is based on. The top left figure shows the way in which beams can be produced from multiple internal reflections inside the bond layer. The right figure shows the layout of main beams of the model: beams A, B, and C; being the front, middle and back beams. The boundaries are also labeled in terms of the reflection and transmission, e.g. R_{34} is the reflected beam amplitude between the boundaries of refractive index n_3 and n_4 . The small circular image is the place where the phase of the beams is combined to calculate the interference.

3.4.2 Mathematical description

The Fresnel equations 3.15 and 3.16 were used to calculate the reflection and transmission coefficients of the light amplitudes at each boundary. Due to the refractive index of the material being dependent on the wavelength, these were wavelength specific values. Here n_1 & n_2 are the refractive index of the air & front silica disc, θ_i & θ_t are the angle of incidence & transmission and R_{12} & T_{12} (see figure 3.6) are the reflection & transmission amplitude coefficients for

the interface at which beam A is reflected.

$$R_{12} = \frac{n_1 \cos \theta_i - n_2 \cos \theta_t}{n_1 \cos \theta_i + n_2 \cos \theta_t} \quad (3.15)$$

$$T_{12} = \frac{2n_1 \cos \theta_i}{n_1 \cos \theta_i + n_2 \cos \theta_t} \quad (3.16)$$

The first boundary seen on the right in figure 3.6 labeled R_{12}/T_{12} is the reflection/transmission coefficients of light power calculated from the amplitude coefficients. The subscripts indicate the refractive index boundaries with which the light interacts. e.g. R_{12} is the reflection coefficient of light passing from refractive index one into refractive index two. By the symmetry of the sample, the coefficients of reflection for R_{43} and R_{32} , are equal to R_{23} and R_{34} respectively. These were then used to calculate the amplitude of the various beams which were reflected towards the photodiode.

As can be seen in figure 3.6 the bond layer is modeled to have reflections at both interfaces, resulting in two main beams close together; with a phase difference dependent on the thickness and refractive index of the bond. The beams that make up beam 'B', both with amplitude coefficients of $\sim 3 \times 10^{-4}$ with respect to the input beam; shall be known as Ψ_{B1} and Ψ_{B2} . Subsequent higher order reflections – beams which totally internally reflect in the bond layer – shall be ignored due to their expected amplitude coefficients being $\sim 1 \times 10^{-7}$ or less. The phase difference between these beams Ψ_{B1} and Ψ_{B2} , is dependent on: the refractive index mismatch between the bond and disks, the angle of the incidence, and the bond thickness. As is shown in the left and the lower part of figure 3.6, one beam travels an extra distance ' d ' in air, while the other travels twice through the bond layer which has a thickness ' L ' at an angle ' γ '.

$$\Psi_{B1} \text{ path length} = d \cdot n_1$$

$$\Psi_{B2} \text{ path length} = \frac{2L}{\cos \gamma} \cdot n_3$$

The separation between the beams at the boundaries of, the bond layer to silica, and the silica to air, are equal and denoted ‘ D ’, see figure 3.6. The angle of incidence in n_1 is denoted ‘ α ’, in n_2 is ‘ β ’, and in n_3 is ‘ γ ’.

Given that:

$$D = 2L \tan \gamma \quad (3.17)$$

And:

$$d = D \sin \alpha \quad (3.18)$$

It follows that:

$$d = 2L \tan \gamma \sin \alpha \quad (3.19)$$

The phase ‘ ϕ ’ between the beams Ψ_{B1} and Ψ_{B2} is:

$$\phi = \left\{ d \cdot n_1 - \frac{2L}{\cos \gamma} \cdot n_3 \right\} \cdot \frac{2\pi}{\lambda} \quad (3.20)$$

Now that we know the phase difference between beams Ψ_{B1} and Ψ_{B2} , so beams ‘ A ’, ‘ B ’, and ‘ C ’ in figure 3.6, can be written as follows:

$$\text{Beam A} = R_{12} \quad (3.21)$$

$$\text{Beam B} = T_{12}R_{23}T_{21} + T_{12}T_{23}R_{34}T_{32}T_{21}e^{i\phi} \quad (3.22)$$

$$\text{Beam C} = T_{12}T_{23}T_{34}R_{45}T_{43}T_{32}T_{21} \quad (3.23)$$

Where beam ‘ B ’ is the sum of the amplitudes of beams Ψ_{B1} and Ψ_{B2} . As expected from equations 3.21 - 3.23 that only beam ‘ B ’ has a dependency on the bond thickness.

3.4.3 Model variables and outputs

The bond layer can be modelled as having any refractive index and the amplitude of the light reflected at the bond layer interfaces is due primarily to the

index change, and the angle of incidence. The phase change is dependent on: refractive index of the bond, the thickness, and the angle of incidence. The phase has a point of maximum constructive interference, which means that it has an upper limit. This upper limit means that for a given refractive index change, the largest reflected signal is limited by the phase of the two beams from the bond layer. A larger signal requires a larger step change in the refractive index between bond layer and substrate. This allows a lower limit to be set on the difference between the refractive index of the substrate and that of the bond layer.

The model of the bond layer took the following fixed input parameters: polarisation of the beams, refractive index of the two disks, wavelength of the light and angle of incidence onto the sample. The important variables for the model were the bond layer thickness, bond layer refractive index and wavelength of the beams. The angle of incidence was fixed at 45° , the refractive index of fused silica was known for each of the wavelengths of light used, and – post analysis of the model results – the polarisation of the light was settled as only S-Pol. The analysis showed that the front and back beams ‘A’ and ‘C’ were approaching the Brewster’s angle for P-pol light. Thus these beams were significantly weaker than the S-pol beams. The difference in bond reflectivity due to angle of incidence was $\sim 1 - 5\%$ per degree, which produced changes below the noise limit of the experiment in most cases. All the other parameters were assumed constant or were unchanged during the measurements.

Assumptions were made as to reasonable limits on the thickness of a bond layer and the change in refractive index from substrate material to the bond layer. Published data [43, 44, 45] suggests that the thickness of a bond between two disks similar to those under investigation is between the 60 – 100 nm and there is additional evidence that suggests bonds could be both thicker or thinner

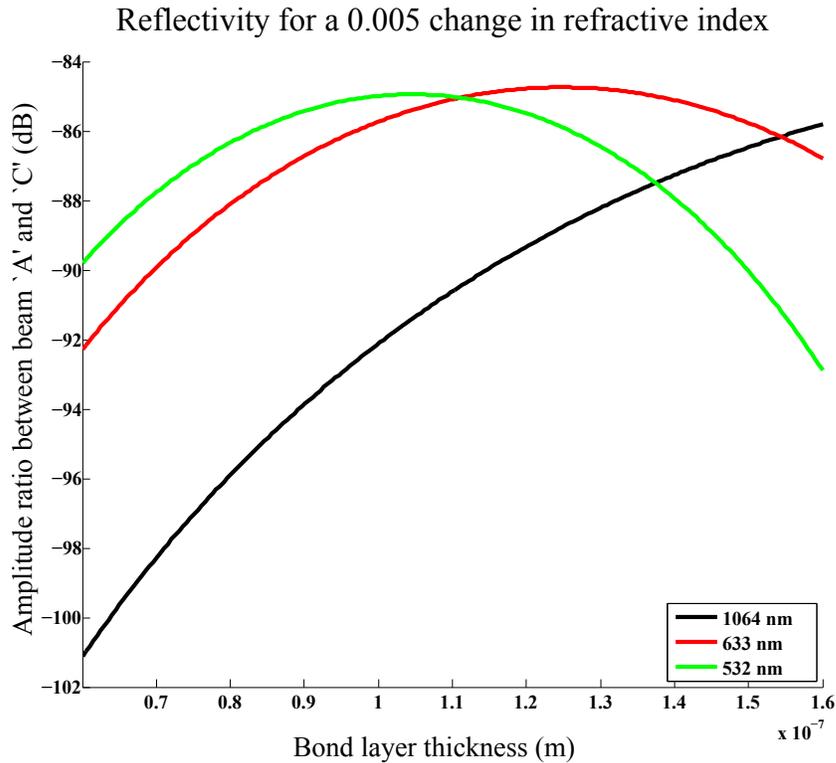


Figure 3.7: *This graph illustrates the different modeled amplitude ratios of the reflected beams for a fixed refractive index change of 0.005 from silica to bond at each wavelength. The natural look of this graph on a linear y-axis scale is a sin wave; however it is more convenient to have the y-axis in dB when drawing comparisons between the different parameters.*

than even that range. The initial range of thicknesses chosen for the model was 60 – 160 nm however this was expanded to include thicknesses down to zero, after results suggested the bond was possibly thinner than 60 nm.

3.4.4 Model results

Plotting equations 3.21 – 3.23 for fixed bond thickness and bond refractive index changes give figures 3.7 & 3.8. For the same bond layer, two different

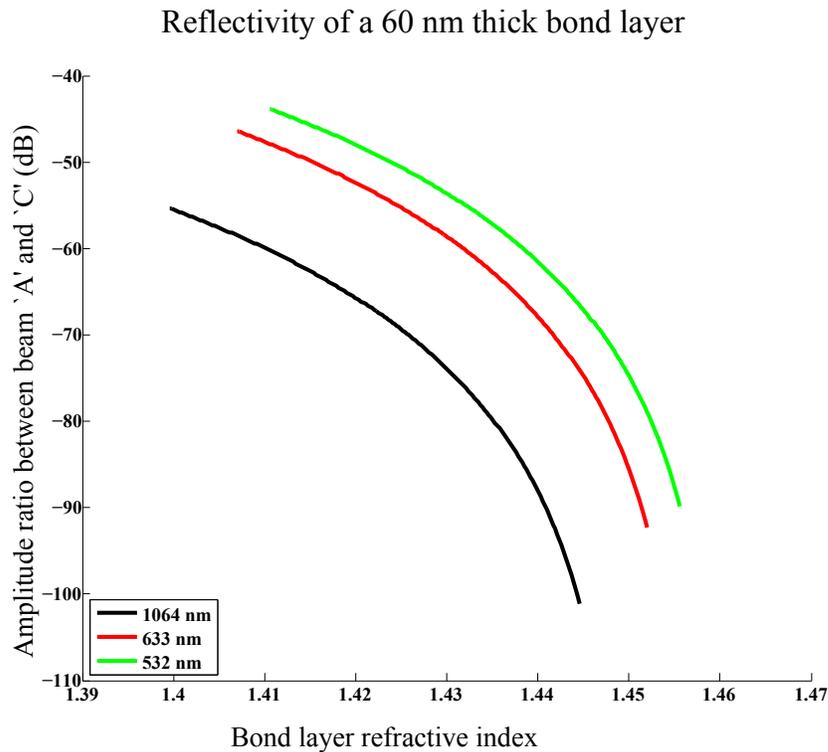


Figure 3.8: *This graph illustrates how the amplitude changes with refractive index for a fixed thickness of bond. The values of refractive index at each wavelength are -0.005 to -0.5 of the value of fused silica at that wavelength. On a linear scale these plots are a parabolic shape centred on the refractive index of silica at each wavelength. Due to this, there is a strong correlation between amplitude size and the bond refractive index change for all wavelengths.*

wavelengths, should have different phase changes. The modeling showed this to have a more noticeable effect on the reflected amplitude – see figures 3.7 & 3.8 – than any of the changes seen due to different polarisations or angles of incidence.

In figure 3.7 a selection of bond thicknesses are shown, with the results dependent only the phase of the signals due to the fixed refractive index change. In the case of constructive interference, the reflection amplitude is then limited

by the refractive index change.

Figure 3.8 would show a symmetry about the refractive index of the substrate material – fused silica – had the refractive index range of the bond layer been extended to include higher refractive indexes. The bond thickness was the same physical size; however its optical path length was dependent on the wavelength/refractive index. The amplitude was strongly dependent on the thickness of the bond; however was not independent of the refractive index, due to the optical path length’s dependency on the wavelength/refractive index.

Maps of the reflected bond layer beam amplitudes compared to beam A

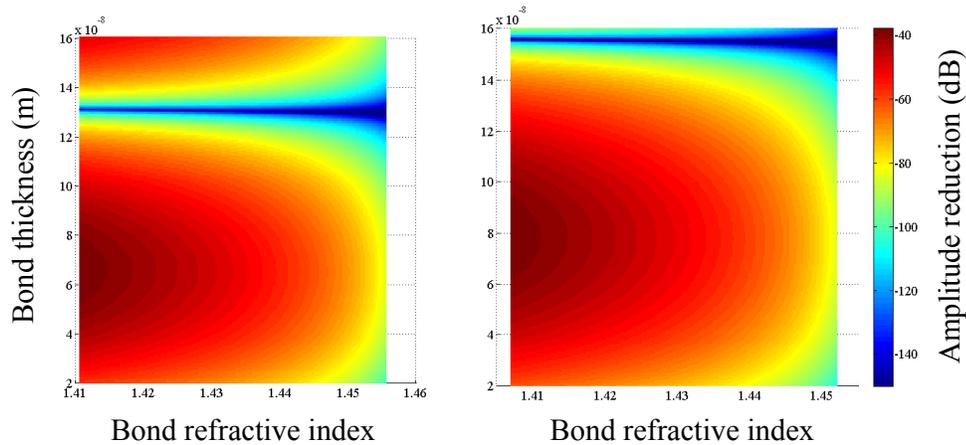


Figure 3.9: *Amplitude of reflected beams from the bond layer compared to beam ‘A’.* The left graph is 532 nm while the right is 633 nm. The X-axis in both is the refractive index of the bond layer, with a range of values which are -0.005 to -0.5 of the value of fused silica at that wavelength. The Y-axis is the bond layer thickness.

Figure 3.9 shows the amplitude of the bond layer beam for all the range of the modeled thicknesses and refractive indexes. These amplitudes of the bond layer

beams are scaled by comparing them to the amplitude of the beam reflected from the front surface of the component. Figures 3.7 and 3.8 are cross sections of the maps in y & x respectively.

3.5 Final results

The combination of the experimental results of the bond layer measurement with that of the model, allowed for a probability map of the likely parameters of the bond layer, for a given measurement position (see figure 3.10). The results of the bond measurements – plus their uncertainty values – were used to create a Gaussian probability map in both 633 nm and 532 nm. The amplitude reduction ratio provided a central value for the Gaussian, while the uncertainty was used to scale the Gaussian width parameter sigma. By combining the probability values of both wavelengths figure 3.10 was produced. The probability has been normalised to one, over the available parameters used.

The most likely parameters of the bond are a thickness of 120 ± 20 nm and a refractive index of -0.024 ± 0.002 relative to fused silica (1.457 at 633 nm). Indicated by the probability maps (figure 3.10), is that the bond layer refractive index is lower than the substrate material. There is no probable parameter set with refractive indexes, of the bond layer, larger than the substrate material. Some less likely results for thicker bonds do exist due to the periodic nature of the ‘ ϕ ’ term in equation 3.22.

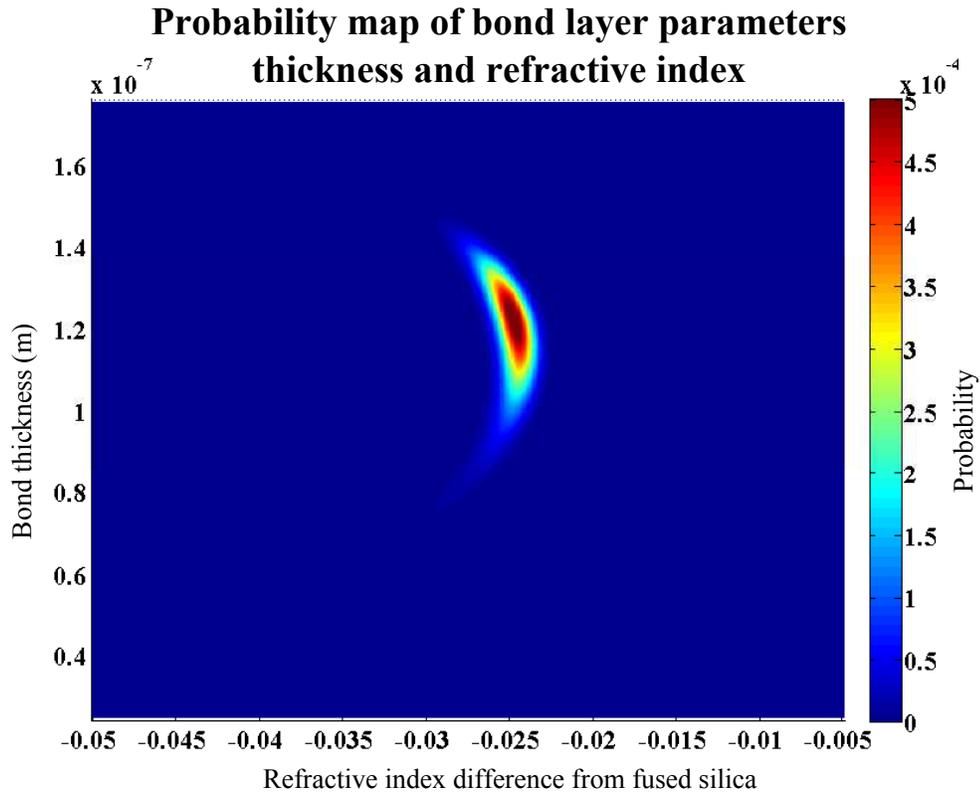


Figure 3.10: *This graphic represents the probability of any pair of parameters shown, given measured values of bond layer reflectivity at both 633 nm and 532 nm. The measured values were given a Gaussian uncertainty distribution, the sigma of which come from the spread of the measured results, and was then compared to similar results at a different wavelength. The refractive index is with respect to fused silica. The sum of the probabilities was normalised to one, for all the possible results in this parameter space.*

3.6 Conclusion

This experiment has demonstrated a possible method of non-destructive testing of fundamental aspects of hydroxide-catalysis bonding. Due to the lack of other opinions for non-destructive testing, this could have significant relevance to many areas of research and industrial applications. The sample investigated showed a thickness of 120 ± 20 nm and a refractive index of 1.433 ± 0.002 at 633 nm. These values are within the range of published data for bond between material of this type. An extension of this work to larger sample size including: different materials, bonding fluids, bonding fluid concentrations and volumes, etc, would provide an enormously useful body of knowledge.

An interesting extension of this experiment could be to monitor the bond properties as it cures, which has been shown to take approximately four months at room temperature for a bond of the type studied here. The large dynamic range of heterodyne detection could effectively cover the likely large ~ 100 dB change in signal size. Measuring the reflectivity of a bond and coupling this with knowledge of strength testing could lead to improved quality control for future bonds. It is possible that a way to correlate strength with bond refractive index or thickness could be found. The likelihood is that the water is still being removed from the polymer chains and the process of dehydration continues long after the bond has become a firm join between the two substrates. This is consistent with the bond reaching full strength after approximately four months.

Chapter 4

Optical Modeling of Second Generation Fibre Injectors

4.1 Introduction

A fibre coupler is used to transfer the propagation of light between free space and an optical fibre. Their fundamental function is to hold a fibre end at the focal point of a lens. Light passing from the fibre is focused into a collimated beam at the output face of the lens, and vice versa, a collimated beam incident on the lens is focused into the fibre.

In LISA PathFinder (LPF) and in all current and foreseeable LISA-type missions, fibres are used to transfer light between: the lasers and the optical modulators, and the modulators and the optical benches. As a result there is a need for a fibre coupled beam delivery system capable of producing an output beam vector with positional and angular requirements. These are generally of the order $\pm 10 \mu\text{m}$ and $\pm 20 \mu\text{radians}$, depending the requirements set

by the optical sensing or the telescope. The changes in the output beam vector are typically thermally driven and due to this all commercial fibre couplers are generally unsuitable. They are typically made from metal and as a result they don't have the required thermal stability to satisfy the needs of LPF or LISA-type missions.

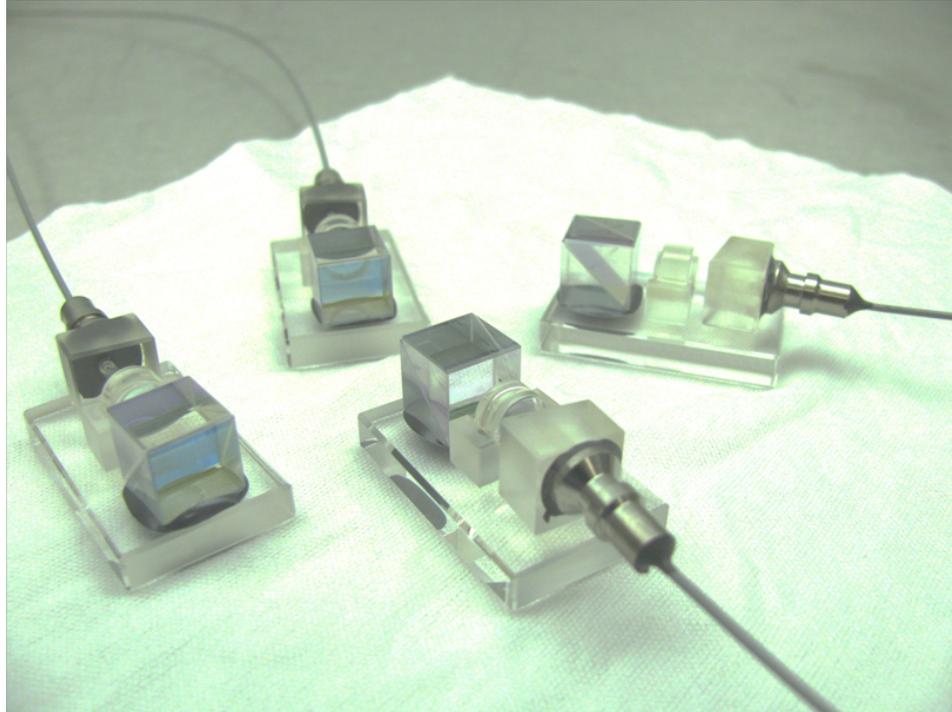


Figure 4.1: A picture of four LPF FIOS. Each FIOS shown has: a fibre holding block, a mounted lens, and a polarising cube, all bonded to a common baseplate. The fibre emerging from the back of the glass block is protected by a metal guard and a rubber guard (not shown). The end of the fibre is flush to the surface of the block facing the lens. The air gap between this face and the lens allows the beam to expand before being collimated by the lens. The polarising cube then cleans the polarisation of the output beam.

The LPF project solved this problem by designing and building a fibre coupler made primarily of glass (see figure 4.1), which was called a Fibre Injector

Optical Sub-assembly (FIOS) [41] . The glass used was fused silica with a Coefficient of Thermal Expansion (CTE) $5.5 \times 10^{-7} \text{ K}^{-1}$, rather than steel with a CTE of $1.1 \times 10^{-5} \text{ K}^{-1}$ for the construction material of the fibre coupler, greatly enhanced the thermal stability [46] These first generation FIOS, were built and successfully demonstrated for the LPF project; however their design needed to be upgraded to meet the stricter requirements of the eLISA project [47]. Presented are the designs considered, the testing conducted, and the output of the completed second generation LISA FIOS.

4.2 Design concepts

4.2.1 eLISA FIOS requirements

The design goals of the upgraded eLISA FIOS were: improved thermal stability, 1 mm radius beams collimated on output from the lens, a wavefront flatness better than $\frac{\lambda}{30}$ RMS, and low risk of failure. This is in addition to the LPF FIOS requirements of precision alignment, space-flight qualified design and being made from non-ferromagnetic material. The improved stability was to come from careful consideration of the material CTE throughout the design. The output beam was controlled by the focal length of collimating optic, and its focal position with respect to the fibre end. The wavefront quality was defined by the surface quality of the fibre end, the output face of the lens figure and the homogeneity of the glass the beam propagated through. Finally the risk of failure of the FIOS was to be reduced through the use of a transparent cover, bonded over the exposed fibre end, in the fibre holding surface. This allows the beam from the fibre to expand up before leaving the glass, which means that the risk of the particulate contamination causing a failure of the FIOS is reduced.

4.2.2 Design options

During the process of updating the FIOS from LPF, three main designs were considered. The optical design of each can be seen in figure 4.2.

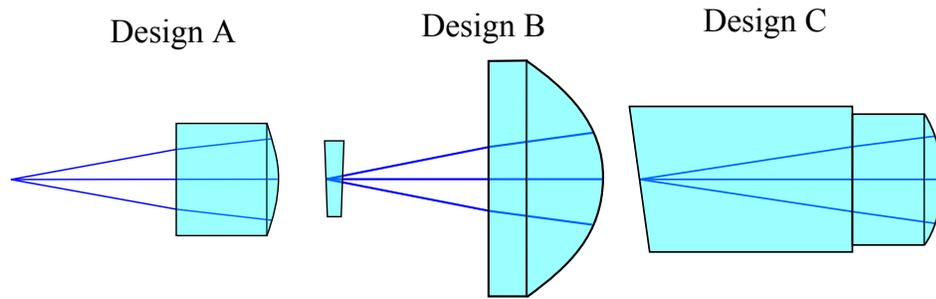


Figure 4.2: A selection of Zemax models showing the various designs considered. Design ‘A’, is similar to the LPF FIOS in design. Design ‘B’ is all fused silica, including the lens and cover slip. Design ‘C’ uses a fused silica spacer between fibre end and lens. Each design is modelled from the virtual waist of a fibre end (not shown).

Design A

This design is similar to the LPF FIOS, though the lens is a different focal length to generate a larger beam, and the glass of the lens is a lead free version of that used for LPF. The design is low risk to build as it has already been demonstrated for LPF with success. It does not mitigate the issues generated by differential thermal expansion, illustrated by figure 4.3, and does not cover the fibre end.

Design B

Design ‘B’ has a wedged cover slip and fused silica asphere (see figure 4.2). This design was expected to produce the required stability, through the use of all fused silica material for both optics and structure. The cover slip avoided the risk of contamination landing on the bare fibre end, which has the potential to cause a failure. The fibre end radius of $4\ \mu\text{m}$ makes this unlikely; however

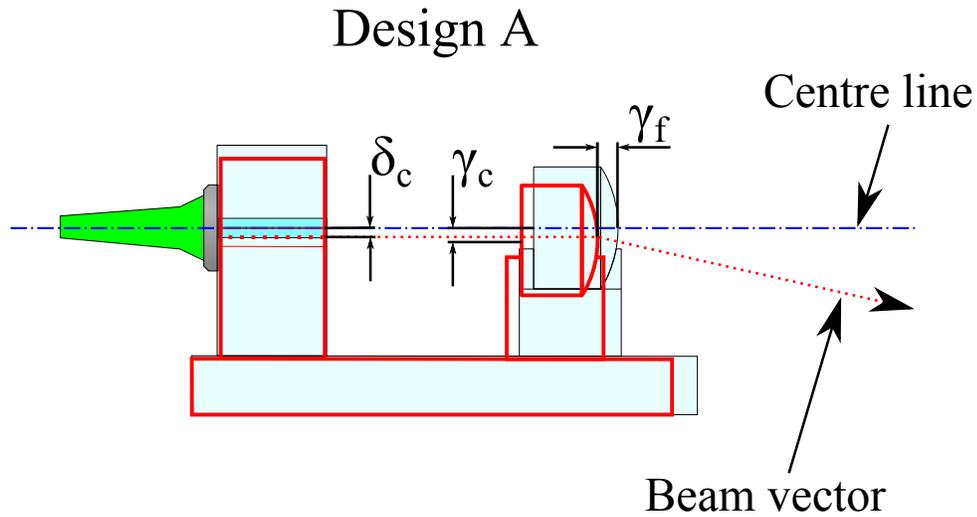


Figure 4.3: *Illustration of thermally driven effects on the LPF FIOS and on Design ‘A’. The blue centre line of the design, coaxial with the fibre and the lens at construction temperature. After a drop in temperature, an exaggerated effect on the dimensions is shown by the red outlines. The change in the fibre end height is denoted ‘ δ_c ’. The shorter physical distance ‘ γ_f ’ leads to a defocused fibre end and so less collimated output. The change in lens position ‘ γ_c ’ is not equal to ‘ δ_c ’, which results in the output vector being angled with respect to the centre line and is due to the CTE of the lens being different to that of fused silica.*

this also means that sub-micron particles could block large areas of the beam emerging from the fibre end, thus causing a critical failure of the FIOS.

The reason the design was dropped was because fused silica lenses of the appropriate focal length proved difficult, or prohibitively expensive to source. Lenses with longer focal lengths were modelled to have favourable tolerances; however experimental testing of the lenses showed significant beam aberrations, which varied lens-to-lens (see figure 4.4). The experimental measurements of the variation in output beam quality was attributed to the figure error of the

aspheres.

Lens-to-lens variation

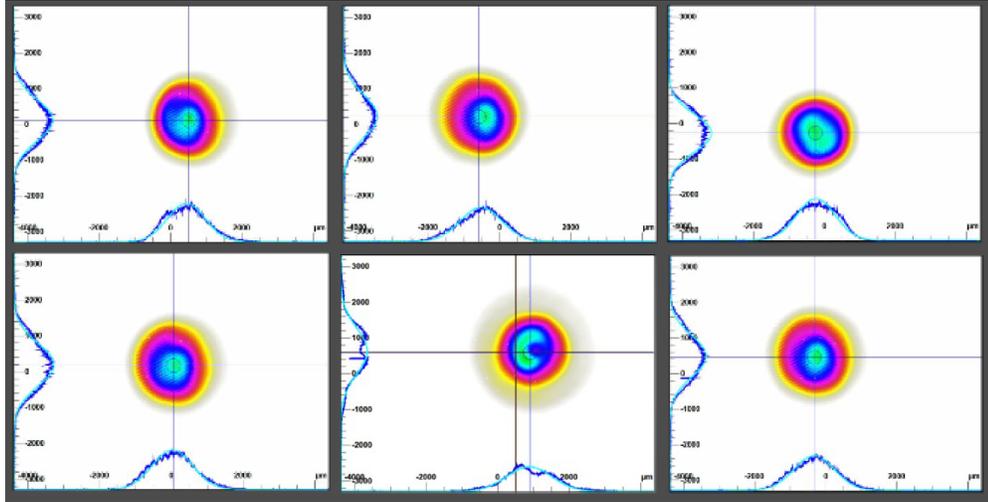


Figure 4.4: Presented is a selection of collimated beam profiles from nominally identical fused silica aspheres. Each beam originated from a cleaved fibre, which had been shown to give acceptable results with other lens including the LPF FIOS lens.

Design C

The design of FIOS settled upon was a quasi-monolithic design with a fibre holder, spacer and lens all bonded together (see figure 4.2). The design provided a stable beam through the coaxial joining of the optical components. The components in the beam path were fixed so their centre of thermal expansion was in line with the beam (see figure 4.5).

This design means all thermal changes are symmetrical around the beam, thus maintaining alignment of the fibre end to that of the lens centre. This allowed for the use of other materials for the lens with higher CTE and therefore

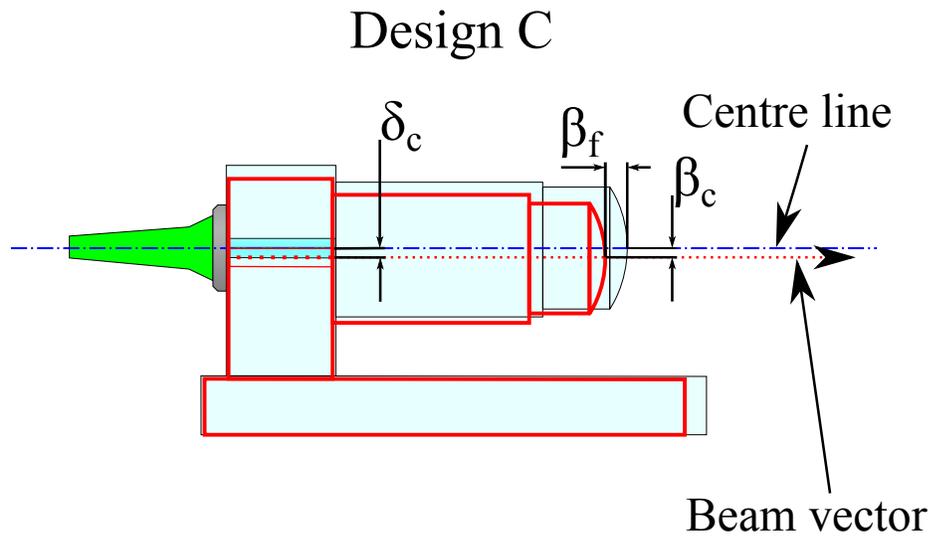


Figure 4.5: *Illustration of thermally driven effects on design ‘C’. The blue centre line of the design, coaxial with the fibre and with the lens at construction temperature. After a drop in temperature, an exaggerated effect on the dimensions is shown by the red outlines. The change in the fibre end height is denoted ‘ δ_c ’. The shorter physical distance ‘ β_f ’ leads to a defocused fibre end and so less collimated output. The change in lens position is ‘ β_c ’ is equal to ‘ δ_c ’, so the output beam vector is parallel to the centre line, resulting in no angular deviation.*

greatly increased the number of possible lenses. The result of thermal expansion modeling showed that to first order the only thermally driven change was the beam radius, which was due to the change in optical path length. This was the design chosen to be carried forward for eLISA, as it achieves all the required goals set for eLISA, including the covered fibre end and the improved thermal stability of the output beam vector.

4.3 FIOS components

4.3.1 FMA and fibre guard

The Fibre Mounted Assembly (FMA) and the fibre guard are inherited directly from LPF. The only changes are that the fibre is from a different source, while the FMA block is a little taller, to accommodate the fused silica spacer. The fibre ferrule is the same as the LPF design [41].

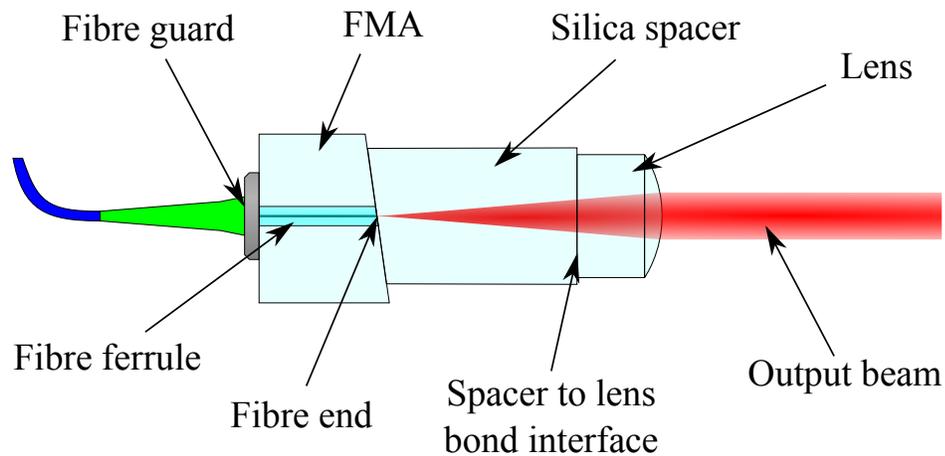


Figure 4.6: *Illustration of the quasi-monolithic FIOS design. Labeled are all the key parts, including the new item, the fused silica spacer.*

4.3.2 Silica spacer

The spacer between the FMA and the lens was made from fused silica and the beam would pass through the spacer before reaching the lens, where it would be collimated. The spacer provided protection for the expanding beam from particulate contamination, by allowing the beam to expand in a medium. The expansion in glass meant that the spacer was longer than the air gap

distance, due to the higher refractive index. The nominal spacer length was determined from the Zemax model to be 11.698 mm for a refractive index of fused silica at 1064 nm of 1.445. This value was dependent on the thickness of the lens and did take into consideration the refractive effects of going from fused silica to the glass of the lens. The interface between the spacer and the FMA face containing the fibre end, was angled at 8° to eliminate potential reflections returning down the fibre, and the spacer was angled to compliment this. By angling the face there was a potential for a trade off between the lateral position of the spacer and the thickness of the spacer as seen by the beam.

4.3.3 Spacer to lens interface

The spacer to lens interface was modelled in detail using Zemax. One of the issues foreseen was the effect on the focal position of a lens, when the light passes from silica into the lens material – of different refractive index – rather than from free space for which the lens was designed.

The small refraction of light as it crosses the boundary between the medium of air and the medium of the lens would be greater than the effect of light traveling from silica to the lens. It was thought that a flat glass surface could refract a curved wavefront – like that of the expanding beam from the fibre end – in a manner that could be described by a ‘weak’ lens. This effect would be reduced if the index of the spacer and lens were more evenly matched than the index of the air to lens.

To measure this effect a test was performed using Zemax and Matlab. A beam from a waist $4\ \mu\text{m}$ – representing the beam from a fibre end – was propagated through two different setups. The first beam was allowed to expand in free

space without passing through anything else. The second expanded up to a size which was ~ 0.5 mm in radius, before passing through a 2 mm thick piece of fused silica glass with flat sides. The beam radius for both simulations was measured – post glass interface – and plotted in Matlab to give the predicted ω_0 value of the beam waist and its position.

The results showed an effect – a slight focusing of the light – but it was negligible and was treated by Zemax as an aberration of the beam. The approximate change in waist size ω_0 was less than a part in 10^4 , while the position of the waist moved by $< 0.1 \mu\text{m}$. When this effect was looked at in the FIOS design, both were negligible compared to the tolerances placed on: the lens focal position ($< 0.1 \mu\text{m}$ compared to $\pm 10 \mu\text{m}$) and the output beam size ($0.8\text{mm} - 0.01\%$ compared to $0.8 \pm 0.1\text{mm}$).

4.3.4 Lens options

All of the experimentally measured beams from readily available fused silica aspheric lenses showed variable beam output quality (see figure 4.4). Many had gross beam aberrations and these defects varied from lens to lens. It was decided that although it would be best to have a fully monolithic fused silica FIOS, the prescription of silica lens was not available in the required quality. To have such lenses made, would have been prohibitively expensive with no guarantee of success. As a result other types of glass were considered before settling upon the same type as used in the LPF FIOS. The lens chosen for the eLISA FIOS was a Geltech 352220 [48], which has the benefits of being space-flight qualified for LPF and a lens, of similar prescription, was used successfully in the LPF FIOS [41].

The discrepancy in the CTE of the lens compared to fused silica raised two

Table 4.1: *Geltech 352220 lens specification*

Material	ECO-550
Effective focal length	11.21 mm* $\pm 1\%$
Diameter tolerance	± 0.015 mm
Thickness tolerance	± 0.040 mm
Surface quality	40/20
RMS wavefront error	Diffraction limited
Clear aperture	5.50 mm
Refractive index	1.592*
ECO-550 CTE	$11.5 \times 10^{-6} \text{ K}^{-1}$

concerns regarding thermal effect over the range of temperatures it might face in launch and operation: firstly, would a bond between two materials of dissimilar CTE hold; secondly, would the lens figure change significantly?

The first thermal problem was tested using a bonded sample and an oven. The sample was a fused silica disk, bonded to a Geltech lens. After allowing the bond to cure for a few days it was thermally cycled over a temperature range larger than expected during launch or operation. The outcome of this test showed no visible signs of bond degradation.

The second was investigated by finite element modeling the FIOS, heated by 20°C. The model showed where the lens would expand and by how much. The lens when heated began to mushroom out at, as the base was fixed to the fused silica spacer. The lens became 2 μm larger in diameter, and 1.6 μm thicker through the centre of the lens with a thickness change of only 1.5 μm at the edges of the lens (see figure 4.7). This was then used to produce an adjusted lens prescription, which was loaded into ZEMAX for testing. This modified prescription showed no significant new beam aberrations or changes

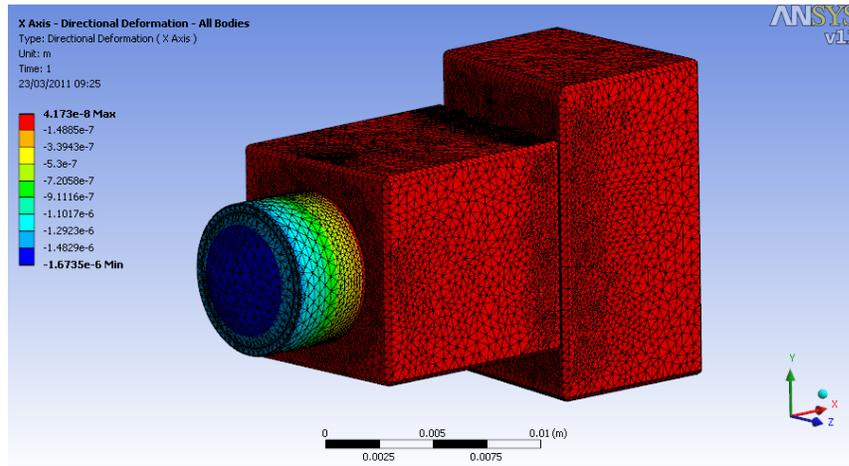


Figure 4.7: *Finite element analysis of the FIOS design by M. Perreur-Lloyd.*

in the output waist size or location. Thus the beam was judged to be within the specification required.

4.3.5 FIOS focal length

The length of the spacer is one of the biggest challenges in the realisation of the design. Although the optical design shows that it must be 11.698 mm for a lens of nominal thickness, it has an angled face. The bonding of the spacer onto the face of the FMA sets the distance between the virtual waist of the fibre end and the output face of the lens. This distance must be equal to the lens focal length and sets the output waist of the FIOS beam and therefore the collimation of the output beam. The spacer must position the lens focal plane to less than $\pm 10 \mu\text{m}$ of the virtual waist of the fibre, in order to collimate the output beam correctly (output waist position $\pm 1 \text{ m}$ from the lens). The angled face of the spacer which is bonded to the FMA is the only point of adjustability in this optical design in terms of focus. This makes this the critical step in the FIOS construction.

4.4 Optical performance

4.4.1 Modelled beam parameters

The modelled output beam parameters of the eLISA FIOS can be seen in table 4.2 and in figure 4.8. The output beam waist is dependent on the virtual waist of the fibre, which was taken to be $4.5 \mu\text{m}$ from precision measurements of a beam from a similar fibre end. For a virtual waist of this size the output beam waist is 0.791 mm, which yields the profile seen in table 4.2. The virtual waist size is set by the mode field diameter (MFD) of the fibre. This varies by $4.5 \pm 1 \mu\text{m}$ between batches and could significantly affect the output beam size. To counter this all fibres used in this project were from the same batch with an expected variation of MFD of $4.8 \pm 0.05 \mu\text{m}$.

Table 4.2: *modelled beam profile from eLISA FIOS*

Distance from lens	50 mm	260 mm	800 mm	2000 mm	4500 mm
Beam radius	0.79 mm	0.80 mm	0.86 mm	1.17 mm	2.08 mm

To ensure that the beam quality of the FIOS was maintained to a suitable level a tolerance analysis was performed in Zemax. The following tolerance values ensured that the output beam was $> 95\%$ Gaussian and had $> 90\%$ ellipticity: focal position $\pm 10 \mu\text{m}$, lateral position $\pm 200 \mu\text{m}$, and angle of incidence $\pm 1^\circ$.

Figure 4.8, shows the phasefront and beam intensity profiles over ~ 4 beam radii. The virtual waist from a fibre is a Bessel function, which in the far field approaches a Gaussian profile. In the model an approximation is made by using a Gaussian profile for the virtual waist of the beam from the fibre.

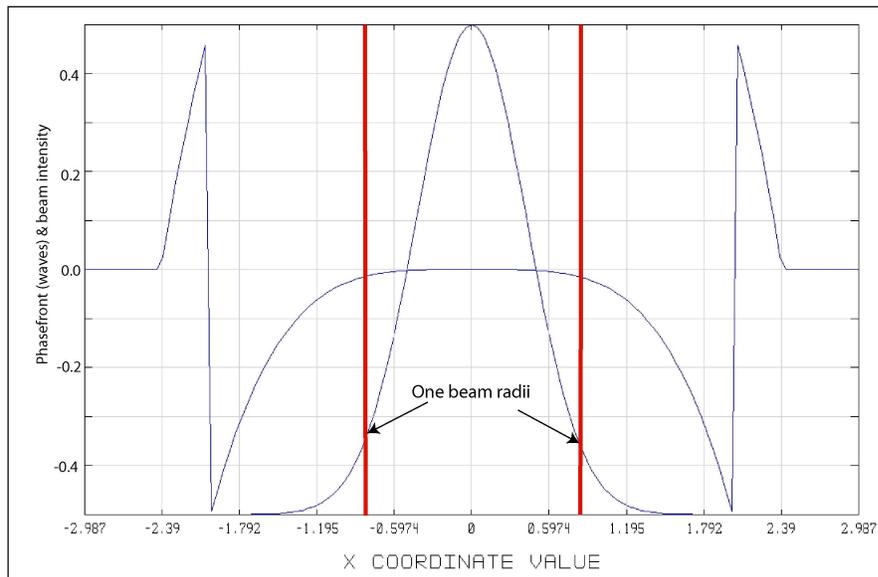


Figure 4.8: *Graphic from Zemax showing the profile of the intensity profile with the phasefront of the beam at the lens output face. Marked onto the graph are two lines denoting $1/e^2$ or one beam radii of 0.79 mm. The scale of the phasefront plot is 10^{th} of a wave per division, while the x-axis is in millimeters.*

4.4.2 Experimental results

Experimental testing of the both the lenses and the FIOS post-construction were conducted. The lenses were mounted and actuated by a Hexapod, which was used to align them to a fibre end with a precision of $\pm 1 \mu\text{m}$ and $\pm 0.01^\circ$. Table 4.3 shows the measured collimated beam profile of a fibre end, using two different Geltech lens C & D and an air gap. This is then compared to the modelled values for a given input virtual waist from the fibre of $4.5 \mu\text{m}$. The beam profiler used to measure beam radii was subject to an error of $\sim 10\%$.

Table 4.3: *Measured beam profile of a collimated output beam from two nominally identical Geltech lenses C & D , as compared to the model predictions where the waist is 0 m or 1 m from the lens.*

Distance from lens	Geltech lens C		Geltech lens D		Model, waist	
	beam radii	ellipticity	beam radii	ellipticity	at 0 m	at 1 m
30 mm	0.92 mm	99.7%	0.92 mm	99.9%	0.84 mm	0.84 mm
130 mm	0.92 mm	98.0%	0.90 mm	97.9%	0.85 mm	0.83 mm
1000 mm	0.90 mm	96.6%	0.88 mm	98.9%	0.89 mm	0.80 mm
2000 mm	1.08 mm	97.5%	1.04 mm	97.0%	1.17 mm	0.99 mm

The beam profiles show good agreement with those of the model; however there is a uncertainty on the absolute size of the beam radius due to the profiler measurement error.

Following the profile of the lenses, a fully bonded eLISA FIOS was assembled, shown in figure 4.9. The output beam profiles of the first two FIOS to be made are shown in figure 4.10. The measured output beam profiles have been fitted to give the position and size of the output beam waist. The output waist ω_0 of FIOS #1 is $690 \mu\text{m}$ and it positioned 990 mm behind the FIOS. The output

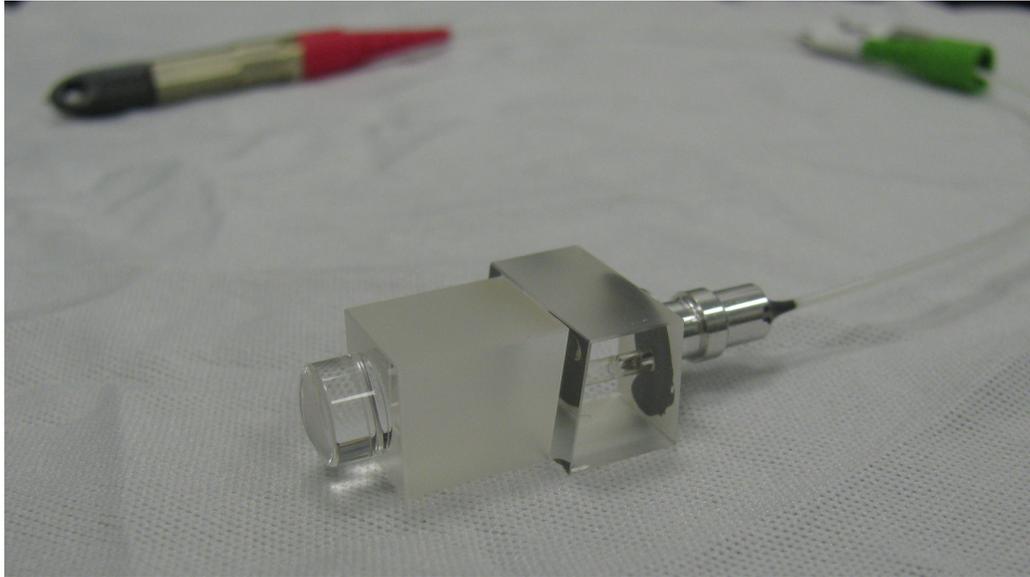


Figure 4.9: *Picture of the first fully bonded eLISA FIOS*

waist of FIOS #2 is $900\ \mu\text{m}$ and it positioned 650 mm behind the FIOS.

The difference in the FIOS output profiles was attributed to a combination of effects associated with the fibre end, only after further testing of the lenses was conducted, see chapter 5. A topic of continued investigation, these include: bonding onto a fibre end inducing a change of the output beam position or size, and the discrepancy between the virtual waist of a beam output from a fibre end and the physical location of the fibre end.

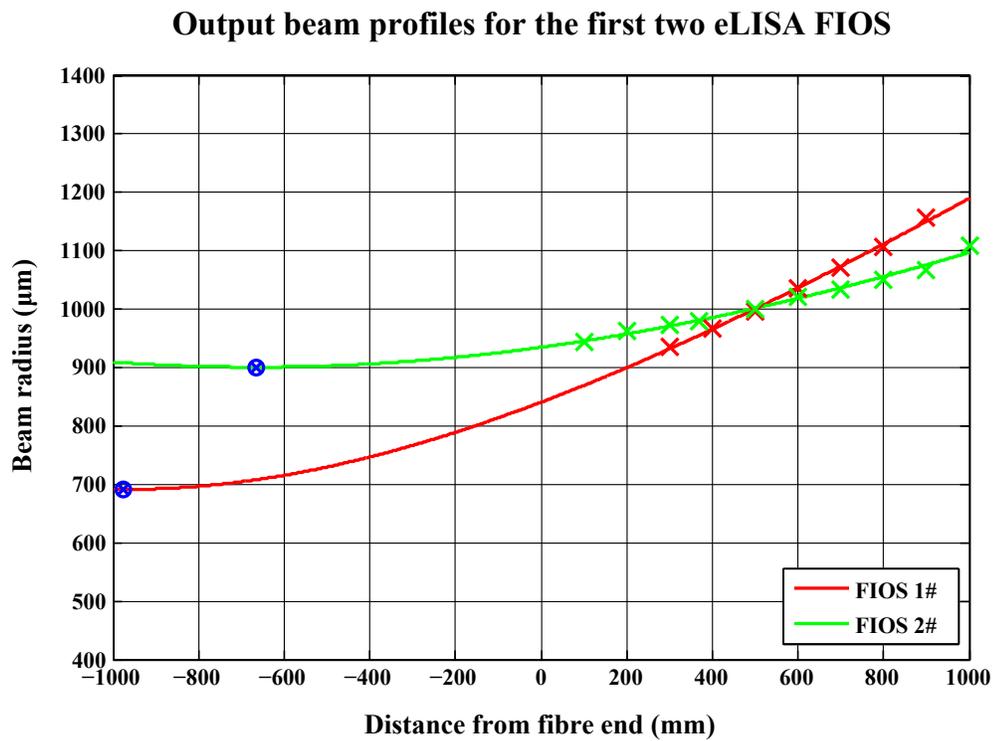


Figure 4.10: Measured beam profiles of the first two eLISA FIOS. Shown are the measured values (\times), the best fit Gaussian beam profiles ($-$) and their ω_0 (\otimes).

4.5 Conclusion

The eLISA FIOS was designed to produce a beam of higher thermal stability and larger beam radius than that of the LPF FIOS. The modelled thermal sensitivity of the eLISA FIOS to focus position is approximately twice that of the LPF FIOS: LPF FIOS is $\sim 3 \times 10^{-8}$ m/K, while the eLISA FIOS is $\sim 6 \times 10^{-8}$ m/K. However over the expected operating temperature range of eLISA ($20 \pm 10^\circ\text{C}$), that is a change on position of less than a 10^{th} of the total tolerance of focal position so this is not an issue.

The thermal stability improvements of the eLISA FIOS in lateral position of the lens centre to the beam centre is expected to be much better than for LPF (see figures 4.5 and 4.3). Determination of an exact coupling factor is difficult to define from the model, as the focal position of the virtual waist and the displaced centre of the lens from the waist centre, combine to give the resulting output beam vector of the FIOS. Future testing of the eLISA FIOS's long term stability during thermal cycling, would provide a measure of the coupling factor. This is expected to be more than sufficient for the purposes of eLISA.

Chapter 5

Optical Testing of Short Focal Length Components

5.1 Introduction

Verification of optical components – through physical testing – is essential in ensuring confidence in the designs which utilise them. This acts as a cross-check between the modeled performance and the experimental results to ensure that the performance is as expected. Important factors to address are the tolerance of the optic placement that ensure the required results and that their finish is of the required standard. Examples being: the tolerance on alignment of a lens, and surface figure error of a lens.

To deal with the issues which arise from using short focal length optics, high precision mountings and adjusters are required. Much of the work carried out and presented in this thesis is possible only through the use of devices capable of moving/measuring at the micron level. These include but are not

limited to: Hexapods, a Coordinate Measuring Machine (CMM), a range of piezo electric actuators, and Vernier scaled mountings. There are also beam measuring devices like: CCD cameras, quadrant photodiodes, and wavefront cameras capable of measuring the beam interacting with such components.

After completing the two proto-type eLISA FIOS, there was an immediate issue with the collimation of the beam. The first FIOS produced a beam which had a smaller output waist size than planned. The output beam modeled as being a possible $\pm 100 \mu\text{m}$ variation in the focal position of the lens. This meant the beam did not have the desired collimation; nor did it look to be within the expected tolerances, of the components used. As part of the investigation into the cause of these unexpected beam parameters seen figure 4.10, a dedicated experiment to measure the focal lengths of the lenses was conducted.

5.2 Method

The lens in question was the Geltech 352220 asphere (see table 4.1), which was similar to the LPF FIOS lenses in focal length and in the glass from which it was made. The lens was chosen because the figure quality was more consistent lens to lens, than any of the fused-silica lenses which were considered. This was the most important factor in making the decision on which lens type to procure as repeatability was a key factor. The reason is due to its construction; made in a mold from ‘softer’ glass improved the accuracy of the figure, over a fused-silica lens which must be polished to shape; a process which is much more prone to figure error. Although it is possible to get the required figure of fused silica asphere, the costs were beyond the project and therefore the molded lenses were used.

5.2.1 Measurement principles

As with many experiments there were two possible measurement types, an absolute measurement of the lens focal length or a relative measurement among multiple lenses. The methods for measuring the absolute focal length of a lens are varied, and many ideas were considered. The most sensitive output parameter which can be measured is the location of the output beam waist. Therefore a method to determine the focal length is to change the separation between an input beam waist and the output face of the lens by a known amount. The resulting change in output beam waist position allows for back calculation of the focal length.

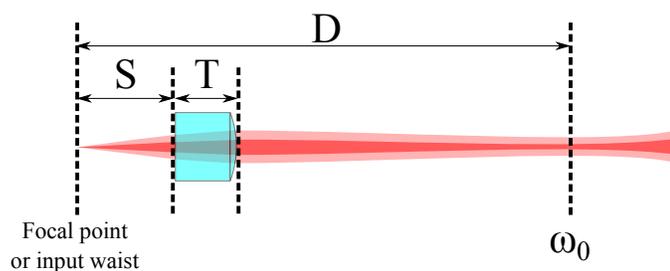


Figure 5.1: *This illustration shows the basic variables of the measurement system. The separation of the fibre end and the lens rear surface ‘S’, is the key variable, as the thickness ‘T’ of each of the lenses was already known.*

5.2.2 Methods of absolute measurement considered

Fluid measurements

Looking first at the absolute measurements, one method was to measure the beam output of a fixed lens at a distance from a fibre end in air and then flood the gap ‘S’ between them with a fluid, of known refractive index. By comparing the output beam’s parameters, before and after the fluid was added, the focal length of the lens could be calculated. The fluid would yield an extra optical path length which could be either: modeled to find the focal length which would produce such output beam positions after a change of separation, or simply compensated by changing the separation of the fibre end and lens until the beam waist was in the same place again.

The issues were: keeping the fluid in the right place so as not to distort the output face of the lens; a small effect due to the change of refractive index as the light entered the back of the lens; ensuring the fluid was cleaned away fully before a new measurement was taken, as well as the practical challenges involved. The small refractivity difference between the fluid and non-fluid

measurements is the boundary at the back of the lens. The presence of a higher refractive index medium means the refraction of the light will be slightly less than for air. The magnitude of the refraction effect is dependent on the refractive index of the fluid. To correct for such an effect optical modeling would be needed to produce a correction factor.

Looking at figure 5.1 the following parameters can be assigned in the case of using a fluid to change the separation: ‘ D ’ as the distance between the fibre end and the output waist ‘ ω_0 ’ of the lens; ‘ Δd ’ as the change of waist position before and after applying the fluid; ‘ S ’ is the separation of fibre end and lens; ‘ ΔS ’ is the change of separation, after the fluid’s application, to bring the output waist back to the same position as before; ‘ n_{fluid} ’ is the refractive index of the fluid; ‘ n_{lens} ’ is the refractive index of the lens; ‘ T ’ is the thickness of the lens. In this situation the following statements describe the system before and after the fluid is applied.

Before fluid and assuming the refractive index of air to be 1:

$$S + \frac{T}{n_{lens}} \rightarrow \omega_0 \text{ at } D$$

and after the fluid:

$$\frac{S}{n_{fluid}} + \frac{T}{n_{lens}} \rightarrow \omega_0 \text{ at } D - \Delta D.$$

The Separation changed to reposition waist:

$$\frac{S}{n_{fluid}} + \frac{\Delta S}{n_{fluid}} + \frac{T}{n_{lens}} \rightarrow \omega_0 \text{ at } D$$

by equating the first and last statements the following equation can be gained:

$$S = \frac{S}{n_{fluid}} + \frac{\Delta S}{n_{fluid}} \tag{5.1}$$

rearranging to give:

$$S = \frac{\Delta S}{(n_{fluid} - 1)} \quad (5.2)$$

This allows for the absolute size of the separation between the fibre end and the lens rear face to be determined from the fluid's refractive index and the reposition distance alone.

The Effective Focal Length (EFL) is the focal length of an optic or combined focal length of a series of optics in air. In this case the lens optical thickness must be converted to the corresponding length in air. Combining this with the thicknesses of the lenses and their refractive index, gives the EFL as:

$$EFL = S + \frac{T}{n_{lens}} \quad (5.3)$$

Given the accuracy with which the refractive index of fluids and glasses are known (1 part in 10^6 or more), the uncertainty in the value of '*EFL*' is limited by the accuracy of the distance '*ΔS*' and thickness '*T*' and the ability to determine the position of the output waist. Although as mentioned before thermally driven effects could have proven problematic in this measurement. The accuracy of these distances was measured to sub-micron using a Hexapod (*S*), CMM (*T*) and a wavefront camera (*D*). Thus the absolute '*EFL*' could be measured to an accuracy of a micron or less, using this method.

Fibre touch measurements

A second method of absolute measurement was to measure the displacement between the lens back surface and the fibre end '*S*'. By starting with the lens touching the fibre end, the fibre end could be moved away until the output waist of the lens was at the appropriate position. This would provide a direct

measurement of ‘ S ’ with an uncertainty dependent on the accuracy of the Hexapod and the ability to judge when the lens and fibre end were in contact. This allowed the EFL to be calculated by taking the gap between fibre end and the back face of the lens ‘ S ’, adding on the optical path length from the specific thickness ‘ T ’ of the lens

An absolute measurement method considered involved using a fluid to change the refractive index of the gap between the lens back and the fibre end or beam source. This was impractical to implement, due to the need to mount the fibre end and lens facing vertical, while using a considerable volume of fluid ~ 2 litres for each measurement. Ultimately these fluid measurements were not conducted due in part to time constraints but mainly the impractical nature of the mechanical assemblies required to conduct the tests. Of note is also the thermally induced refractive index changes of the fluid, induced by the beam passing through it.

Absolute results

Several absolute measurements were made using the method where the fibre end was moved until in contact with the back of the lens. These results were limited by significant variation, which were most likely due to the uncertainty in judging the point of contact between the fibre end and the lens. The final value was an average absolute EFL of 11.24 ± 0.1 mm.

5.2.3 Relative measurements

A third method was to use a fixed lens position and fixed output waist target to measure the relative focal lengths of the lenses. The input waist would be adjusted to position the output waist on a wavefront camera, with the input waist position referenced to the lens rear surface, as shown in figure 5.2. This method used a fixed coordinate system of the Hexapod to record the position of the lenses when they produced the desired output beam waist position. The results of which would provide relative measurements of the focal lengths. This method relied on the lenses being replaced into the holder to sub-micron accuracy.

If the measurements of the lens focal position are made relative to one another then the following change is made to the variable ‘ ΔS ’. It becomes the change in separation of input waist and lens position, when comparing one lens to another. Given equation 5.3, it follows that ‘ ΔEFL ’ between lenses is:

$$\Delta EFL = \Delta S + \frac{\Delta T}{n_{lens}} \quad (5.4)$$

Where ‘ ΔT ’ is the difference in thickness between the lenses. This is the variation of EFL for one lens compared to another, or the overall average.

5.3 Experimental method and setup

The chosen method of beam detection was to use a wavefront camera and then position the lens with respect to the fibre end so as to produce a waist on the wavefront camera. The wavefront camera was the most sensitive detector available to measure the waist position of a beam, the waist position could be measured with sub-micron precision.

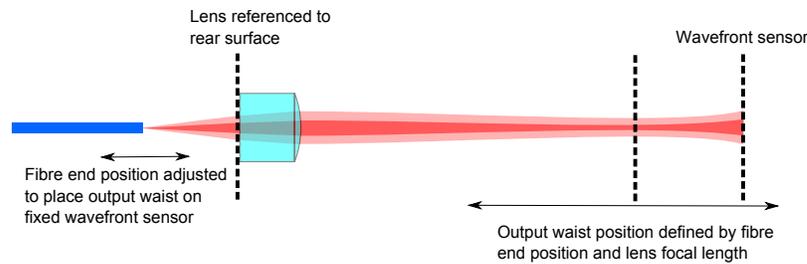


Figure 5.2: *This illustration shows the variables and the reference points of the chosen measurement system. Changes of the separation between fibre end and lens moves the output beam waist position, which was measured by a wavefront sensor at a fixed distance from the fibre end.*

5.3.1 Optical setup

The optical setup as illustrated in figure 5.2 was used to measure the focal lengths of the FIOS lenses. The optical system was a simple object projection, where the object – the fibre end – was projected onto the wavefront sensor. To adjust the position of the projection on to the wavefront sensor, the fibre-to-lens distance was adjusted by moving the fibre end.

The measurements taken were the position of the fibre end in the Hexapod coordinate frame, and the distance between the fibre end and the wavefront

camera. The distance between the wavefront camera and the fibre end was measured by removing the lens and using the wavefront camera to measure the curvature of the beam from the fibre end. The camera was able to measure this distance to less than 1 mm. The fibre end movement was accurate to $\pm 0.5 \mu\text{m}$, limited by the Hexapod accuracy [49].

5.3.2 Lens placement and orientation

The position and repositioning of the lenses was highly important in ensuring the repeatability of the results from lens to lens. The lenses were placed onto the front of the U-groove and then pushed firmly back against the ball bearing as shown in figure 5.3. The U-groove had a rough machine finish which seemed to catch the lenses. In the first few attempts to replace a single lens, there were a number of instances where it was out of position by tens of microns. To help mitigate this problem, a fluid was used to lubricate the motion of the lens. A few drops of methonal were placed on the U-groove before the lens. This fluid reduced the chances of the lens catching on the rough surface and also evaporated entirely, helping ensure it was not affecting the results.

Another parameter which was kept constant was the lens orientation during each measurement. A small spot was marked on the side of the barrel. This spot was kept at the top of the lens to maintain the same rotational placement in the axis of the barrel/U-groove of the holder for each measurement. If the lens face was not centred in the barrel of the lens, this would avoid beam misalignment on repeat measurements, in fact the lens faces were centred in the barrel of the lens to better than one micron. This was judged by the deviation of the beam vector during rotation of the lens in the U-groove.

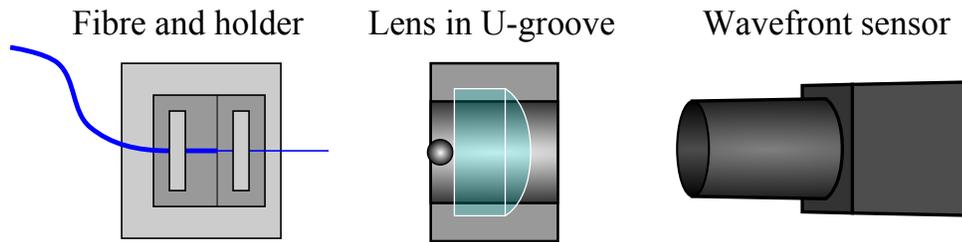


Figure 5.3: Graphic to illustrate the details of the optical setup. The fibre was held by two clamps, which pushed it into a groove on the metal plate. The end of the fibre protruded a distance of 25 mm to allow it to be held close to the lens. The lens holder is a U-groove, and had a ball bearing glued at one side. The wavefront camera was atop a lab-jack which was adjusted in order to centre the beam on the wavefront camera.

5.3.3 Wavefront sensor alignment

The wavefront sensor was adjusted by gentle tapping of the lab-jack to ensure that it could be centred on the beam to less than $100\ \mu\text{m}$, using the lab-jacks inertia to achieve small movements. This method was used to centre the beam in the plane of the bench, after the jack height had been moved up or down to centre the beam. The distance between the wavefront camera was maintained at $\sim 1.1\ \text{m} \pm 1\ \text{mm}$ from the fibre end. The distance was remeasured between each lens focal position measurements.

5.4 Results

5.4.1 Measured focal positions

The results for the lenses were combined with the pre-measured lens thicknesses. The focal position results in table 5.1 are the relative separations ‘ ΔS ’ of the lens rear surface and the fibre end.

The values of ‘ ΔT ’ are measured relative to the average measured lens thickness ‘ T ’ (5.0125 ± 0.0005 mm). The values of ‘ ΔS ’ are measured relative to average focal position of ‘ S ’ as defined in figure 5.1. This means that the ‘ ΔEFL ’ values are relative to the average EFL of this group of lenses. Using equation 5.4 the ‘ ΔEFL ’ was calculated.

Table 5.1: *Relative EFL of selected Geltech 352220 aspheres*

Lens name	Δ Focal position (mm)	Δ Lens thickness (mm)	Δ EFL (mm)
Test lens	$+0.003 \pm 0.004$	$+0.018$	$+0.014 \pm 0.006$
SN 06	-0.006 ± 0.010	-0.001	-0.006 ± 0.010
SN 07	-0.024 ± 0.019	$+0.016$	-0.015 ± 0.019
SN 08	$+0.002 \pm 0.004$	$+0.005$	$+0.006 \pm 0.005$
SN 09	$+0.012 \pm 0.009$	-0.029	-0.006 ± 0.010
SN 10	$+0.021 \pm 0.007$	-0.033	$+0.001 \pm 0.008$
SN 11	-0.022 ± 0.005	$+0.022$	-0.008 ± 0.007
SN 12	-0.005 ± 0.012	$+0.012$	$+0.002 \pm 0.012$
SN 13	$+0.019 \pm 0.005$	-0.010	$+0.013 \pm 0.006$

5.4.2 Expected uncertainty

The air conditioning in the lab was in need of some maintenance at the time this experiment was being conducted. The air temperature underwent large fluctuations which resulted in temperature drifts of 3° per hour. This was a problem, as the goal of this experiment was to measure the focal length of the lenses to \sim micron accuracy which required a series of measurements over the course of an hour or longer. Given the distances between the fibre end and the lens mountings of ~ 0.1 m it was possible that changes of the order of $\sim 1 \mu\text{m}$ could have occurred between measurements.

One of the factors which was dealt with before this experiment was a measurement of the thicknesses of the lenses. This important parameter was measured using a CMM (Coordinate Measuring Machine) which gave the results in table 5.1. This was needed in the production of the FIOS (see section 4.3.5) to ensure the correct optical path length between fibre end and lens face.

5.5 Conclusion

The reason for this test was to check the EFL of the Geltech 352220 aspheres, to see if that could explain the issues seen with the first two proto-type eLISA FIOS.

The results of relative measurements between multiple lenses (see table 5.1) showed that the variation of the lens EFL was typically less than $\pm 35 \mu\text{m}$ from an average value. Comparing this to the $\sim 100 \mu\text{m}$ modeled variation of EFL between the two lenses used in the first two proto-type eLISA FIOS, the measurements show that if the lenses used for these eLISA FIOS were typical, then they were not the reason for the output beam variation.

The absolute measurement result of $\sim 11.24 \pm 0.1 \text{ mm}$ and the relative measurement result of $\pm 35 \mu\text{m}$, are both within the quoted manufacturer tolerance of $11.21 \pm 1\% \text{ mm}$.

To summarise, the work conducted and presented in chapters 4 & 5 was successful in developing and troubleshooting the new design for the eLISA FIOS. As of writing multiple FIOS have been successfully produced, now that the initial problems in manufacture have been investigated.

Chapter 6

Performance testing of photodiodes

6.1 Introduction

Photodiodes are the main method of detection of science signals on a LISA PathFinder (LPF) and any future eLISA/LISA-type missions. Their role on the optical benches is to detect both the interferometric signals and measure the optical power levels. The photodiodes are the first part in each of the signal measurement chain and play an important role in the overall detection sensitivity. Due to their position and importance, LPF and all current LISA type missions have a built in redundancy in case a photodiode were to fail. The redundancy is implemented by both output ports of the beam combiners being populated by a photodiode. LISA-type missions measure not only the phase of the heterodyne signal, but also utilise what is called Differential Wavefront Sensing (DWS) [50]. This yields valuable information about the alignment in terms of angle, between the two interfering beams. By detecting the relative

phase of the signal in each quadrant of a quadrant photodiode (QPD) the relative angular alignment of the two optical beams can be determined. In the case of LPF, where there are freely floating test masses, this provides a measure of the test mass angular alignment in two degrees of freedom. One of the goals of LPF was to reach a noise floor of $6.3 \text{ pm}/\sqrt{\text{Hz}}$ in the longitudinal readout measurement in a band from 0.1 mHz to 100 mHz. This and other goals placed strict requirements on the photodiodes in terms of bandwidth and spatial uniformity of response. There are also other requirements on how well the photodiodes have to be aligned to the optical beams, but these are not dealt with in this chapter.

This chapter will detail the results of a series of tests performed on several types of photodiode to determine their suitability for both LPF and LISA-type missions. These tests include: spatial uniformity response of the signal magnitude, spatial uniformity response of the signal phase, variation of signal magnitude with frequency, and variation of signal phase with frequency. The photodiodes tested are two different types of 5 mm diameter quadrant photodiodes. The first is an InGaAs quadrant photodiode made by OSI and known as the “InGaAs photodiode”. The second is Silicon quadrant photodiode made by Silicon Sensor and tested in two forms: the commercially available version known as the Silicon Sensor photodiode or “QP22-E”, and the other packaged by lewicki and known as the “Lewicki photodiode”. The Lewicki photodiode has an integrated thermo-electric coupler (TEC), and comes pre-assembled, with long cables. The Silicon Sensor group are now part of First Sensor, and the QP22-E photodiodes are now sold under the name “QP22-Q”.

Table 6.1: Specifications of InGaAs and QP22-E photodiodes

Photodiode	Diameter	Gap size	Amp/Watt	Rev. bias	Capacity
InGaAs	5 mm	45 μm	0.8	5 v (max)	250 pf
QP22-E	5.3 mm	70 μm	0.6*	150 v	3* pf

* values for QP22-E photodiode while under 150 volts reverse bias

6.1.1 OSI InGaAs photodiodes

The LPF flight optical bench has ten photodiodes. There are eight quadrant photodiodes measuring the science outputs of the four interferometers. Each of the four interferometers has a quadrant photodiode on both of its output ports for redundancy. Additionally there are two photodiodes which measure the power levels of the input beams to the optical bench and have different requirements to the other photodiodes.

The InGaAs photodiodes originally procured for that bench were produced by OSI, in their now closed down European operations. This was the second order of these photodiodes as they had been used with success on the engineering model optical bench. They are relatively large InGaAs photodiodes with a diameter of 5 mm and a gap between the quadrants of $\sim 45 \mu\text{m}$. These photodiodes met all of the requirements placed upon them in both the engineering model and the flight model procurements.

The OSI photodiodes suffered a high failure rate during procurement of the photodiodes for the two flight optical benches. Indeed the procurement of photodiodes for the the flight benches started with the first batch of photodiodes being rejected entirely. The photodiodes that are currently fitted to flight bench are from a second batch of photodiodes for the flight benches. A failure rate of +50% was seen at the manufacturing level which continued

to climb during integration of the cabling and subsequent testing which lead to what is a final yield of $\sim 10\%$. This turned out to be just enough functional InGaAs photodiodes to populate the flight bench but not the flight spare bench. This meant that in order to provide photodiodes for the flight spare, new photodiodes from a different source had to be procured.

The list of things that have gone wrong with these photodiodes is extensive. The photodiodes suffered electrical failures pre-delivery to the University of Birmingham which meant there was only around a 50% yield from the wafer. The next major failure of photodiodes came sometime after delivery (it is thought) where the photodiodes suffered an electrostatic discharge. This discharge was enough to damage many of the quadrants and caused holes to appear in the surface of the quadrants which were of the order of a hundred microns in size. Other issues not related to the physical structure of the photodiode have reduced the yield to a level where there was not enough to populate both the LPF flight and flight spare optical benches. Despite the high failure rate the surviving functional InGaAs photodiodes have shown excellent performance in each of the tests conducted here and elsewhere, over the last ~ 6 years.

6.1.2 Lewicki/Silicon Sensor photodiodes

Due to the high failure rate of the OSI InGaAs photodiodes, an inability to reprocure, barely any spares for the flight bench and not even a set for the flight spare bench, a procurement of different photodiodes was made by the UK. The chosen supplier Lewicki, packaged Silicon Sensor model QP22-E photodiodes. These photodiodes have what Silicon Sensor call enhanced sensitivity to NIR light. The photodiodes have a responsivity of 0.6 Amp/Watt

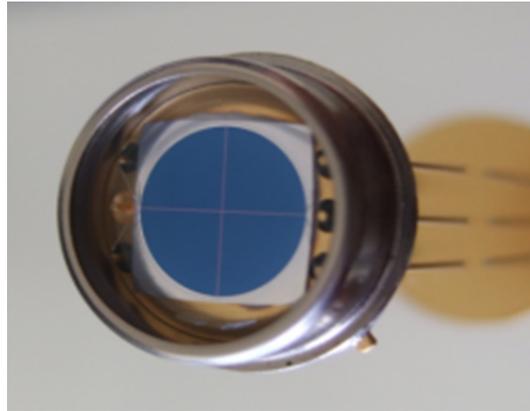


Figure 6.1: *Picture of the Silicon Sensor QP22-E photodiode.*

at 1064 nm compared with less than 0.2 A/W for typical Silicon photodiodes and 0.8 A/W for the OSI InGaAs photodiodes. The reason why Silicon was chosen over InGaAs was that it was seen to be a more robust material, and so less susceptible to damage from electrostatic discharge.

6.2 Photon detection by InGaAs and Silicon

6.2.1 Silicon photodiodes in the visible spectrum

A Silicon photodiode typically operates in the visible spectrum without bias. The photodiode has a p-n junction around which a depletion region of high field forms. Photons of wavelength 633 nm are absorbed after less than $\sim 3 \mu\text{m}$ of Silicon. Photo-electrons formed in the depletion region are quickly swept out by the high field which ensures that the signals are readout promptly.

To improve the performance of a photodiode in terms of time between absorption of the photon and the photo-electron being readout, reverse bias can be applied. This bias voltage is applied across the p-n junction and increases the strength of the electric field which leads to an increase in the depth of the depletion region. This in turn means a shorter travel time for a larger number of photo-electrons to the p-n junction. Statically it is unlikely that a visible spectrum photon would be absorbed outside of the depletion region, thus producing a photo-electron that would wander slowly through the Silicon before reaching the depletion region and then p-n junction where it would be readout. Any photo-electrons generated outside the depletion region will wander slowly through the substrate until reaching the depletion region where it will be readout.

6.2.2 InGaAs photodiodes in near infrared spectrum

The operation of an InGaAs photodiode such as the OSI InGaAs photodiodes at 1064 nm is similar to that of a typical Silicon photodiode with visible light. The InGaAs photodiodes tend to not require bias for operation, however a reverse bias voltage can help to improve the response just like in the case for

Silicon. The absorption depth of a 1064 nm photon in InGaAs is $\sim 1 \mu\text{m}$. This is again like the Silicon case very close to if not inside the depletion region of the photodiode. The photo-electrons have a high mobility and are promptly readout from the photodiode shortly after absorption.

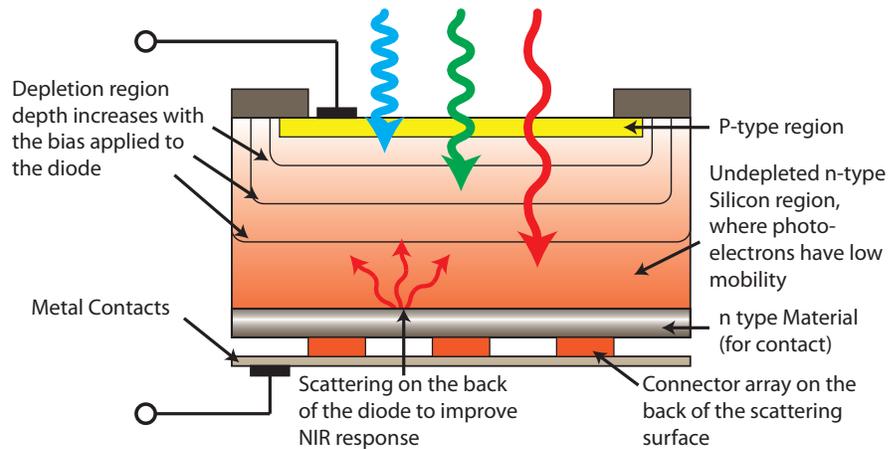


Figure 6.2: This diagram shows the some of the specialised features of the QP22-E type photodiodes. Illustrated are the variable absorption depths for different wavelengths of light. Though not to scale it is clear to see that short wavelength light (blue) is absorbed in a shallow depth of silicon near to the p-n junction and inside the depletion region. Shown is the connector array on the backside of the photodiode. These are thought to be the cause of the array of spots seen on all of the 2D scans of the photodiodes. Note also the back scattered light. This is done to raise the Amps/Watt rating of the photodiode by reflecting unabsorbed light back through the device.

note: this does not represent the actual QP22-E photodiode configuration and is for illustrative purposes only

6.2.3 Silicon photodiodes in near infrared spectrum

As illustrated in figure 6.2, longer wavelengths of light are absorbed much deeper. The absorption length of 1064 nm light in silicon is ~ 1 mm which is much deeper than the typical depth of the depletion region and in fact the overall thickness of the photodiode. As much as 60% of the near infrared (NIR) light incident on a Si photodiode doesn't generate a photo-current. The thickness of the QP22-E photodiodes substrate is only $\sim 380 \mu\text{m}$. Part of the reason for not making thicker photodiodes, is that for every 100 nm of photodiode thickness, there is an additional 1 ns of response delay. The production of photo-electrons outside the depletion region has two effects on the response of the photodiode. The first effect is a slower response of the photodiode to a signal. The second is the ability of the photo-electrons – outside of the depletion region – to diffuse laterally in the photodiode substrate leading to them potentially being readout in adjacent quadrants. This is called cross-talk and is to be avoided as it reduces a quadrant photodiode's ability to resolve the relative phase of signals in each quadrant. This is a particular problem for Dynamic Wavefront Sensing (DWS).

6.2.4 Enhancing Silicon responsivity to near infrared light

To help improve the performance of the QP22-E photodiodes at 1064 nm a number of changes to the silicon photodiodes are made during fabrication. The photodiodes are thicker to improve the chances of the photons being absorbed, which increases the responsivity of the photodiode. The photodiodes are more robust to be capable of having 150 V reverse bias applied. This is in fact their nominal operating point for 1064 nm light. By having such a

large reverse bias applied, the depth of the depletion region is vastly increased, which improves the speed at which the photo-electrons are readout and increases the likelihood of the photons being absorbed in the depletion region. It also reduces the chances of cross-talk between the quadrants as there is less chance of a photo-electron wandering unguided by an electric field into another quadrant. The final part of the customisation is a layer at the back of the photodiode which scatters the photons which are not absorbed back through the photodiode substrate. This increases the chances of the photons being absorbed by effectively increasing the thickness of the Silicon substrate. This scattering does however lead to what is termed ‘optical cross-talk’ where photons are scattered into adjacent quadrants before being absorbed. This optical cross-talk is unaffected by increasing the reverse bias.

6.3 Photodiode characterisations

The tests conducted on the photodiodes were: the spatial uniformities of responsivity and phase lag, and their transfer functions for signal amplitude and phase lag. The spatial uniformity tests are relevant to all LISA-type mission including LPF and eLISA. The signal frequency dependent tests are split between LPF requirements and eLISA/LISA-type missions. Not all of these tests were performed in the one place, a subset of them were performed over four different groups: The University of Birmingham, University of Glasgow, Albert Einstein Institute Hannover, and Astrium GmbH Friedrichshafen.

6.3.1 Spatial uniformity testing

The spatial uniformities of responsivity and phase lag were measured using beams of 1064 nm light. The beam was focused to a spot with a beam radii $\sim 100 \mu\text{m}$. The size of the beam sets the resolution of the spatial testing. Either the photodiode or the beam can be moved to produce a map of spot measurements covering the full active area of the photodiode. Some of the tests performed produced a full 2D map of a photodiode's active area; however for small beams the time taken to make a full 2D map is very much longer due to the squared dependence on the beam radius. The author and others at the University of Glasgow, performed 1D scans over key photodiodes regions with a beam of $20 \mu\text{m}$ in diameter. This allowed for high spatial resolution to be achieved over the region of interest, in a one hour time frame compared to a day or so for a full 2D map.

A limit of less than 3% variation in responsivity over the active area of the quadrants was the requirement for LPF. The origin of this requirement is not

fully understood by the author and the subsequent results in this chapter, combined with a simulation performed by E.Fitzsimons, showed it to be non-critical. As a reference, eLISA inherited a similar value on uniformity.

6.3.2 Frequency response testing

The QP22-E photodiode bandwidth was questioned following a series of tests showed higher than expected phase lag. These tests suggested lower than required bandwidth for the eLISA optical bench, where it was planned to have the QP22-E photodiodes. The photodiodes were tested for variation of output signal magnitude and phase to changing input signal frequency. These transfer functions were conducted over a range of heterodyne frequencies from, 1 kHz to ~ 50 MHz. This range fully covered both the heterodyne beat notes of the LPF optical setup of 1 kHz, and beyond the 20 MHz proposed for LISA-type missions.

The QP22-E photodiodes had been considered for eLISA, due to their quoted low capacitance of 3 pF and 12 ns rise time, which implied a cut off frequency safely above the 20 MHz requirement.

6.4 Spatial intensity response

6.4.1 Introduction

Spatial responsivity tests of the photodiodes were carried out by the University of Birmingham during the photodiode qualification. After uniformity issues with the Lewicki Silicon photodiodes were found, Glasgow University and then Astrium conducted scans too. Glasgow University produced high resolution 1D scans of intensity response; while Astrium GmbH performed both intensity response and phase lag 2D scans of the QP22-E photodiodes.

Spatial uniformity circuit diagram

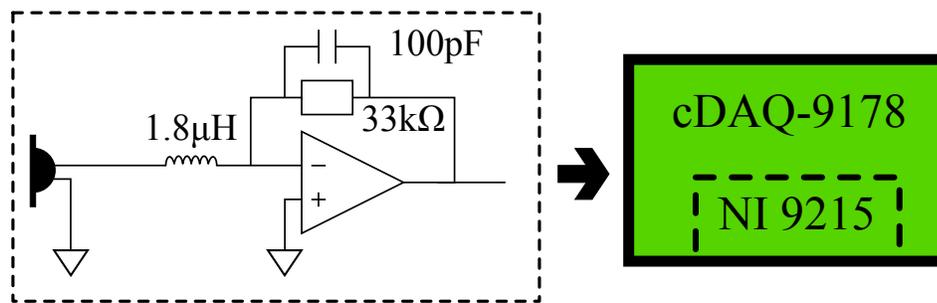


Figure 6.3: *This is the circuit diagram of the readout chain for each quadrant of the photodiodes tested for spatial intensity response. In this example, the photodiode has no reverse bias applied; however later scans did employ reverse bias. The front end electronics of this setup were made to be identical to the flight electronics of the LPF mission, and the signals are feed into a N.I. 9215 16-bit adc module, in a cDAQ-9178.*

Key factors to ensuring good results in these tests are: the repeatability of the beam size on the photodiodes; consistency of the dc light power level or depth of amplitude modulation; distance between measurement points; identical readout of quadrants; corrections for dark current and any off set voltages.

The beam powers used were approximately 20 mW or less in dc power, with tests done to check the linearity of the photodiode's response to beams of around this power level.

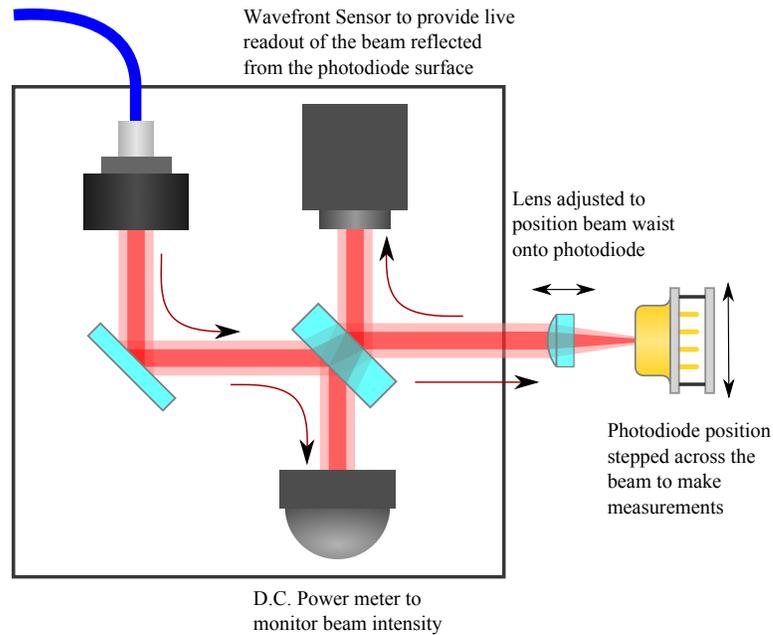


Figure 6.4: *Optical layout used to produce 1D scans at the Univeristy of Glasgow. The light was delivered by a fibre coupler and then aligned onto the middle on the photodiode through the beam splitter. Some of input light, reflected from the beam splitter, was directed to a power meter. Light passing through the beam splitter was focused down onto the surface of the photodiode. The light reflected from the surface of the photodiode, passed back through the lens; was re-collimated, and partially reflected to the wavefront sensor. The radius of curvature of the light reaching the wavefront sensor was therefore a measure of the spacing between the lens and the photodiode surface. For a collimated beam from the fibre coupler the correct spacing of lens-to-photodiode would yield a focused spot on the photodiode surface and a well collimated beam at the wavefront sensor.*

To ensure sufficiently detailed measurements of the gap region between quad-

rants the beam probe must be smaller than the gap. The gaps of the QP22-E and the OSI InGaAs photodiodes are $70\ \mu\text{m}$ and $45\ \mu\text{m}$ respectively and a beam of $20\ \mu\text{m}$ diameter was used. Smaller beams require more points to cover the full photodiodes so it was seen as a good idea to do 1D high res scanning of the photodiodes, rather than 2D maps. Horizontal and vertical scans were performed to check the consistency of the gap size; although they revealed that the gap thickness was not equal, and varied from place to place.

6.4.2 Birmingham University scans

Work done at the University of Birmingham as part of the photodiode qualification procedure, confirmed that the InGaAs photodiodes met the uniformity requirements. The test measured the magnitude response of the photodiodes to 1064 nm light over their active area and showed that they were uniform to better than $\pm 3\%$. The responsivity of the photodiodes was measured using a 175 μm diameter beam with 100 μm steps between measurements. The photodiodes were probed with 1064 nm light from a power stabilised laser. Each sample point on the photodiode was the average of a 100 measurements, taken at 1 kHz. These points were then plotted on a graph, showing the four separate quadrants (figure 6.5) [51].

The size of the beam was much larger than that which could be used to measure the profile of the photodiode's gap region and the step size was larger than the gap so the resolution was poor.

University of Birmingham QP22-E photodiode scans

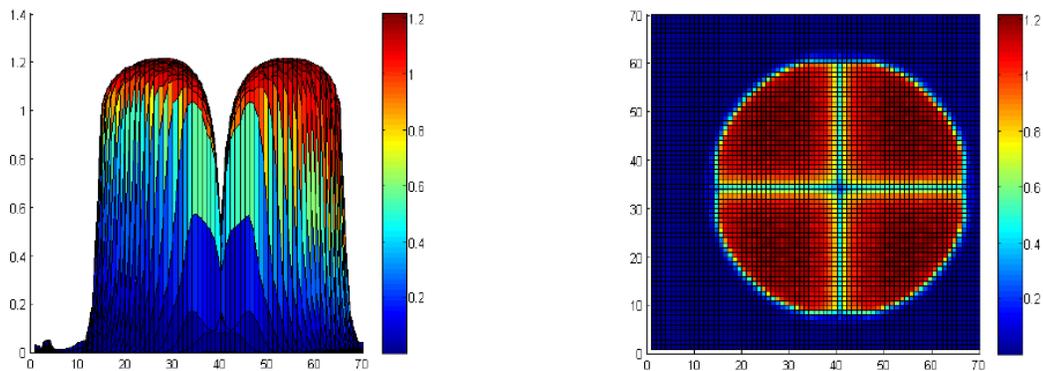


Figure 6.5: Shown is the University of Birmingham scans of a QP22-E photodiode showing large scale non-uniformity of the response.

The results of the scans can be seen in figure 6.5. There is a large scale non-

uniformity clearly visible in the quadrant response which should ideally be flat over the active area. There was overlap at the gap regions due to the large beam size compared to the gap of $70\ \mu\text{m}$. This meant that the gap is undefined in the scans. There is also a clear spot pattern in the quadrants. This was suspected of being the connector array illustrated in figure 6.2.

6.4.3 Glasgow University scans

Experimental setup

The beam size of the setup (figure 6.4) was measured to be $20\ \mu\text{m}$ in diameter and verification of this can be seen in some scans where narrow features are present and from figure 6.6. With a small beam like this it is important to maintain its focus position onto the photodiodes surface. Failure to monitor the spot size could lead to larger than expected beams scanning the photodiode.

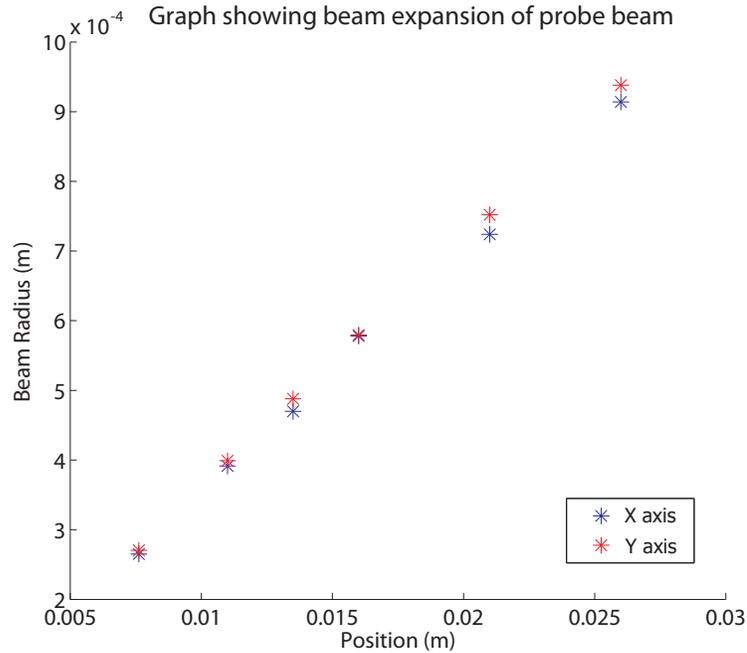


Figure 6.6: The graph shown here is the plots of both the X and Y axis of the beam used to probe the photodiode. A straight line fitting of these points gave an ω_o of $19.5 \pm 2\ \mu\text{m}$.

Beam size

The spot size or focus on the photodiode was monitored using light reflected from photodiode surface and measured by a wavefront sensor (see figure 6.4). To ensure the spot was focused to better than the Rayleigh range onto the surface of the photodiode, the radius of curvature at the wavefront sensor had to be kept between -1000 mm and 500 mm. This translated into positioning of the photodiode to a few microns with respect to the lens used to focus the light. During the testing, a collimated beam of < 3 m radius of curvature, ensured that the desired spot size on the the photodiode was maintained.

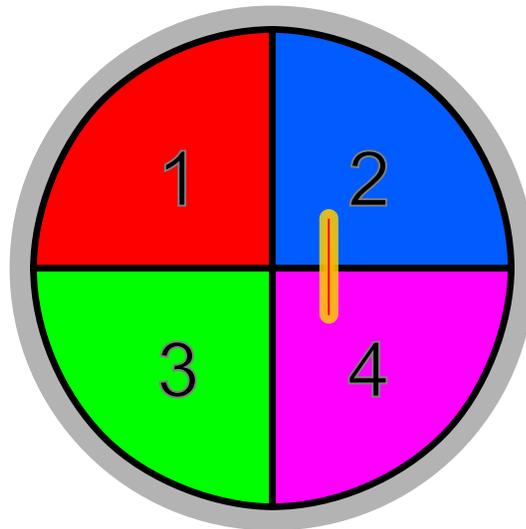


Figure 6.7: The figure shows the region over which a typical 1-D scan of a QP22-E photodiode was conducted. The scan runs parallel to the vertical gap of the photodiode, but at a distance from it to ensure there is no significant signal from quadrants 1 and 3. The diagram is to scale: the black lines are the thickness of the gap, the quadrants are the correct area, and the thickness of the red line in the orange strip, is the width of the beam used to probe the photodiode.

The photodiode under test and the lens used to focus the light were both mounted on Hexapods. These enabled control of angles to a milli-degree, and positional control at the sub-micron level. The positioning was repeatable to sub-micron and realignment was only required due to changes of photodiode or other components in the beam path moving.

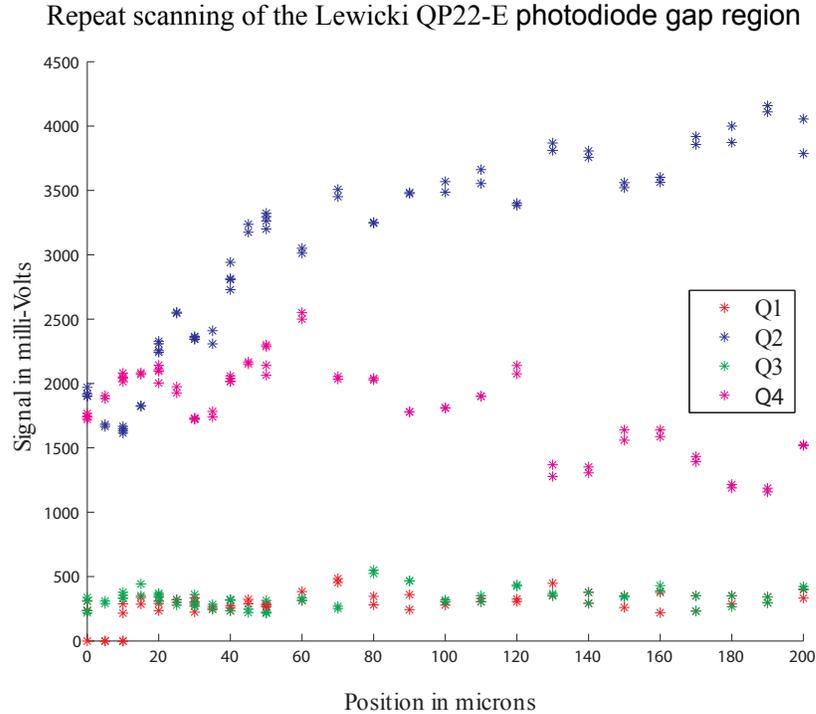


Figure 6.8: *Graph of repeatability of the measurements. Taken by rescanning the same path of a QP22-E photodiode with a 5 – 10 minute break between scans. Each repeated data point, shows variation at the $\sim 3\%$ level which is about the same as the error in the measurements due to the correction for the power variation. The scan starts each time in the centre of the gap between quadrants one and two. And shows an arbitrary, repeatable shape to the responsivity of the photodiode. The irregularities of the photodiode, are much larger than the $\pm 3\%$ uniformity of response requirement which is needed over the active area of each quadrant.*

Measurement repeatability

Repeat measurements of the same photodiode were made over the course of several minutes to show the long term repeatability of the set up (see figure 6.8). This does not prove that it was repeatable day-to-day, but ensured that over the course of a single photodiode scan, any variation was due to photodiode features and not noise in the experiment. There can also be some level of confidence that this was not a profound effect or one which changed day-to-day, due to the consistency of the results as a whole. These measurements were corrected for beam power variation, using the power monitor values taken at the time of the measurement.

6.4.4 Results

The plots of an OSI InGaAs photodiode and a Lewicki and Silicon Sensor photodiodes can be seen on figures (6.9, 6.10, 6.11, 6.12). Note the striking difference in uniformity of response. The InGaAs photodiode shows well defined, flat topped quadrants and a clear gap region (see figure 6.9). The gap is a steep sided, fall off of response, as the beam moves from one quadrant to another. There is minimal cross-talk, and the slight rise of signal in other quadrants – while the beam is in the gap – which could be from scattered light due to the, possibly rough nature, of the gap's surface.

The QP22-E photodiodes show a sloped response with a dip towards the centre of the photodiode and a roll off towards the edges (see figure 6.10). The drop in signal starts before the beam gets close to the edge of the quadrant; even after moving across the gap into a new quadrant signals continued to be recorded in the previous quadrant for some $\sim 300 \mu\text{m}$. The gap region

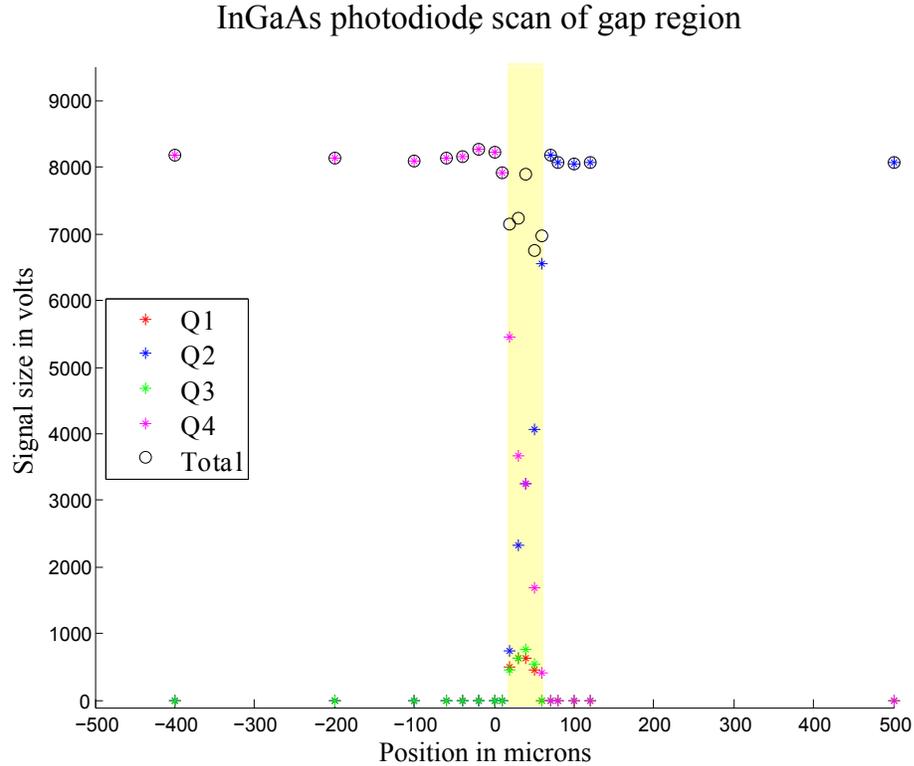


Figure 6.9: Graph showing the spacial response of a OSI InGaAs photodiode. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the sharp transition between the quadrants. The data shows variation at the $\sim 3\%$ level which is about the same as the error in the measurements due to the correction for the power variation.

is almost insignificant in its effect on the overall response of the photodiode. There is a drop in the overall response, when compared to an area, close to the middle of a quadrant; however each channel shows just as much signal, as it does cross-talk, in the quadrant areas, close to the gap.

The application of 9 V reverse bias, to the QP22-E or Lewicki photodiodes, improved the shape of the response. The fall off as the beam approached the gap got better, and the signals dropped away faster; however cross-talk, could still be seen $\sim 250 \mu\text{m}$ into the adjacent quadrant.

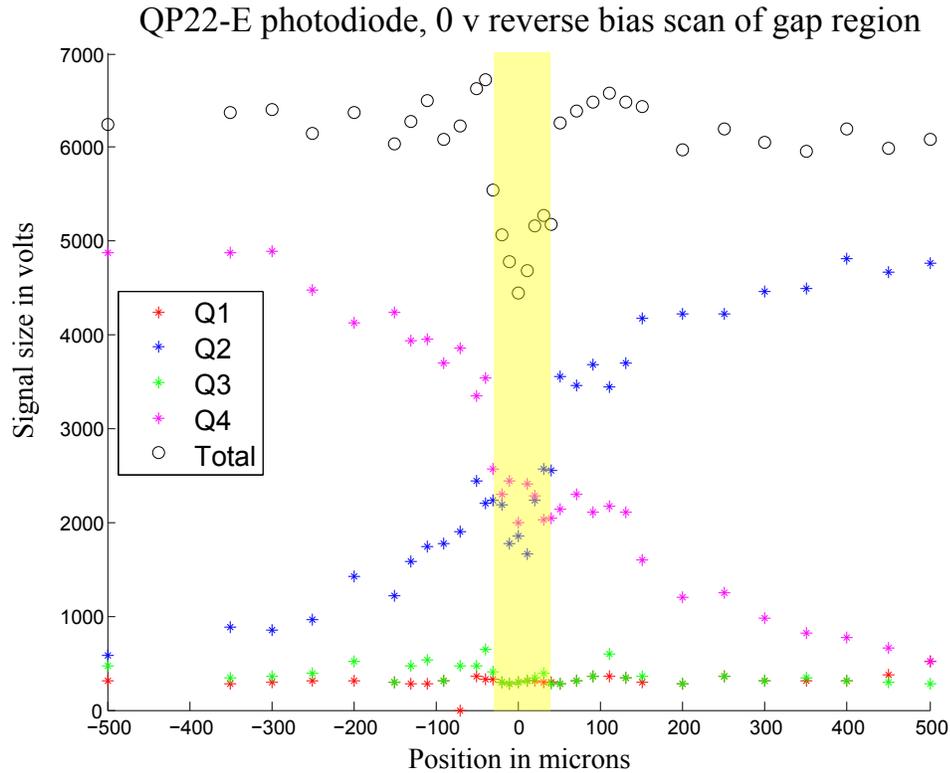


Figure 6.10: Graph showing the spacial response of a Lewicki QP22-E Photodiode without reverse bias. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.

Figure 6.12 shows the improvements that 100 V of reverse bias were able to achieve. The fall off on approach to the gap, is now limited to less than $100 \mu\text{m}$ and is a steep fall off. The cross-talk; however is still visible – to a lower level – $\sim 200 \mu\text{m}$ into the next quadrant.

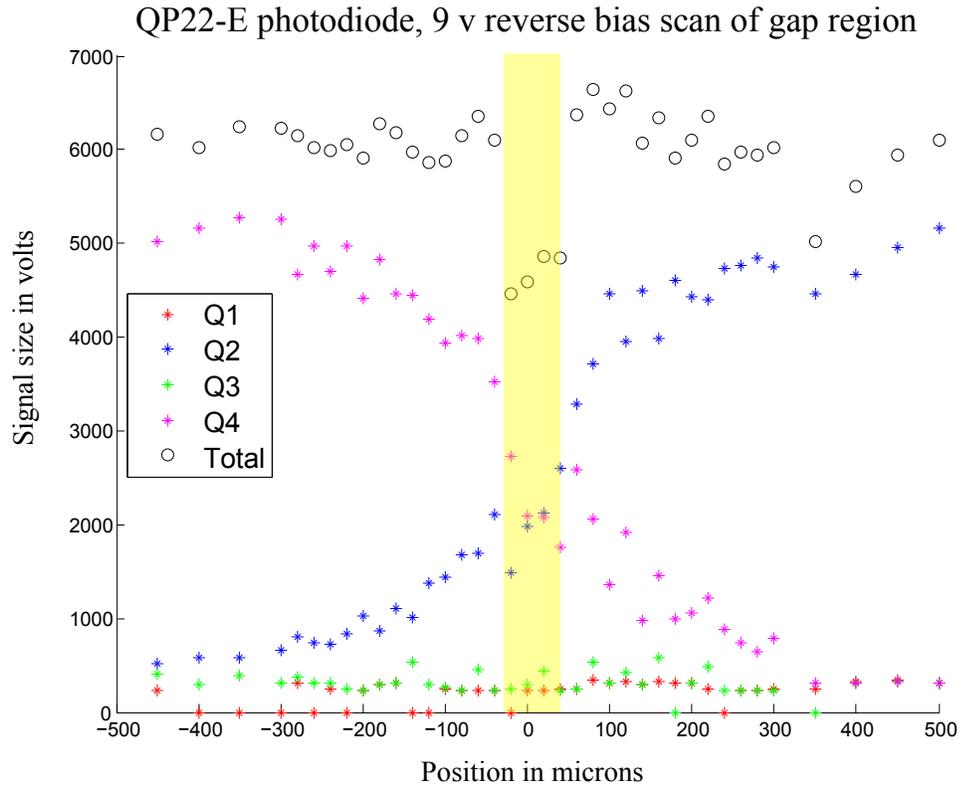


Figure 6.11: Graph showing the spacial response of a Lewicki QP22-E Photodiode with a reverse bias of 9 V. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.

Astrium scan of spatial phase response

A 2D scan of the phase lag response of a QP22-E photodiode showed decreasing phase lag as you move radially away from the centre of the photodiode (figure 6.13). The phase lag was greatest towards the gap and there was also a matrix of dots which appeared in a regular pattern over the active area of the photodiode. These dots although not easily found on any of the 1-D scans, showed clearly on the intensity response maps too. These are thought to be

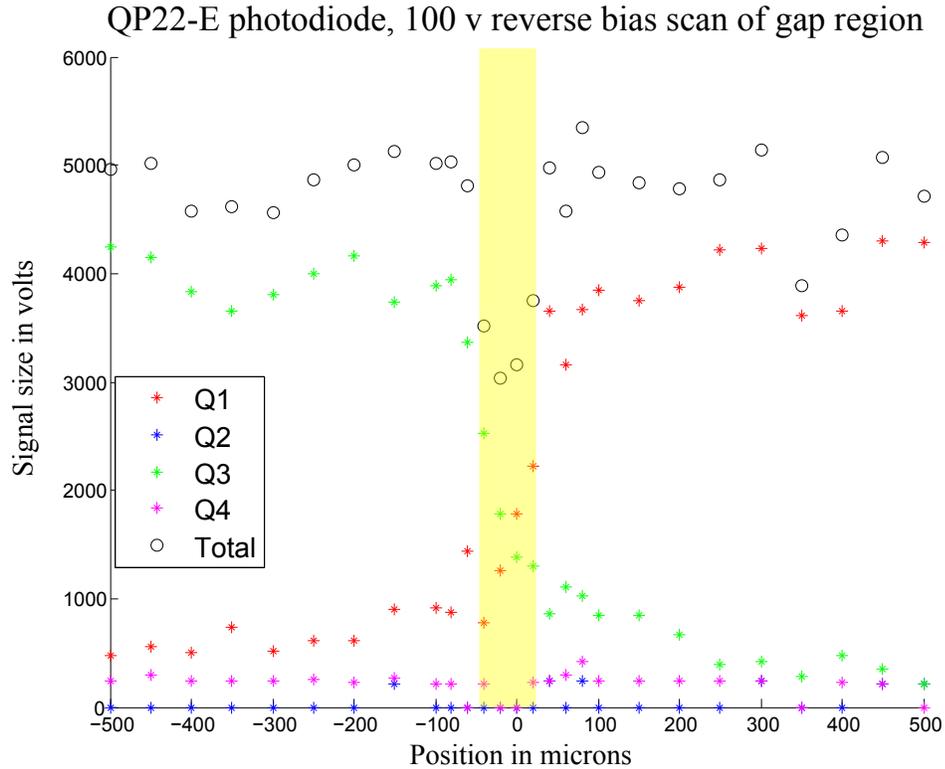


Figure 6.12: *Graph showing the spacial response of a Lewicki QP22-E Photodiode with a reverse bias of 100 V. The scan done at Glasgow University detailed the gap region of the photodiode, and shows the transition between the quadrants. The data shows variation greater than $\sim 3\%$. This is above the error in the measurements due to the correction for the power variation.*

features from the connector array on the back of the photodiode substrate, seen in the responsivity maps produced by the University of Birmingham.

6.4.5 Conclusion

The cross-talk of the photodiodes is caused by the incoming photons penetrating deeply into the silicon photodiode (see figure 6.2). This means that many of them will generate photo-electrons outside of the depletion region where there

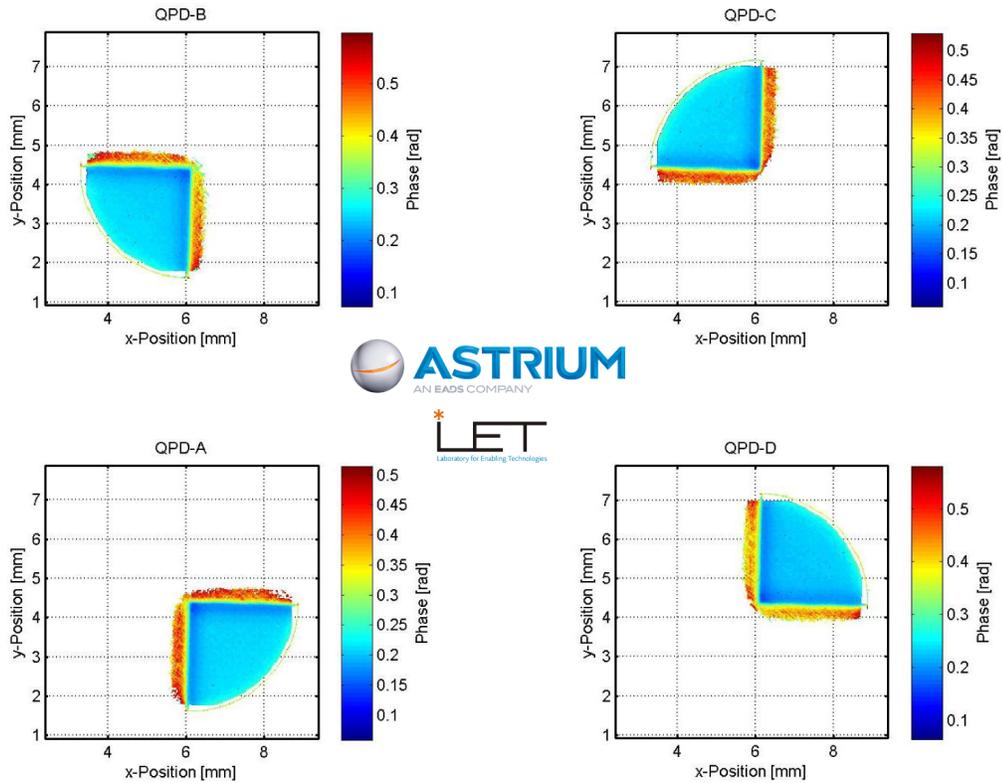


Figure 6.13: Graph produced by Astrium GmbH Friedrichshafen [52], measuring the spatial phase response of the QP22-E photodiodes. The measurements were done using 1064 nm light on an unbiased photodiode.

is no sizable electric field to move them towards the p-n junction of the photodiode. This causes the travel time to be longer and allows the photo-electrons the ability to diffuse laterally in the photodiode leading to them being readout in adjacent quadrants. Applying more reverse bias increases the depth of the depletion region, which helps limit the cross-talk. The results of these scans of the QP22-E photodiodes with various reverse bias voltages applied show that cross-talk is significantly improved by higher reverse bias. The photodiodes are designed to be operated with high reverse bias. The specification table for the QP22-E photodiodes shows clearly that many of the photodiode tests were

performed at 150 volts reverse bias. The residual cross-talk of the photodiodes at 150 V reverse bias is due to what is termed optical cross-talk. This is caused by photons scattering off a special surface at the back of the photodiode. These photons pass through the photodiode substrate, without being absorbed, and are reflected back into the photodiode substrate by a scattering surface, see figure 6.2. By scattering the photons back through the substrate their chances of being absorbed are increased; however they can be scattered into the adjacent quadrants. This effect is independent of the reverse bias applied to the photodiode.

6.5 Frequency response

6.5.1 Introduction

Following the finding of the phase lag across the photodiode quadrants (figure 6.13) it was decided to test the bandwidth of the silicon photodiodes. The size of the phase lag in those scans gave cause for concern over the bandwidth of the QP22-E photodiodes. The goal was to analyse the Silicon photodiodes response to higher frequency signals. The method was to produce transfer functions of the photodiodes response in terms of both signal magnitude and phase, for different values of the frequency of the input signal and the reverse bias applied to the photodiode. The signal magnitude and phase would be measured relative to a reference photodiode. As a reference to measure against the University of Glasgow used a single element InGaAs photodiode (C30619 PerkinElmre).

Frequency testing circuit diagram

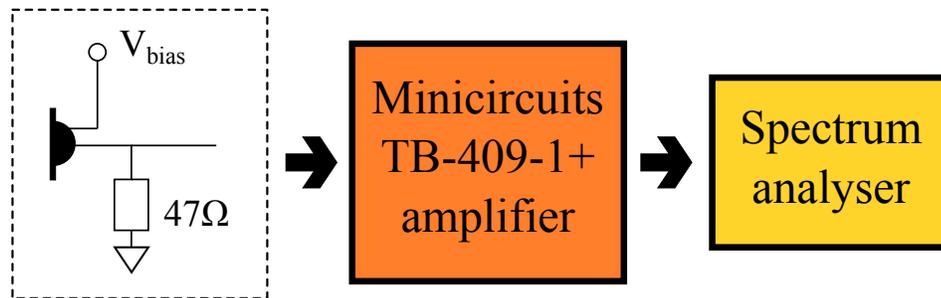


Figure 6.14: The circuit diagram of the readout chain for a single photodiode quadrant or photodiode, during the tested for frequency response. The non-recorded quadrants, were connected to ground through 47Ω . The minicircuits TB-409-1+ contains a GALI 1+, DC to 8 GHz amplifier. The output of which was fed into one of two Spectrum Analysers, depending on the experiment.

6.5.2 Low frequency transfer functions

The first tests performed on the QP22-E photodiodes was comparing their signal magnitude response compared with an InGaAs photodiode. This first provided information on the use of the QP22-E photodiodes on the LPF optical bench, for which the QP22-E photodiodes had been procured. The tests used the setup seen in figure 6.15 which was a heterodyne beat note produced by twin AOMs. The contrast of the interferometer setup was $\sim 70\%$, with mean power levels of ~ 10 mW. The power levels were not constant over the full range of frequencies. The frequency range explored with this setup was limited to 100 kHz by the FFT Analyser (Stanford research systems model SR780 signal analyzer).

The InGaAs photodiode was quoted as having a bandwidth of 350 MHz [53], allowing it to act as a reference. By making the measurements of QP22-E photodiodes relative to the InGaAs photodiode, a measure of the frequency response of the QP22-E photodiode was produced which was independent of the changing light powers in the interferometer.

The results of this experiment showed a marked drop in the signal magnitude response as you went up in frequency, figure 6.16. Without any reverse bias applied the signals were down by a factor of two at ~ 25 kHz when compared to 1 – 2 kHz signal. A reverse bias of 9 V improved the response a little but it was still a factor of two down by ~ 50 kHz. Looking now to the 100 V reverse bias plot, it can be seen that it is flat all the way out to 100 kHz. To further understand why we look towards the phase lag plots.

The graph of phase lag showed similar results to that of the magnitude response, seen in figure 6.17. The graph here is the relative phases of the signals coming from the two photodiodes. Again the plot without reverse bias shows a

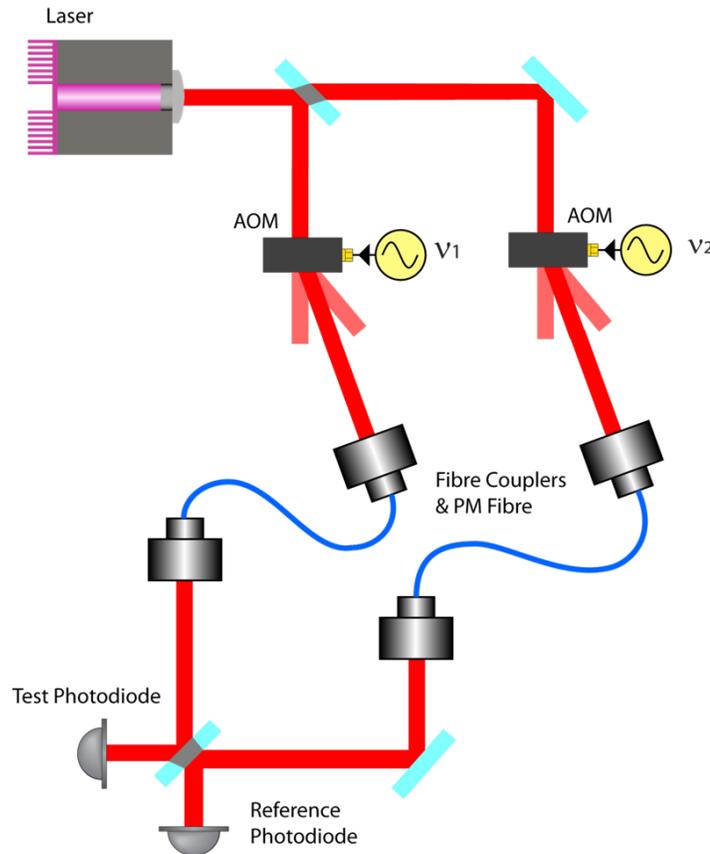


Figure 6.15: The figure shows the optical setup used to test the QP22-E photodiodes over the frequency range of 1 – 100 kHz. The light from a laser was split into two AOMs driven at 80 MHz. The difference in heterodyne beat note was controlled by changing $\nu 2$.

fast drop off in response which starts at a lower frequency than 1 kHz. As can be seen it reaches almost 45 degrees behind in the mid to high tens of kHz. The 9 V reverse bias case shows marked improvement over the photodiode without bias; however it is still a long way short of what is required. Note also the hint of an up turn from the 30 kHz onwards. This was a puzzling artifact which is seen in all later transfer functions performed at the University of Glasgow. The plots of the 100 V reverse bias is again a flat line out to 100 kHz.

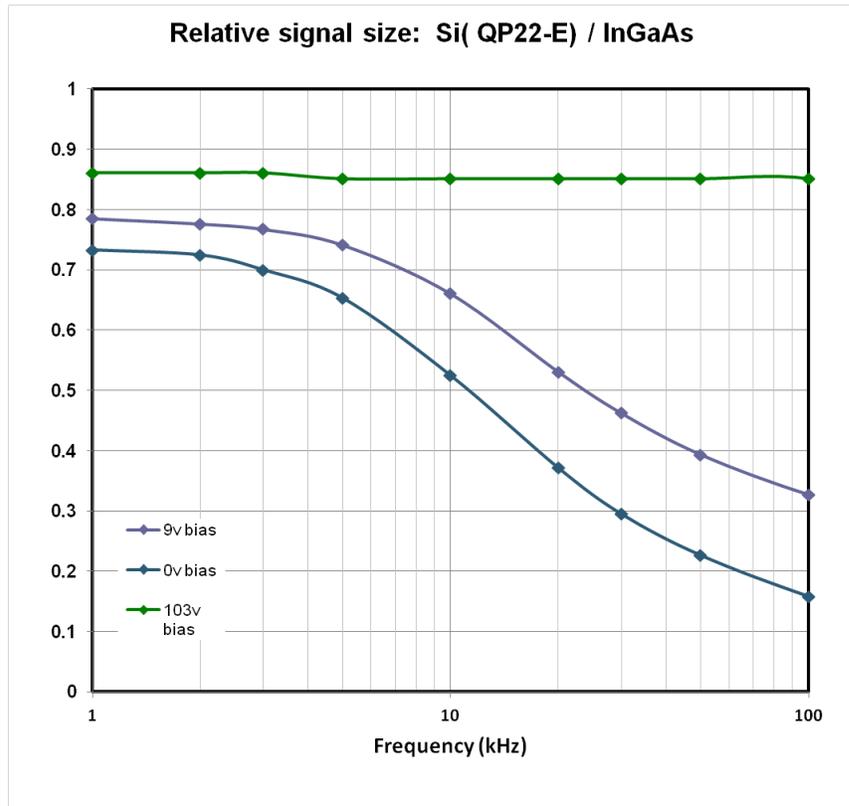


Figure 6.16: The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at three different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the three different reverse bias voltages.

Both of these results are consistent with the ideas laid out for how photoelectrons produced outside of the depletion region act (see section 6.2).

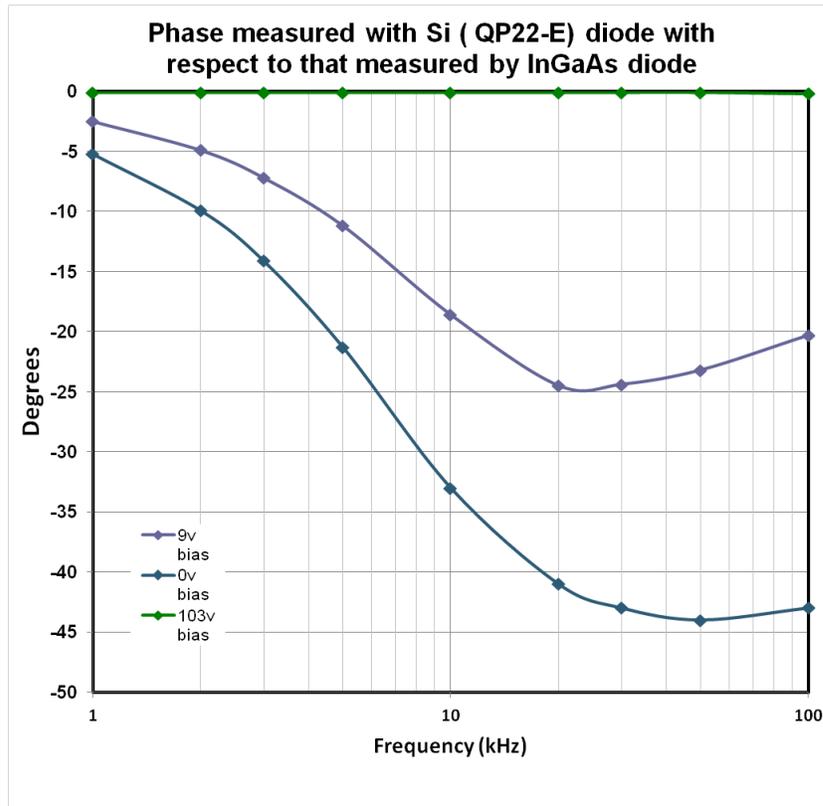


Figure 6.17: The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at three different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the three different reverse bias voltages.

6.5.3 High frequency transfer functions

In order to test the QP22-E photodiodes over both the LISA pathfinder and LISA frequency band, a pair of offset locked lasers were used. This set up provided a stable heterodyne beat note between ~ 1 kHz and ~ 100 MHz. The offset locking frequency was supplied by an FFT analyser, allowing the heterodyne frequency to be swept over the range of frequencies. The analyser changed

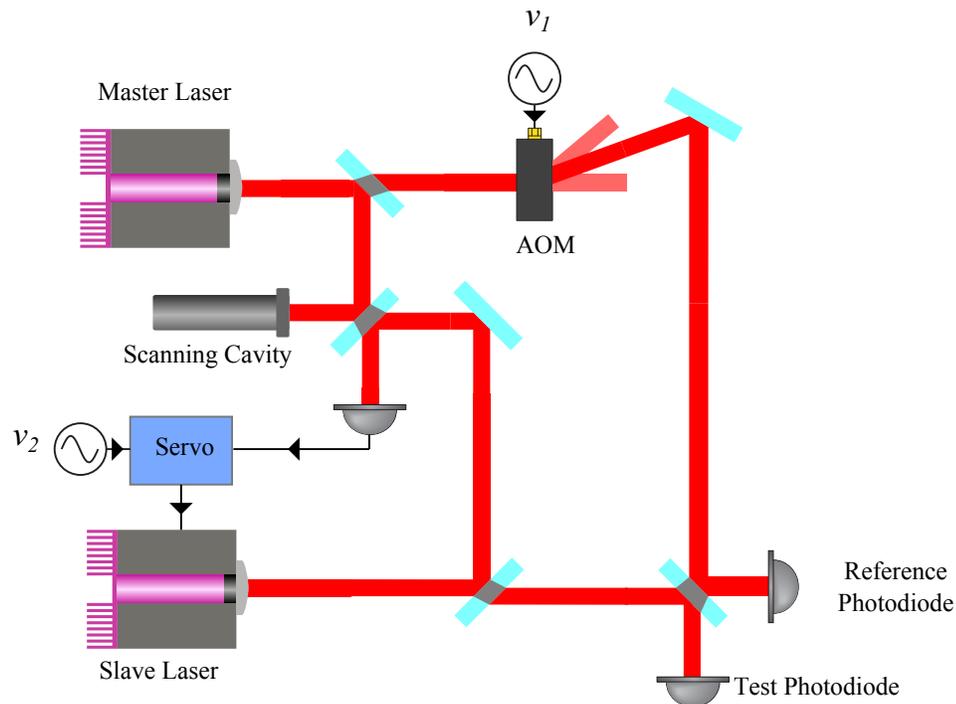


Figure 6.18: The figure shows the optical setup used to test the QP22-E photodiodes over the frequency range of ~ 1 kHz to ~ 100 MHz. The light from the lasers was passed through faraday isolators before being directed into a scanning cavity and a high bandwidth InGaAs photodiode. This photodiode was used to detect the beatnote between the two lasers. The frequency difference between the lasers was set to ν_2 , and was maintained via a servo which used feed back from the photodiode to drive the slave laser to follow the the master at the correct offset. The heterodyne beatnote at the test photodiodes was ν_2 minus ν_1 where ν_1 is 80 MHz. The heterodyne frequency was changed by adjustment of ν_2 . The reference photodiode was again smaller in area than the QP22-E photodiode, so all results are again relative and not absolute in there magnitude. The measured phase difference was unaffected by this.

the heterodyne frequency by changing the offset locking point of the lasers ν_2 . The resulting signals from both the InGaAs and QP22-E photodiodes, were input to the FFT analyser. The analyser could only go up to 20 MHz so all

data taken beyond this frequency were done with spot measurements.

The beatnote between the lasers was unstable at the low frequency end of the test. At a beatnote frequency on the test photodiodes of ~ 1 kHz, the signal phase was very unstable; however, because we had a perfectly good measurement set for these lower frequencies this was not required.

The data appeared to show a minor dip below 1 MHz when comparing any QP22-E with an InGaAs photodiode. The reason for this dip was not fully understood. A series of cross checks to the signal chains, changes of light levels, and even different photodiodes all showed a dip. Over the course of these checks and tests we were able to show that it was: independent of what readout chain was used, that the light levels did not change its size or depth, it had slight variation over the different QP22-E photodiodes tested and that the dip was not seen when comparing two InGaAs photodiodes.

Work done by our colleges at AEI didn't show signs of this dip; however their setup was not identical. Their signal was produced by amplitude modulation of a laser diode, where the built in reference photodiode was used to compare the measured signal magnitudes and phase from the tested QP22-E photodiodes. Their amplitude modulated signals were small compared to the dc light levels on the photodiodes which we assumed to be the only difference. Comparatively the two experiments had an approximately equal ac signal size but our setup had much lower dc light levels.

The OSI InGaAs photodiode and the Lewicki photodiode were both untested due to their long pre-assembled cables. These prohibited the use of intimate readout electronics, and made signal readout chain length matching – to avoid signal lag from unmatched signal readout chains at high frequencies – very difficult.

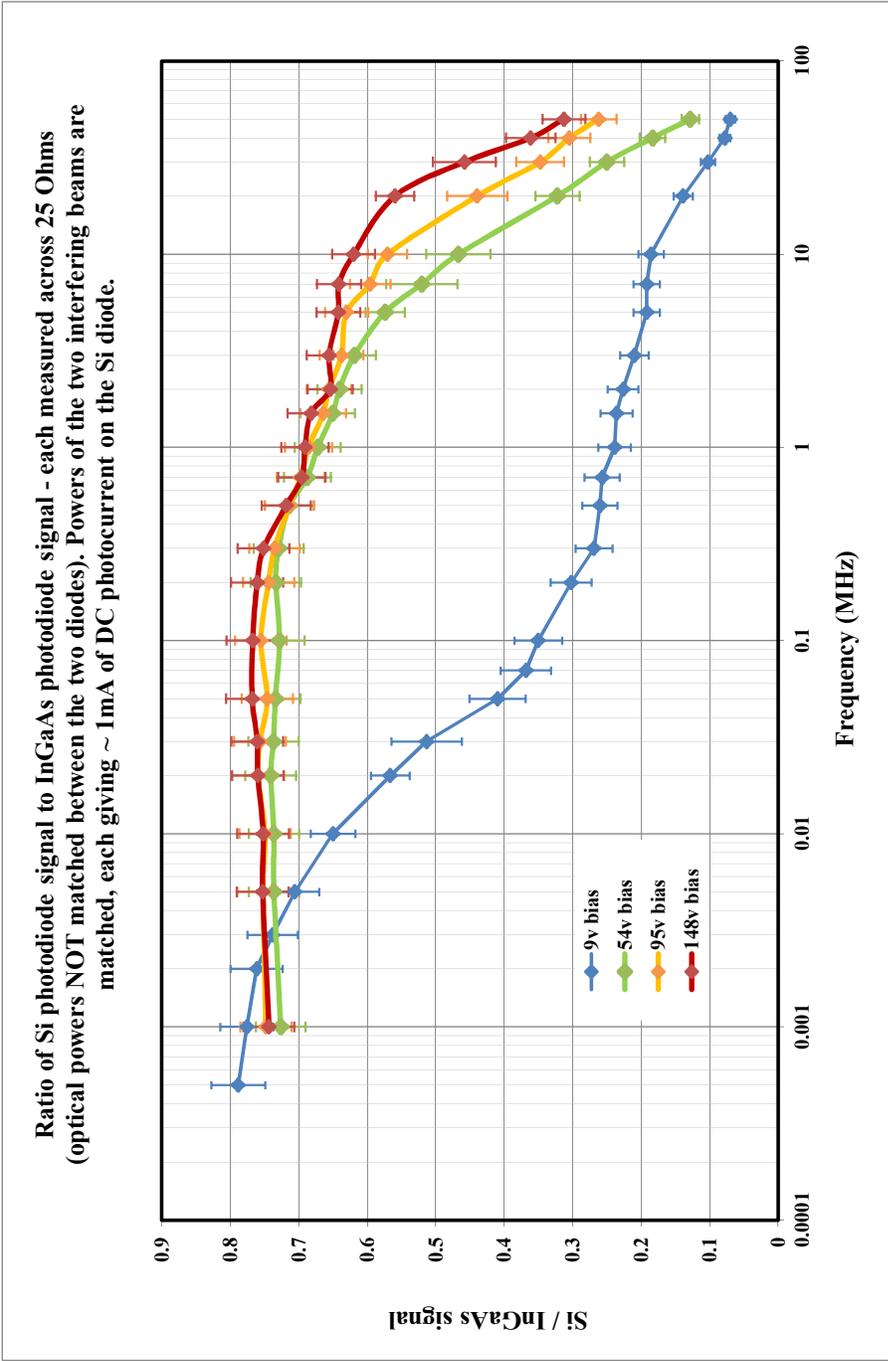


Figure 6.19: The graph shows the ability of the QP22-E to transfer an input light signal (the heterodyne beatnote) into an output voltage at four different reverse bias voltages. The y-axis is the ratio of signal amplitudes readout from the InGaAs photodiode vs the signal amplitudes readout from the QP22-E. The over and under filling of the InGaAs photodiode causes the different dc levels that appear for the four different reverse bias voltages.

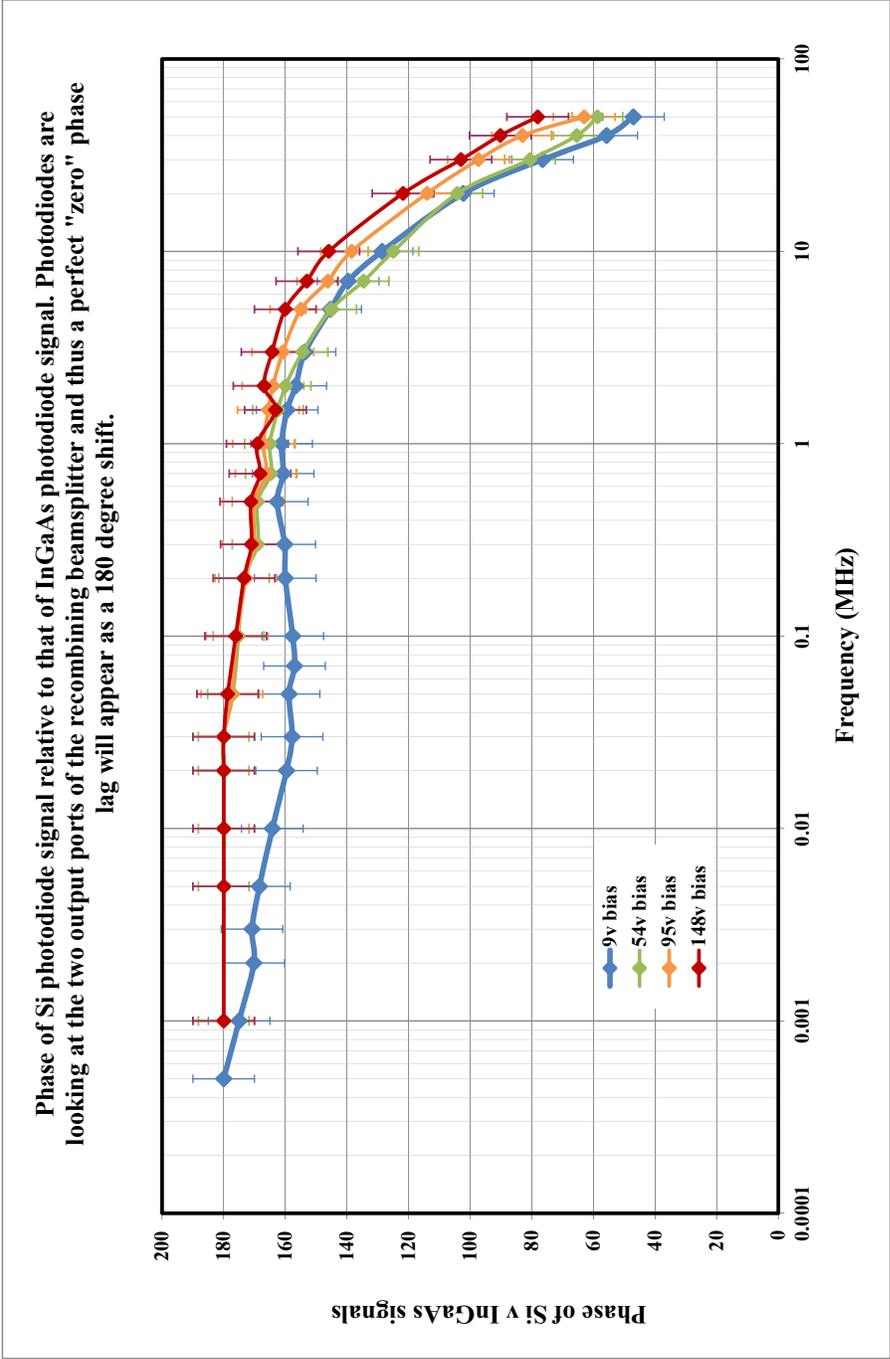


Figure 6.20: The graph shows the phase lag of the QP22-E compared to an InGaAs photodiode. The QP22-E was measured at four different reverse bias voltages. The y-axis phase of the readouts from the two photodiodes.

6.5.4 Conclusions

The conclusion of this work was that the QP22-E photodiodes were unsuitable for use with LISA-type missions. The reason was their poor magnitude response and large phase lag (figures 6.19, 6.20) at the max heterodyne beatnote frequency which was planned to be at 20 MHz. Although there was the possible fix of applying a large 150 V reverse bias to help solve these issues, such large reverse bias on photodiodes would produce significant power dissipation. This could affect the required temperature stability on the optical bench. In order to maintain the required thermal stability, the photodiode power would need to be stabilised.

6.6 Pulse response times of photodiodes

6.6.1 Introduction

Another area of photodiode input to the LPF mission is the power stabilisation servo that monitors dc power levels on the optical bench. This servo uses the light power levels as measured by the photodiodes and feeds back to AOM drivers, correcting for beam power variations. This servo requires a bandwidth of 100 kHz from the photodiodes, which is higher than the required bandwidth for the science output of the interferometers. From the previous work done to analyse the frequency response of the QP22-E photodiodes, it was known that without a reverse bias they did not have the required bandwidth. An experiment was undertaken to look for the response of the silicon photodiodes to light pulses to further characterise this time response.

A follow up on frequency domain analysis of the QP22-E photodiodes for phase response in the time domain was chosen. The experiment used an AM beam which required rapid switching between on/off to produce a step change in light power. The time taken between the on and off states needed to be less than the ‘rise time’ of the QP22-E ($\sim 12 \mu\text{s}$). By comparing the response of the QP22-E photodiode to the step change, with that of the InGaAs photodiode (C30619), an understanding of the profile of the response could be gained. The bias applied to the photodiode during this investigation was only that which could be supplied by the LPF electronics, which is limited to 5 volts. The time constants of these two effects were what we wanted to determine, such that they could be used as inputs to model the stability of the power monitor servo.

6.6.2 Theory

The response of a typical photodiode is separated into two different regimes. These regimes are the slow and prompt response. A prompt response is one in which the time between a photon arriving at the photodiode and the resulting photo-electrons moving across the p-n junction is short. This prompt response comes from the production of photo-electrons in a depleted region close to the p-n junction. An InGaAs photodiode working in the NIR region has a response to a pulse or step change in light power which could be said to be prompt. This is due to the generation of photo-electrons close to the p-n junction.

In the QP22-E photodiodes the generation of many of the photo-electrons is outside of the depletion region, far from the p-n junction. This leads to longer times between photon arrival and photo-electron read out. The so called slow response.

6.6.3 Experimental setup

To ensure that the test was as representative as possible, the dc light powers of the beams were set to ~ 1 mW which is close to the expected 0.8 mW. The light used was amplitude modulated light from an A&A AOM. The drive to the AOM was via a Mini-circuits ZLW -6+ mixer to produce fast switching. Using the mixer, switching times of ~ 10 ns were achieved. The modulation of the light was applied by a square wave with its amplitude and offset voltage determined the size of the light power and the depth of the modulation.

The reference photodiode used in this setup was a PerkinElmer C30619G In-GaAs single of diameter 0.5 mm which had a lens in front of it to help ensure all of the light was incident on the photodiode and not lost. This photodiode

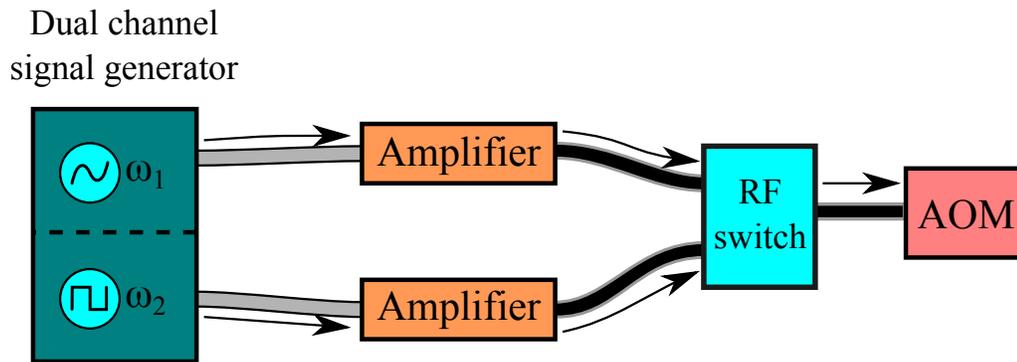


Figure 6.21: *Light step generation setup. The RF switch inputs were: ω_1 an 80 MHz signal amplitude modulated at 1 kHz, and ω_2 a square wave modulation signal. The output of which drove an AOM which produced a beam of light, amplitude modulated at 1 kHz, and periodically switching on and off with frequency ω_2 .*

had been used before to demonstrate the ability of the QP22-E with 100 V bias to be used in the multi-MHz region without any problems.

Pulse testing circuit diagram

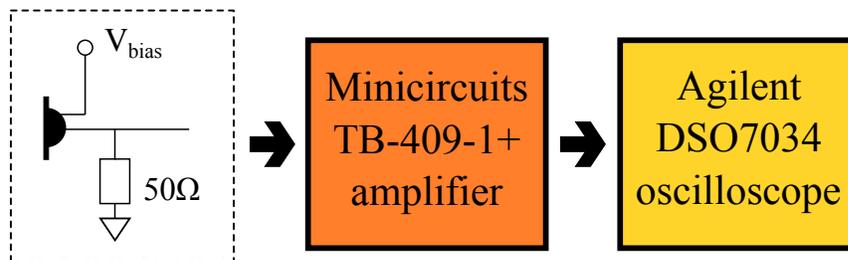


Figure 6.22: *Shown is the readout chain for the photodiodes used in the pulse response experiment. The amplifier was rated as DC to 8 GHz while the oscilloscope had a 350 MHz bandwidth.*

Both the reference InGaAs photodiode and the QP22-E were biased with 5 V. The photodiodes were connected to identical read out chains which were as simple and intimate as possible, see figure 6.22. The unrecorded signals

from three of the quadrants, were readout across $47\ \Omega$ to ground. Note that the oscilloscope channel was set to $50\ \Omega$, while there was an additional $50\ \Omega$ to ground connection. This meant the signals were effectively readout across $25\ \Omega$. This was due to a ‘T’-piece being used between the photodiode and the Minicircuits amplifier. This did improve the bandwidth without causing issues to the signal size as there was ample light power available.

As it turned out the propagation of the acoustic waves through the crystal within the AOM was the source of some delay in the system response and also hide the prompt response of the diode. The waves travel at $\sim 4\ \text{km/s}$ through the crystal and the crystal is has $\sim 4\ \text{mm}$ of crystal before the beam entrance, followed by the $1\ \text{mm}$ where the beam is present. This $4\ \text{mm}$ of crystal showed up as a lag between input signal and any electrical signal from either photodiode. The $750\ \text{ns}$ delay was followed by a sharp rise in signal levels in both diodes for $\sim 150\ \text{ns}$ which has been deduced to be the effect of the first acoustic waves passing through the beam in the crystal.

6.6.4 Prompt response

All photodiodes have a prompt and slow response. The prompt response is the readout of photo-electrons in the depletion region, while the slow response is the readout of the photo-electrons generated in the non-depleted region

Shown in the figure 6.23 is the prompt response of the the QP22-E photodiode and the InGaAs photodiode. There it is a delay between the RF switching on and the photodiodes responding. The reason for the delay is the propagation time of the acoustic waves in the AOM crystal. This is why the InGaAs response has the shape that it has. It is a measure of the fraction of the power of the light beam in the AOM which had received the signal. It is the integral



Figure 6.23: Shown is the fast and slow response of both an InGaAs photodiode and a QP22-E photodiode. The InGaAs response is limited by the time taken for the acoustic wave to travel across the beam in the crystal. The prompt response of both is evident, with the QP22-E photodiode showing a slow rise over a $10\text{s } \mu\text{s}$.

from one side to the other of the power of the Gaussian beam, which has the shape of the error function. The same shape is seen in the QP22-E photodiode

response only it has a 'slow' response part which tails off from it. The 200 ns taken by the waves crossing the beam hides the 'true' prompt response of the photodiodes.

6.6.5 Slow response

In terms of its amplitude, this limited prompt response is seen to be anywhere between a half to a third of the full beam power step size. This is an instantaneous response compared to the following slow response of the photodiode as it rises with a time constant of $12\ \mu\text{s}$ (see figure 6.23). It is clear to see the sharp almost immediate response of the InGaAs photodiode and the contrasting sharp then slow response of the QP22-E. This step change in light power was done with a 95% amplitude modulation to show the effect clearly. The timescale for the signal recovery is about $36\ \mu\text{s}$.

6.6.6 Conclusion

The results of the stepped light powers, showed that the QP22-E photodiodes has a prompt response, followed by a slower response which has a time constant of $12\ \mu\text{s}$. Around 2/5 of the QP22-E response signal is prompt like that of the InGaAs photodiode. The slow response of the photodiode would greatly impede the stability of the power monitor servo, and it was decided post analysis to use InGaAs photodiodes.

6.7 Summary

The results of these tests conducted on the QP22-E photodiodes has profiled their response to 1064 nm light. These results show that compared with an InGaAs photodiode, these modified Silicon photodiodes are not suitable for many of the purposes a LISA-type mission would require them to fill. The photodiodes do perform at a standard which is border line usable when they are applied with a reverse bias of more than 100 V. Without this reverse bias the photodiodes have vast overlap between quadrants, and act as a low pass filter to signals as low as 1 kHz. Their uniformity of response over the active area of the photodiode is also very poor. It has been measured that the uniformity is at best $\pm 10\%$ of the average and at worst can be 50% lower at the gap region than at the middle of a quadrant.

The reason for this is partly that in order to improve the Amps per Watt of these photodiodes beyond that of a normal silicon photodiode. They have caused the spatial response of the photodiodes to become very irregular and scattered. Through the use of the scattering to effectively double pass the light through the photodiode, the spatial response is much worse than a typical Silicon photodiode.

Chapter 7

Induced Longitudinal Path Length Noise

7.1 Introduction

Following the work described in chapter 6, on the QP22-E photodiode characterisation, a potential noise source was identified and measured. The noise – an apparent longitudinal pathlength change – was predicted to appear when a non-stationary optical beam was detected on a QP22-E photodiode. As the beam’s centre, moves around the photodiode’s active area, so the non-uniformity of the photodiode response to both signal magnitude and phase, generates an apparent longitudinal pathlength change. This unexpected source of noise is potentially important, because longitudinal pathlength is the primary measurement in LISA PathFinder (LPF) and all LISA-type missions. Although the flight optical bench of LPF was fitted with InGaAs photodiodes, the flight spare bench had to be populated with the QP22-E photodiodes due to the InGaAs photodiodes being unavailable. To analyse the significance of

the effect, further understanding and investigation was required.

In LPF the beams reflect from free floating test masses in geodesic motion. The relative position and angular noise between the spacecraft and test masses cause the beams to move. The residual motion between the test masses and the optical bench is expected to be of the order 10^{-7} radians/ $\sqrt{\text{Hz}}$. Therefore in the interferometers which use the test masses as mirrors there are non-stationary beams, which move across the photodiode's active area.

To determine the magnitude of the issue a numerical simulation was performed, the results of which prompted a follow up experimental investigation. The model was based on the results of several different tests performed on multiple QP22-E photodiodes. The phase lag, cross-talk, and other non-uniform response facets of the QP22-E were combined into a model which was used to simulate the output of each quadrant as a beam was moved across it.

The experimental test used one of the QP22-E photodiodes and a single beam of amplitude modulated (AM) light. The position of the beam's centre was radially displaced from the centre of the photodiode.

7.2 Model of QP22-E characteristics

The model of the photodiode was a simple set of matrices created in Matlab. The matrices were 1000×1000 , with each entry corresponding to a $5 \times 5 \mu\text{m}$ square of the photodiode's active area. Four matrices, each describing the response of a single quadrant as; a coefficient between the input light signal, and the output photo-current; while a second four matrices, gave the delay between the signal's arrival, and the output of a photo-current. All of these were based on scans of QP22-E photodiodes in operation without reverse bias. The beam was modeled as a single 1000×1000 matrix which contained a Gaussian intensity pattern of radius 0.8 mm. The beam's centre was moved radially from the photodiode's centre, to a displacement of up to $20 \mu\text{m}$.

7.2.1 Intensity response of QP22-E

Figure 7.1 shows the cropped simulated response of a single quadrant of the photodiode in the model. This was based on the profile of the 1-D scans of QP22-E photodiodes.

The response of the photodiode was modeled on the large scale changes in response seen in figure 6.10, and not the small scale ($10 - 20 \mu\text{m}$) variations evident in figure 6.8. The efficiency of the photodiode at converting incident 1064 nm light, into a photo-current is measured as a coefficient. There is a clear fall off towards the edge of the quadrant which starts $\sim 400 \mu\text{m}$ inside the quadrant and continues for $\sim 800 \mu\text{m}$. The profile of this was modeled as a Gaussian cumulative distribution function (cdf), which was subtracted from a uniform photodiode response.

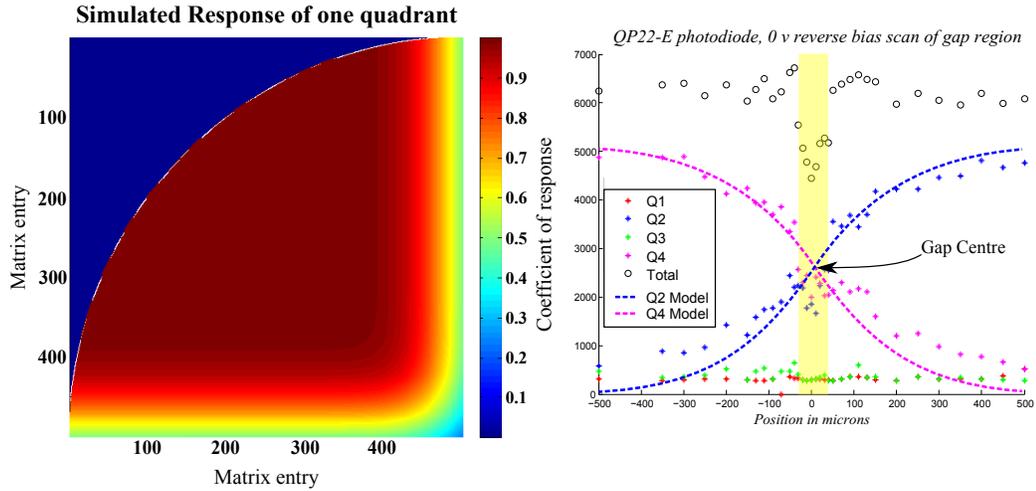


Figure 7.1: Response of a single quadrant, modeled in Matlab. The shape of the response is based on 1-D scans performed at the University of Glasgow. The graph is cropped to show a 500×500 section over which the photodiode would ideally have a flat response.

7.2.2 Phase lag of QP22-E

The phase lag is caused by the delay between, the signal arriving, and a photocurrent being read out. This was measured as the angle ' ϕ ', where ' ϕ ' is the phase lag for an input signal of frequency ' ν '. For this modeling a typical LPF heterodyne frequency of 1 kHz was used throughout. This avoids any complications when comparing with frequency dependent phase shifts.

$$\text{Phase lag } \phi = \frac{2\pi}{\nu}$$

The phase response of the QP22-E photodiode was measured by Astrium (see figures 6.13 & 7.2). We were allowed access to this data which showed that the phase lag had an unexpected spatial variation. The lag was shortest close to the edge between the quadrant and the gap; rose up to a flat value in the central area of the quadrant, before rising further at the outer edge of the photodiode.

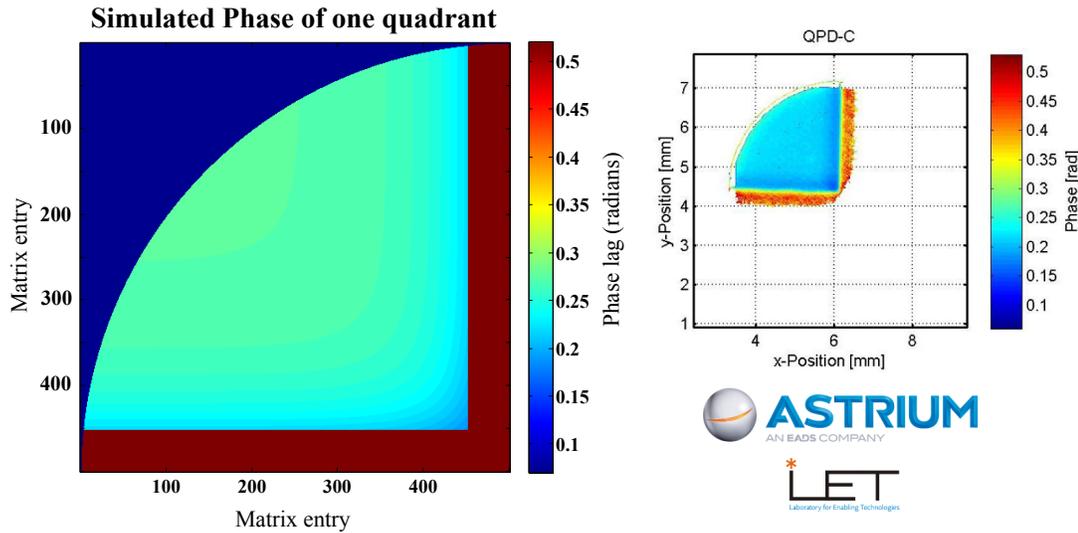


Figure 7.2: Phase response of a single quadrant, modeled in Matlab. The shape of the phase response is based on the Astrium scans [52].

As seen in figure 7.2 the modeled phase lag approaches 0.2 radians at the edge of the quadrant beside the gap. The central region of the modeled quadrant was 0.27 radians. The Astrium data suggested that the phase of the cross-talk was measured to be ~ 0.5 radians, and that the gap region had a phase of $0.3 - 0.4$ radians, but both were uncertain. The gap was not included in the model due to the low resolution of the data available; however the cross-talk phase was.

7.2.3 Simulation results

The full simulation including magnitude and phase response was run for a range of beam displacements and the results are shown in figure 7.3.

The matrices were combined in the following way:

Simulation of the coupling of relative motion between diode and beam to apparent longitudinal phase signal size

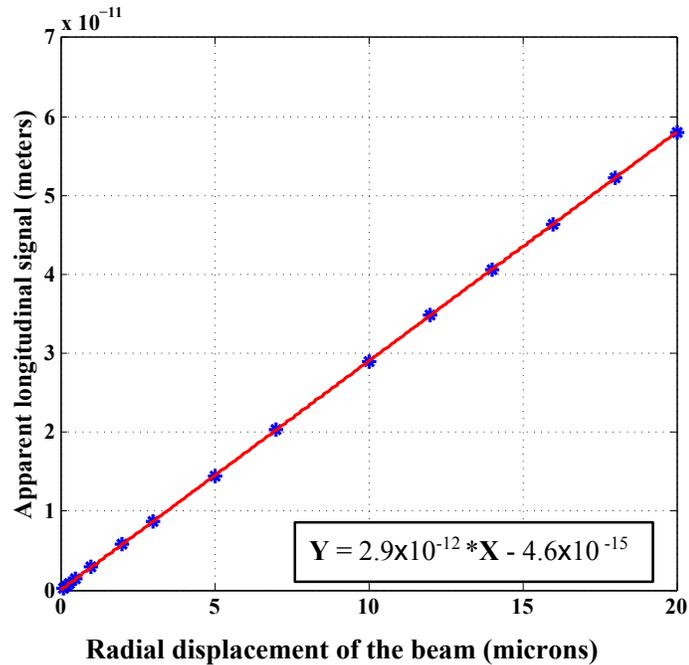


Figure 7.3: Results showing a coupling between, beam motion across the simulated photodiode and the resulting apparent longitudinal signal.

- The 'beam' matrix was multiplied by the intensity response matrices. This produced matrices with values which represent the magnitude of the photo-current.
- This new set of matrices (one for each quadrant) was combined with the phase lag matrices. This changed the values from magnitudes, to complex numbers.
- The next step summed the complex numbers for each quadrant into an overall signal magnitude and phase for each quadrant.
- The final step took the combination of these signal magnitude and phase for the four quadrants, and then (using equation 7.1) calculated a path-

length ' L '.

The phase of the signal for each quadrant changed as the beam moved across the photodiode active area. This change in phase is the unwanted effect that we wished to investigate. The four signals were then summed together to get an overall phase ' θ '. This phase ' θ ' could then be interpreted as:

$$L = \frac{\theta\lambda}{2\pi} \quad (7.1)$$

Where ' λ ' is the wavelength of the light 1064 nm and ' L ' is the apparent longitudinal pathlength change. The model coupling factor – in figure 7.3 – was $2.9 * 10^{-6}$ or 2.9 *picometers of pathlength / micron of beam motion*. An uncertainty on this value is hard to judge as it was based heavily on the model assumptions. The result had a weak dependence on the profile of the intensity response for a quadrant. The phase response showed a variation of 5 – 10% when the phase lag step 'height' of a quadrant was changed by ± 0.02 radians.

7.3 Experimental investigation

An experiment was setup to verify the magnitude of the effect seen in figure 7.3. It consisted of an optical beam, AM at 1 kHz and incident on a quadrant photodiode which was actuated across the beam on a linear stage. The signal magnitude and phase from each quadrant was measured. The beam was designed to remain stationary, while the photodiode moved to avoid potential angular changes in the pointing of the beam as it was actuated. Such angular changes in the output beam vector of the beam delivery system – an LPF FIOS – could result in larger than expected changes of beam position on the photodiode. The position of the beam on the quadrant photodiode could be varied. The setup in figure 7.4, is a picture of the equipment used and its layout.

Each photodiode quadrant was read out by electronics – identical to the flight electronics of LPF – into a commercial cDAQ 16-bit adc system, running under Labview, see figure 6.3. The program recorded the signal magnitude and phase relative to a reference signal. The reference signal was phase locked to both the signal generator which produced the 1 kHz amplitude modulation signal for the light and the 10 Hz modulation signal for the piezo which actuated the photodiode. The piezo was driven at 10 Hz over the range 1 – 20 μm pk-pk.

The hexapod had an uncertainty in absolute position of less than 1 μm over a 15 mm range [49] and this was used to calibrate the linear stage on which the photodiode was mounted. The apparent beam position on the photodiode – as measured by a Labview program used to read out the photodiode quadrants – was checked by displacing the beam on the photodiode using the hexapod. This was then used in turn to calibrate the actuator on which the photodiode was mounted.

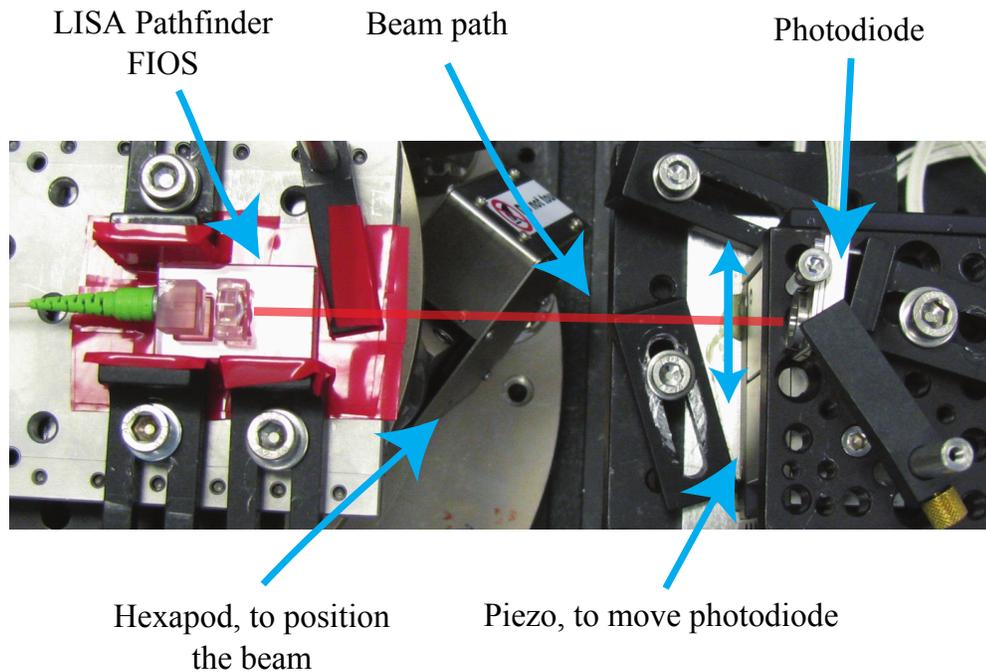


Figure 7.4: *Picture of the experimental setup. Shown in the picture is an initial test, involving a InGaAs photodiode. The LPF FIOS was mounted on a hexapod and used to deliver a geometrically stable beam of AM light which was centred on the photodiode. The photodiode was attach to a linear stage, which was actuated by a piezo, to produce the desired relative motion of beam to photodiode.*

7.4 Experimental results

The Labview output data was a magnitude and phase value for each quadrant. These were then combined in the same way as was done for the simulation, the results of which can be seen in figure 7.5.

The modeled result (figure 7.3) was 2.9×10^{-6} m/m and the experimental result is $2.9 \pm 0.3 \times 10^{-6}$ m/m. The experimental result agrees with that of the simulation; however the uncertainty in the experimental coupling factor is

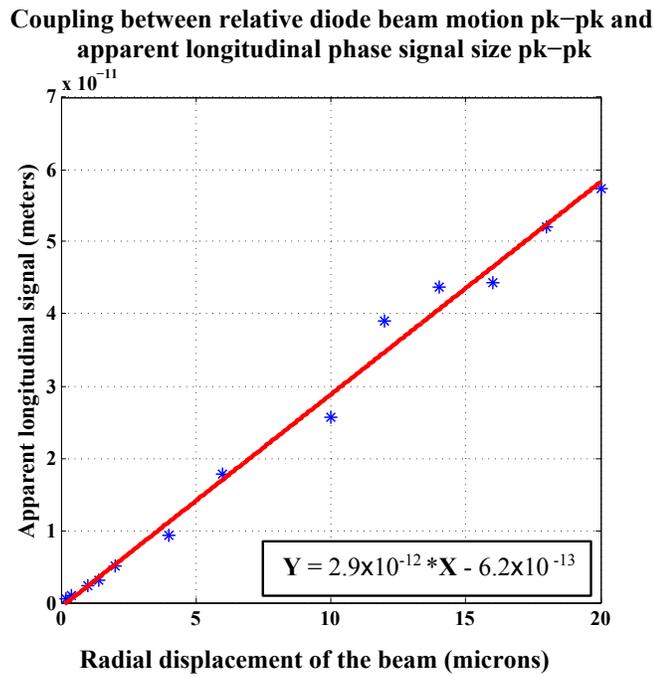


Figure 7.5: *Experimental results showing a coupling between, beam motion across the photodiode and the resulting apparent longitudinal signal. The gradient of the best fit line, is equal to that of the simulation.*

$\pm 0.3 \times 10^{-6}$ m/m. The uncertainty comes mainly from the beam-to-photodiode motion ($\pm 0.1 \mu\text{m}$).

7.5 Conclusion

To summarise the results of both the model and the experimental testing, this investigation provides a modeled coupling factor of 2.9×10^{-6} m/m and an experimental factor of $2.9 \pm 0.3 \times 10^{-6}$ m/m between the beam motion across a QP22-E photodiode, and a resulting apparent longitudinal pathlength signal. This value was measured using fully representative LPF equipment including: the phasemeter, photodiodes, radius of the beam and light power, and signal size and frequency. By using a single beam of AM light instead of two interfering beams, only one of which is non-stationary, many of the potential noise sources associated with interferometry are avoided.

Taking the coupling factor, derived from this experiment and applying it to the expected beam jitter from the test masses in LISA PathFinder, we find that the non-uniformity of the phase of the QP22-E photodiodes is ~ 0.29 pm/ $\sqrt{\text{Hz}}$ in the ‘X1’ science interferometer. With an expected performance goal of 6.3 pm/ $\sqrt{\text{Hz}}$ this a small but potentially significant noise source. The noise is generated via test mass angular tilt with respect to the optical bench. The angular tilt of the test masses with respect to the optical bench is measured through the dynamic wavefront sensor (DWS) [50], so in principle the noise source could be modeled and subtracted out in post processing.

Chapter 8

Beam Dump Optical Design and Testing

8.1 Introduction

In moving forward from LISA Pathfinder (LPF) to eLISA, there is a need to have beam dumps. Whereas LPF has interferometer beam powers of ~ 10 mW in each beam, by contrast eLISA's science interferometer has strongly mismatched beam powers of 1 mW local oscillator beam and a 100 pW received beam from the distant spacecraft. With such low light levels, the risk of phase sensing measurements being contaminated by stray light from unwanted beams is a strong possibility. The unwanted beams are produced from various places, typically beam splitting components, where both outputs ports of the splitting component are not required.

One place where a beam dump is required is the unwanted light from a splitting component in the optical path of the beam which is bound for the distant

spacecraft, see figure 8.1. This is component must reflect the out going beam and transmit the received beam. It achieves this through the use of a dielectric coating, which is designed to primarily reflect S-pol light and transmit P-pol light; however the coating does not produce ideal results and therefore there is some unwanted light. This unwanted light must be dumped. The beam dump must sufficiently suppress the light so that the science interferometer achieves its goal of $1 \text{ pm}/\sqrt{\text{Hz}}$ noise floor.

Commercially available beam dumps, typically used on optical tables are, unsuitable for use in space. This is due to their black powder coating which is not spaceflight compatible. It is feared parts of this black powder coating could flake off during launch – when the spacecraft with undergo violent shaking – and float around in the vacuum of the optical bench housing. If a flake of black absorbing powder were to land on any of the sensitive optical components it could cause a failure. To avoid this risk we have designed a beam dump, and sourced a coating suitable for eLISA.

8.1.1 Stray light from unwanted beams

The received light does not produce any unwanted beams; therefore it places no requirements onto the beam dump design. The local oscillator beam enters initially as a 1 W beam. This beam has several associated but unwanted beams which require suppression. The largest of these is a beam resulting from the beam splitting component, mentioned earlier, with complex coating requirements for both S-pol and P-pol light. The P-pol aspect is relevant to the received beam and not the local oscillator so we shall ignore it for now. The component would ideally reflect all of the incident S-pol light; however it has been measured that the coating is 99.74% reflective to P-pol. Thus an

unwanted S-pol beam of ~ 2.5 mW local oscillator light is left. This is the largest source of potential stray light.

8.2 Light suppression goal

8.2.1 Interferometer requirement

Given that the power of the local beam at the science interferometer is 1 mW, we must calculate the stray light level which could produce a noise level of 1 pm. The science interferometer is a phase sensitive measurement of the heterodyne beat note between the local oscillator and the received beam. The stray light (of the same polarisation) would pull the phase of the local oscillator beam, thus producing an apparent signal. To limit this signal to less than 1 pm, the local beam's phase would have to be pulled by less than 10^{-6} of a cycle because at 1064 nm, 10^{-6} of a cycle is ~ 1 pm. Beginning with an expression for the heterodyne signal in the science interferometer:

$$I_{het} \propto |\Psi_{LO} + \Psi_{SL} + \Psi_{Rx}|^2 \quad (8.1)$$

where ' Ψ_{LO} ' is the amplitude of the local oscillator light, ' Ψ_{SL} ' is the amplitude of the stray light, and ' Ψ_{Rx} ' is the amplitude of the received light from the distant spacecraft. Combining the stray light and local oscillator beam we can denote them as ' Ψ_x ', where:

$$\Psi_x = \Psi_{LO} + \Psi_{SL} \quad (8.2)$$

' Ψ_{LO} ' and ' Ψ_{SL} ' have the same frequency, but the phase between them ' ϕ ' could be expected to be unstable. If the phase was stable then this would produce a dc offset of the measured signal in the interferometer. This would not effect the performance of the experiment in a significant way; however

the stray light is unlikely to have a stable phase. The worst case scenario is where the phase is changing at a frequency in the measurement band of the interferometer. For this derivation we shall ignore the time variant component of ' ϕ ' and look to suppress the phase shift of ' Ψ_x ' with respect to ' Ψ_{LO} ' for worst case values of ' ϕ '.

Defining as follows:

$$\Psi_{LO} = E_{LO} \sin \omega t \quad (8.3)$$

$$\Psi_{SL} = E_{SL} \sin (\omega t + \phi) \quad (8.4)$$

$$\Psi_x = E_{LO} \sin \omega t + E_{SL} \sin (\omega t + \phi) \quad (8.5)$$

Where ' E ' is the amplitude of the wave, ' ω ' is the frequency, ' t ' is the time, and ' ϕ ' is the phase between the ' Ψ_{LO} ' and ' Ψ_{SL} '. Using the double angle formula for $\sin (\omega t + \phi)$ gives the following expression for ' Ψ_x ':

$$\Psi_x = E_{LO} \sin \omega t + E_{SL} [\sin \omega t \cos \phi + \cos \omega t \sin \phi] \quad (8.6)$$

Which leads to this result:

$$\Psi_x = [E_{LO} + E_{SL} \cos \phi] \sin \omega t + [E_{SL} \sin \phi] \cos \omega t \quad (8.7)$$

Using the trig equation 8.8:

$$a \sin x + b \cos x = \sqrt{a^2 + b^2} \sin (x + \theta) \quad (8.8)$$

With:

$$\theta = \tan^{-1} \left(\frac{b}{a} \right) \quad (8.9)$$

This allows ' Ψ_x ' to be expressed as a single sin wave:

$$\Psi_x = \sqrt{[E_{LO} + E_{SL} \cos \phi]^2 + [E_{SL} \sin \phi]^2} \sin(\omega t + \theta) \quad (8.10)$$

Where:

$$a = E_{LO} + E_{SL} \cos \phi \quad (8.11)$$

$$b = E_{SL} \sin \phi \quad (8.12)$$

From equations 8.11 & 8.12 it follows that the phase shift ' θ ' is:

$$\theta = \tan^{-1} \left(\frac{E_{SL} \sin \phi}{E_{LO} + E_{SL} \cos \phi} \right) \quad (8.13)$$

For values of $E_{LO} \gg E_{SL}$ equation 8.13 becomes:

$$\theta \simeq \left(\frac{E_{SL}}{E_{LO}} \right) \sin \phi \quad (8.14)$$

Thus to avoid a phase shift of 10^{-6} of a cycle, the fraction $\frac{E_{SL}}{E_{LO}}$ must be 10^{-6} or smaller i.e. an amplitude suppression of 10^{-6} . When applying this to the light powers in the science interferometer; the local oscillator beam of 1 mW requires that the stray light be less than 1 femtoWatt (10^{-15} W).

8.2.2 Stray light in a LISA-type optical bench

The goal of an amplitude suppression of 10^{-6} is for light in the science interferometer. The layout of a LISA-type optical bench provides additional factors to be considered in calculating the light suppression needed from the beam dump. The largest of these is the dielectric coated beam splitter – component ‘N’ – seen in figure 8.1, which is the source of the unwanted light but it also helps mitigate the light reaching the science interferometer.

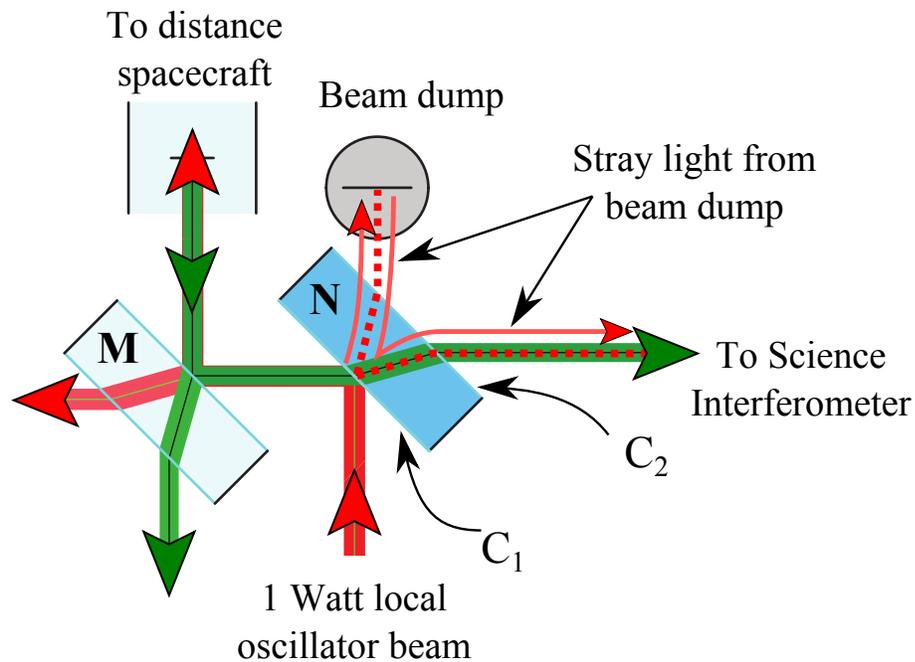


Figure 8.1: Shown is the area where the unwanted beam is produced. Labeled is the local oscillator light in red, and the received beam from the distant spacecraft in green. The coatings of the component ‘N’ are also labeled ‘C₁’ and ‘C₂’. Component ‘M’ does not affect the stray light. The light which is scattered or escapes the beam dump can couple back into the science interferometer along the optical path of the received beam.

The interaction with component ‘N’ and its coatings leaves most of the stray

light passing through the component instead of being reflected towards the science interferometer. The reason for this is that the LISA-type mission arrangement means that the light leaving the spacecraft (red light in figure 8.1) is a different polarisation to that of the light interfering in the science interferometer. The green light in figure 8.1 has been received from the distant spacecraft and passes through the component ‘N’ without much reflection. Thus for stray light of the correct polarisation to reach the science interferometer, then it must be reflected off component ‘N’ which is designed to transmit light of that polarisation in order to allow the received beam to reach the science interferometer. Component ‘N’ has been procured for an engineering level optical bench. The author has measured this coating as having the following reflection and transmission coefficients to S-pol and P-pol light. The coating C_1 is the complex coating and reflects/transmits 99.74%/0.26% of the incident S-pol light, the red local light bound for the distant spacecraft. C_1 was measured as reflecting/transmitting 0.5%/99.5% of the received P-pol light and the stray light scattered from the beam dump with the correct polarisation to effect the science interferometer.

Taking these effects into account the amplitude suppression required by a factor of 14 in amplitude suppression due to the interaction with component ‘N’. This relaxes the amplitude suppression goal at the beam dump to 1.4×10^{-5} , for a given input stray light beam of 2.6 mW from an incident 1 W beam on the coating C_1 .

8.3 Possible beam dump designs

8.3.1 Commercial beam dump design

The first step in the design of a suitable beam dump was to understand the different types of beam dump designs typically seen in a lab. A typical beam dump (see figure 8.2) is a black box, with a hole for light to enter. Inside the box there is a cone pointing towards the hole. Light entering the beam dump strikes the tip of a cone and is scattered onto the insides of the box. A black powder coating acts as a strong absorber to the light with some hemispherical scattering. This design of the beam dump only works as long as this strong absorbing coating is present.

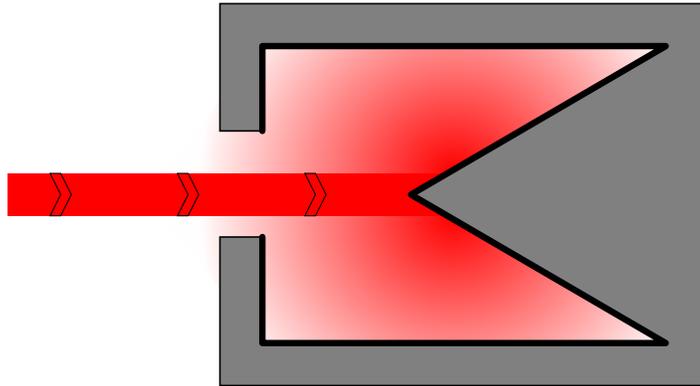


Figure 8.2: *Diagram of a commercial beam dump design. The light entering the beam dump is scattered off the conical point in the centre. The thick black line is the black powder coating which is a strongly absorptive material. The design makes it so that after 3 reflections, light can exit the beam dump along the same beam path it came in on. A typical commercial beam dump of this design has light amplitude suppression of 10^{-5} or less of the input light exiting the beam dump.*

8.3.2 Conic cutout design

Moving from this design, a starting point was to look at the cone shape as a cutout from a block, and the light reflections are internal to the conic shape (see figure 8.3). Without the black powder coating, it was thought that the way to suppress the light power of an incident beam was through multiple reflections. The taper angle ' θ ' of a conic cutout provided a variable to start this investigation. It is clear that the smaller the angle of the conic cutout, the larger the number of reflections needed before the light can exit the design. The goal was to maximise the number of reflections before light rays exit the design.

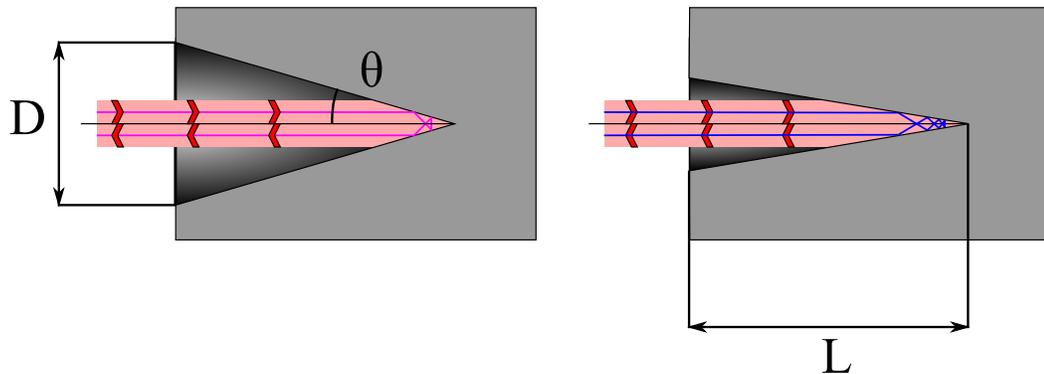


Figure 8.3: *Sectioned views of the conic cutouts. Shown for each design is a single ray entering the beam dump in the top half of the beam (red-pink area), before exiting the beam dump in the lower half of the beam, after multiple reflections in the conic cut out. The number of reflections for a ray traveling parallel to the axis of the conic shape is determined by the angle ' θ '. The ray in the left beam dump design (purple ray) has four reflections, while in the right beam dump design (blue ray), has eight.*

These initial designs tested were basic, and used to explore what could be done to maximise the number of reflections internal to the beam dump. The

analysis of the designs was done in Zemax, with the conic cut outs imported from a CAD design created in Solidworks. These CAD models were then imported to Zemax for analysis. A range of designs were investigated, with each subsequent CAD model had a sharper conic cutout than the last, resulting in more reflections before the light escaped. The length of the conic ' L ', and therefore the beam dump, became longer for smaller angles ' θ '. The diameter ' D ' of the design was 15 mm, which was set to be large enough to include rays, 3 beam radii from the centre of the beam.

Table 8.1: *Conic cutout beam dump results: angle, number of reflections, length and the minimum area of the beam dump foot print assuming a metal thickness of at less 5 mm.*

Angle ' θ '	Reflections	Length ' L ' (mm)	Minimal area (mm ²)
45°	2	7.5	313
30°	3	13.1	451
22.5°	4	18.2	579
15°	6	28.1	826
10°	9	42.6	1190
7.5°	12	57.0	1550
5°	18	85.8	2270
3.75°	24	114.5	2986

Table 8.1, shows the specifications of a number of conic cutout designs. The minimal area is an example of the foot print size that beam dump would have on the optical bench.

There were a number of problems with these conic cutout beam dumps:

- for larger numbers of reflections, the beam dump length became longer;
- after a certain number of reflections, all rays would exit the beam dump;
- scatter from all the reflections, internal to the beam dump, could exit without further reflection;
- the central point of the conic cutout would cause some light to be scattered out prematurely.

For these reasons, alternate solutions were explored. The options were: a spiral beam dump, or a light pipe with which to remove the light from the bench before dumping.

8.3.3 Light pipe

Light pipes are commercially available and seemed to offer the best solution. Any light which enters the light pipe, could be guided out of the spacecraft, into space. This meant that the light suppression would be limited by the coupling into the light pipe, and any scatter from internal reflections. In theory this could provide excellent suppression. However the potential problem in their use – and in any design involving light which is guided – is the possibility that light can be reflected back on itself. This could be caused by a tight bend of the light pipe. Ray-trace simulations showed that there was a clear minimal bend radius which the pipe must have. This radius was dependent on the light pipe diameter, which was chosen as > 3 times the beam radius. The minimum bend radius of the design looked at was > 3 times the diameter equivalent to > 18 times the typical radius of a beam on the eLISA optical bench ~ 50 mm.

Bends tighter than this ran the risk of propagating light back onto the bench. Complexities of integrating a light pipe with the correct bend radius through the spacecraft was another problem. Ultimately this route was judged to be higher risk than a standard beam bump.

8.3.4 Spirals

As an evolution from the light-pipe approach, a series of spiral designs was produced. The designs were spirals of finite length and ‘light traps’. The finite length meant rays would exit the design after a number of reflections. The light traps intended to totally internally reflect the light internal to the beam bump, thus cycling it multiple times. All the designs used a curved spiral entry section which then led to either a cylindrical drum which was a light trap, or simply a further spiral section. Like the light pipe there was a minimum bend radius dependent on the thickness of the spiral section and again this thickness was chosen as > 3 times the beam radius. A spiral-type beam dump was determined to be the best choice and was therefore investigated further.

8.4 Geometric analysis of spiral designs

The initial testing and development of the spiral and light ‘trap’ beam dump designs were their geometric shape and proportions. The analysis was focused on ensuring the beam dumps functioned as intended, through optimisation of their shape to maximise the number of reflections before a ray could escape. The overall volume of the beam dump was limited to $40 \times 60 \times 100$ mm; however the width (40 mm) and length (60 mm) were parameters which were best minimised.

8.4.1 Method

The setup was a simple light source and detector plane which covered the output of the beam dump. These tests used $\sim 10^3 - 10^4$ rays for analysis of each design and their iterations. The simulation was run and a limit to the number of reflects per ray was set at 10. Rays which didn’t reach the detector after 10 reflects didn’t appear; rays which did were drawn in the output. This allowed for investigation of the path taken by the rays (see figure 8.5).

8.4.2 Initial designs

The designs seen in figure 8.4 and figure 8.5, were modeled in Solidworks and then imported to Zemax as a CAD model for further analysis.

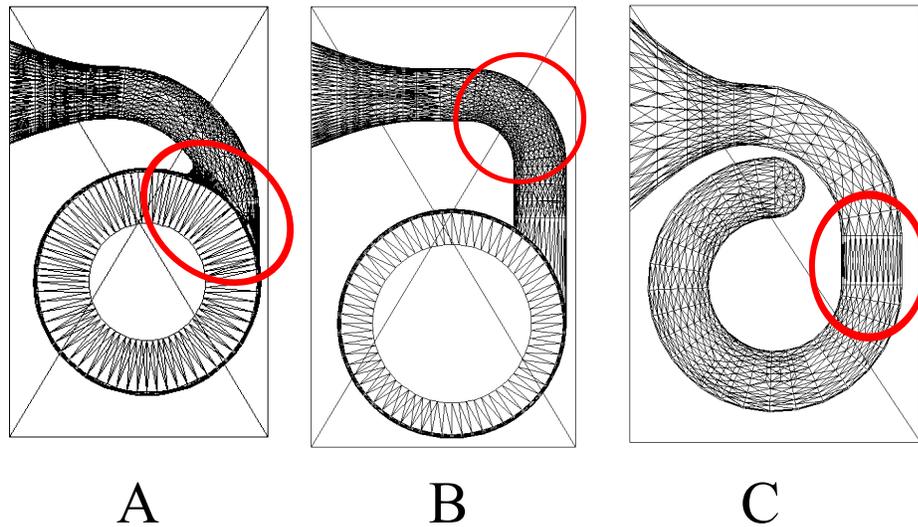


Figure 8.4: Shown are three different beam dump designs. Designs ‘A’ and ‘B’ are both light traps, where light is funneled into a cylinder to cycle. Design ‘C’ is a finite spiral, with an extra large opening diameter. The diagonal lines are artifacts of the Zemax output.

Design A

Design ‘A’ in figure 8.4, is a version of the first beam dump to have an empty cylinder in which the light could cycle around inside. The main issue with this design was the large opening (highlighted in figure 8.4), between the curved entry section and the cylindrical trap. This allowed rays to escape the design after 4 reflections. This number was reduced by narrowing the entry section; however this in turn proved to have problems with delivery of the light into the cylinder. Narrowing a curved entry section had the knock on effect of increased the minimum bend radius. For example a curved section with bend radius 75 mm was narrowed by 30% over a length of 100 mm. The resulting minimum bend radius of such a section increased from 50 mm to ~ 100 mm.

Design B

Design ‘B’ in figure 8.4, addressed the opening to the cylinder problem of design ‘A’ by changing the angle at which the entry section intersected the cylinder. This was made possible by the tighter bend in the entry section (highlighted in figure 8.4). While reducing the possibility of rays to exit the design once inside the cylinder ‘trap’ of the beam dump, the tight bend radius of the entry section caused issues. The input rays – from the edges of the beam profile – were able to make multiple reflections in this tight bend and exit the beam dump.

Design C

Having judged that a cylinder ‘trap’ design had several issues, a move was made to look at a finite spiral design. Design ‘C’ in figure 8.4, is a finite spiral which also has the same common issue of light exiting prematurely. The reason for this was determined to be that the short straight section between the curved entry section and the $\frac{3}{4}$ torus section which made up the rest of the beam dump. The goal of this design was to have light skim of the outer surface of the spiral and not touch the inner surface. The short straight length section (highlighted in figure 8.4), caused light to reflect off the inner surface, which upon entry to the torus section allowed for multiple reflections leading to rays exiting the design prematurely.

Design D

Design ‘D’ also known as the ‘final design’, seen in figure 8.5, came from a series of iterations where the radius of the spiral and the diameter of the spiral and

entry section where adjusted. The optimisation yielded a design where every ray that entered was guided fully to the end of the spiral. Once at the end it would head back out of the beam dump, but only after completing a total of at least 11 internal reflections. Building from the design ‘C’ the straight section was removed and the radius of the torus was reduced to accommodate for this. The reduced size of the spiral meant it had, the smallest footprint of any design capable of ensuring at least 10 reflections, internal to the beam dump. Its foot print on the bench of 25×55 mm (1375 mm²), made it smaller than the 7.5° conic cutout design shown in table 8.1.

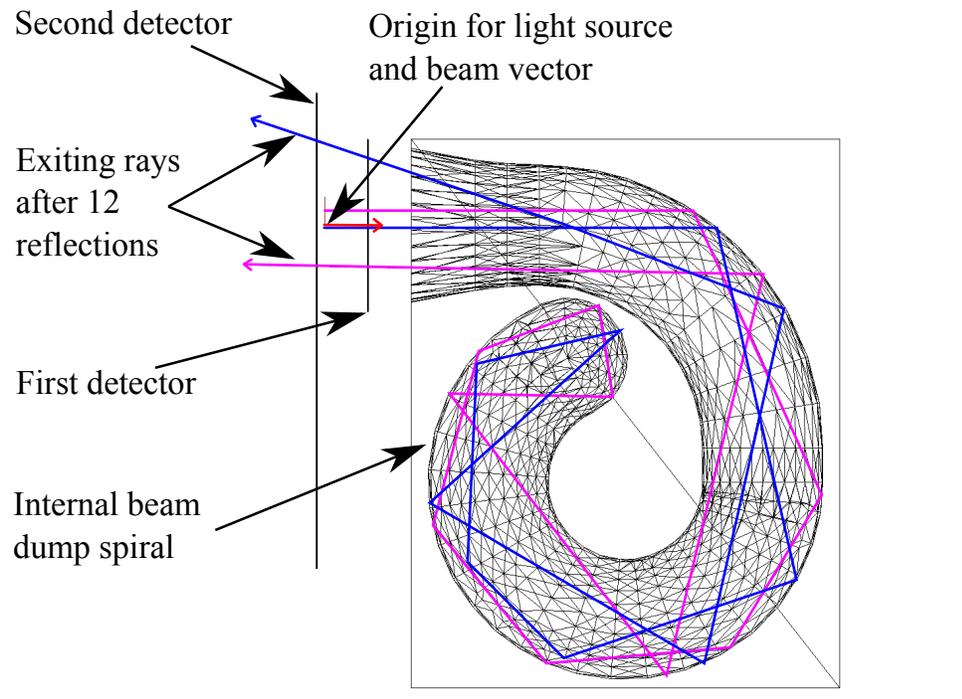


Figure 8.5: Design ‘D’ also called the ‘final design’. Shown is the location of the several objects: the origin of the beam, and two of the three detector planes (mentioned in more detail in the following section). Also shown are two exiting rays, each having made 12 internal reflections. Although it looks like ~ 10 reflections, there are reflections (on both rays) which cannot be seen in this projection.

8.5 Theoretical performance

This section takes the four designs mentioned in the previous section and measures their performance through a more detailed series of simulations. Again this was done through the use of the optical modeling program Zemax; although to analyse their performance $> 10^5$ rays were used.

8.5.1 Method

As can be seen in figure 8.5 there are a number of detector planes, and a beam origin with beam vector. The detectors for this analysis measured the number and position of rays passing through them.

The first detector was positioned to record the profile of the beam input to the design. The origin produced a beam of size and profile defined by the user. In this case a Gaussian beam profile of radius 2.23 mm was used (see figure 8.6). This beam profile was made up of 10^7 rays which then traveled into the beam dump.

The second detector a distance of 30 mm from the design, proved a profile of the rays exiting the beam dump. The profile of the rays which exit the beam dump shown in figure 8.7 was taken using the rays which exit the design after 11 – 13 reflections. This range of reflections was chosen after analysis of the power output from the beam dump was calculated for the coating Acktar ‘Vacuum BlackTM’, which was 98% absorptive (see table 8.2). It was assumed that the rest of the light was absorbed and that there was no hemispherical scatter of the reflected light.

A third detector was placed 60 mm from the design, with a size of only $20 \times$

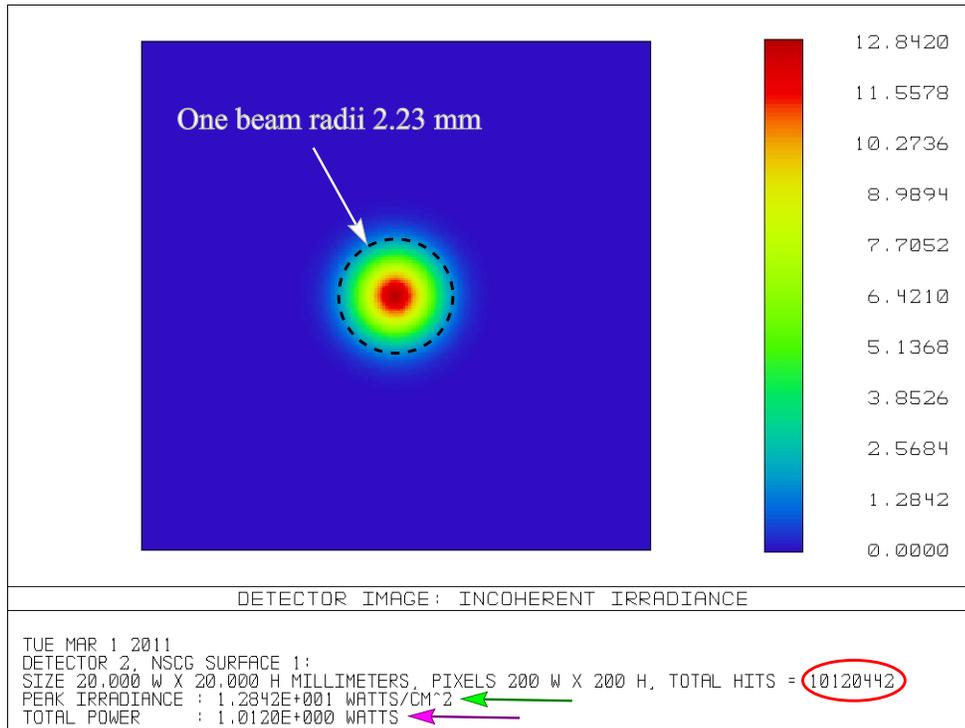


Figure 8.6: *Input beam profile at the origin labeled in figure 8.5. The red circle is the number of rays incident on the detector plane (10^7), the Green arrow is the ‘peak irradiance’ and the units of the colour bar (W/cm^2), and the pink arrow is the total power incident on the detector (1 W).*

20 mm. This third detector was used to estimate the number of rays traveling parallel to the input beam vector, but in the reverse direction.

The table 8.2 shows that the only significant contributions to the light power exiting the design come from the first two groups of rays to exit. It is only when the coating is $\sim 80\%$ absorptive that rays with > 12 reflections need to be considered.

The simulations were done as ray trace due to the incompatibility of the physics optic propagation (used in the FIOS modeling chapter 4) with the non-sequential setup. The beam was not propagated as a Gaussian beam and

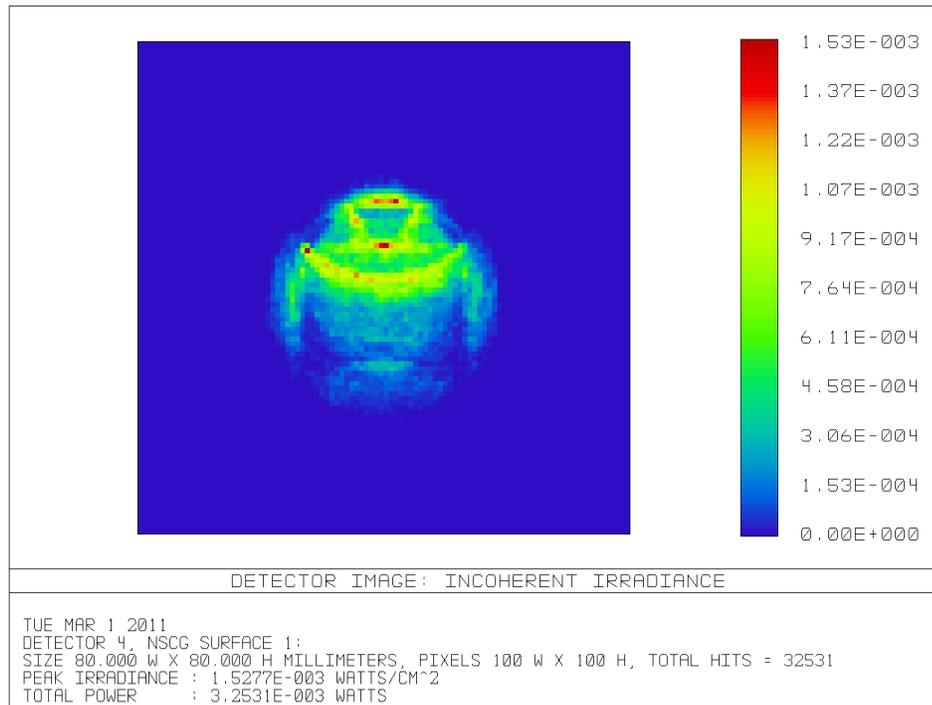


Figure 8.7: *Plot showing the statistical spread of rays exiting from the beam dump after 11 – 13 reflections. The area of this detector is 80 mm square and it is positioned 30 mm from the beam dump exit. Shown on the data for the plot is the number of rays which hit the detector, these number 3.2×10^4 of the 10^7 that entered.*

therefore means the output profile – measured by the detector planes – are likely to be of much higher contrast than the true output. All the simulations were run without any scattering profiles and without any absorption coefficients per reflection. The scattering profile of the intended coating and surface finish was unknown and so no scattering profile was used. All rays also maintained their starting power regardless of the number of reflections. To calculate the power exiting a beam dump, a summation of the number of rays leaving the beam dump design after a certain number of reflections was combined with the knowledge the power reduction per reflection. These setting lowered the

Table 8.2: *Total light power escaping from the ‘Final design’, assuming a total input power of 1 W (each ray being 10^{-7} W), no scatter on reflection and a 98% absorptive coating.*

Reflections	No. of rays	Power per ray	Total output power
11	11	2×10^{-27} W	2.2×10^{-26} W
12	360	4.1×10^{-29} W	3.6×10^{-26} W
13	910	8.2×10^{-31} W	3.65×10^{-26} W
14	3200	1.6×10^{-32} W	3.655×10^{-26} W
15	6800	3.3×10^{-34} W	3.6551×10^{-26} W

computational power required to perform the analysis but at the loss of accuracy. The balance of which was judged to be acceptable for the purposes of measuring the output from the beam dumps assuming no scattering.

8.5.2 Model results

A series of simulations were performed on the four beam dump designs ‘A’, ‘B’, ‘C’ and ‘D’. These simulated the distribution of rays exiting the designs, for a given number of internal reflections, table 8.3 shows the results. The input rays were only 10^5 in number to save computational time. This does mean that there is a possibility that a ray could exit the design ‘D’ after only 11 reflections; however it was not seen in this data set.

The data is graphically displayed in figure 8.8. The data shows that the spiral designs, ‘final’ and ‘C’, suppress the rays possibility to escape at lower numbers of reflections, better than the light traps, designs ‘A’ and ‘B’. However the opposite is true for higher numbers of reflections, with design ‘B’ a being best, and the ‘final’ design spiral the worst. This is however not as important because

Table 8.3: *Number of rays escaped from a design given 10^5 input*

No. of reflections	Final design	Design A	Design B	Design C
10	0	227	10	1
11	0	248	40	4
12	2	274	103	27
14	9	266	361	201
16	258	510	607	815
18	2213	1590	825	2164
20	8131	4042	1178	4208

after 20 reflections the light power is very much smaller than the light power of 10 reflections. Thus the final beam dump design offers better suppression of light power.

8.5.3 Simulation summary

The main results of the simulation suggest that a finite spiral is the best choice of beam dump design. This is due to better retention of light power, internal to the beam dump, over a light trap design. The traps offer lower numbers of rays exiting the design after > 18 reflections (see figure 8.8); however it is the suppression and containment of rays with low numbers of reflections, which carry the most power as shown by table 8.2.

The ‘final design’ CAD render can be seen in figure 8.9. Taking into account the lack of scatter at each reflection, a theoretical power suppression of $\sim 4 \times 10^{-26}$ is calculated for a coating of reflectivity 0.02. A titanium beam dump, of design ‘D’, has a theoretical power suppression of $\sim 5 \times 10^{-6}$ for a reflectivity of 0.6. In reality this is likely to be dominated by the scatter from the first reflection of the beam, assuming a hemispherical scatter of the 2% reflected light power. For a polished surface this could be improved from a hemispherical scatter to low angle scatter; however a model of such a surface would be poorly defined due to the lack of information about the finish of the realised beam dump and that of the coating intended for use on the beam dump.

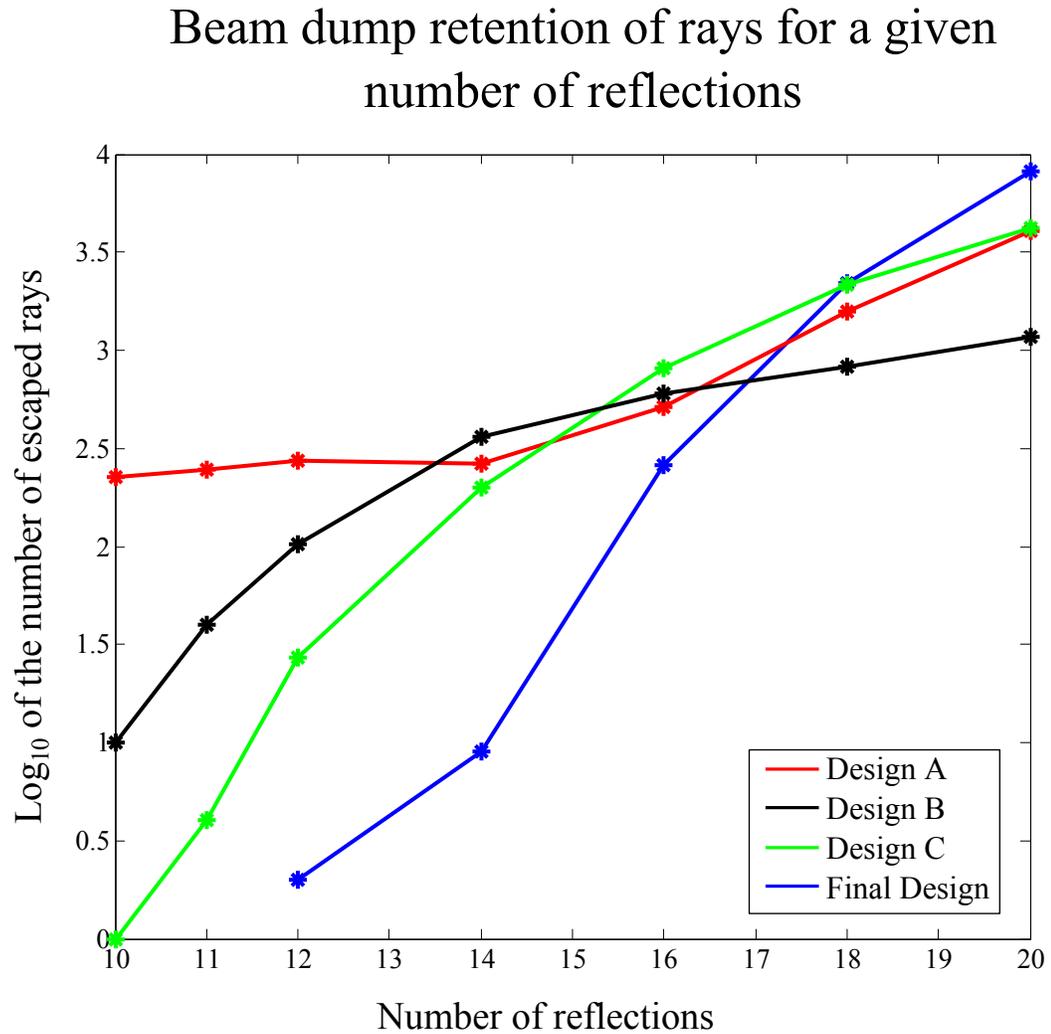


Figure 8.8: Plot of the beam dump designs, showing the variation in performance. Each design shows a different spread of rays exiting after a set number of reflections. Note that it is \log_{10} of the number of rays exiting that is plotted. The number of input rays was 10^5 .



Figure 8.9: A graphical render of the final beam dump design. Shown are the two halves of the beam dump, which allow the spiral to be machined out and the absorbing black coating applied before assembly. A small footprint on the optical bench area is achieved by the spirals, which make use of the available height above the optical bench. CAD model and render produced by Michael Perreur-Lloyd.

8.6 Experimental testing of a prototype beam dump

Two sizes of prototype beam dump design ‘D’ were procured. They were both had a rough machine finish; therefore making them non-representative of the fully realised final design. Made from Aluminum, these beam dumps were 95% reflective to 1064 nm light. Additionally one of each size of beam dump was anodised with a standard black coating.

8.6.1 Method

The method of testing adopted for measurement of the light suppression from the beam dump was a heterodyne interferometer. Interferometry is used to make a measurement of amplitude which is appropriate when light powers are small. Heterodyne interferometry makes for simple signal extraction and measurement.

The light source was an Innolight non-planar ring oscillator (NPRO) Nd:YAG laser, which feed 1064 nm light into a fibre splitter and a pair of fibre AOMs. Both the splitter and the fibre AOMs were from Gooch and Housego. The fibre coupled output of the AOMs feed the fibre couplers, seen in figure 8.10.

8.6.2 Optical setup

The beam dumps were held in a mounting with 5 adjustable axes. The X,Y and Z positions were adjustable, as well as two of the three rotational axes of the object. Rotation around the axis in which the beam was traveling was the

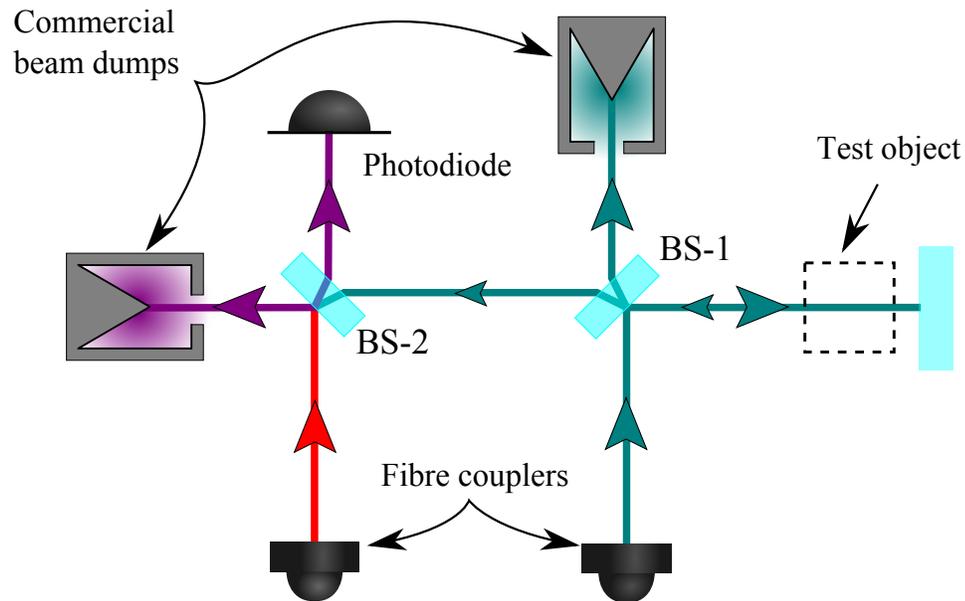


Figure 8.10: *The optical setup for the beam dump testing. Not shown are the fibre AOMs which fed the fibre couplers or a waveplate and polariser pairing which cleaned the output polarisation of each coupler. The red beam is the local oscillator beam ' ω_1 ', the blue beam is the probe beam ' ω_2 ', and the purple beam is the result of their combination into a heterodyne beat note. The unused port of each beam splitter 'BS-1' & 'BS-2', was dumped into a commercial beam dump. The optic behind the test object was a mirror, used to produce a reference heterodyne signal with which to compare the effectiveness of the beam dump.*

one degree of freedom, not possible. This ensured that the test object could be aligned the beam over a range of positions and angles. The angles of the test objects were adjusted over a range of values $\pm 20^\circ$ to find the optimised signal.

The risk of stray light affecting the measurement was reduced through careful suppression of unwanted light. This was done through the use of commercial beam dumps shown in figure 8.10. The optic 'BS-1' was cleaned before each

measurement, as scattering of the probe beam from this optic could reach the photodiode. The photodiode signal was amplified and measured on a Stanford FFT analyser.

8.6.3 RF cross-talk

The offset frequencies were produced by a dual channel signal generator. These frequencies ' ω_1 ' and ' ω_2 ' were 80.01 MHz and 80.02 MHz. This meant the heterodyne beat note was 10 kHz. These signals were then fed into two +30 dB RF amplifiers. The amplifiers could output a maximum of ~ 0.25 W, hence the input signals were ~ 0.25 mW. The output of the amplifiers drove the AOMs. Figure 8.11 shows the key electronics and signal paths between the signal generator and the AOMs.

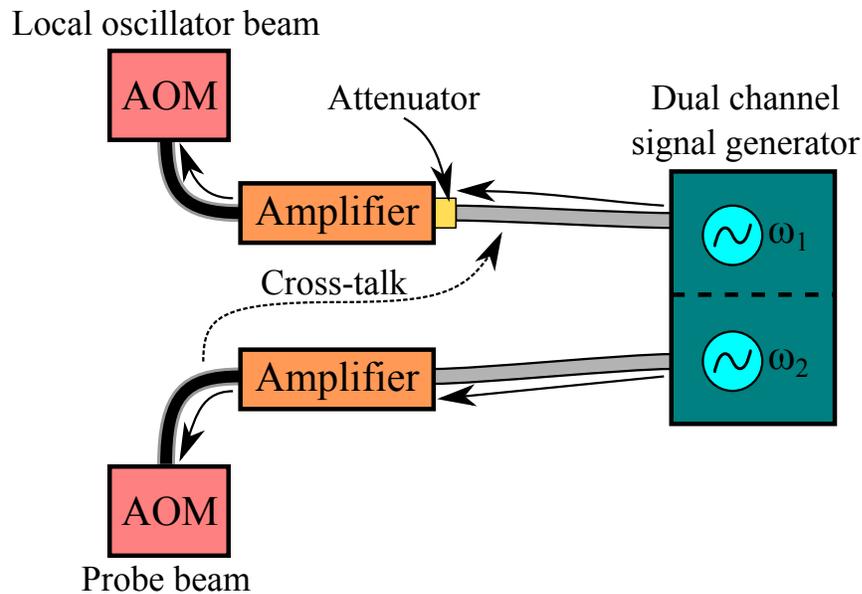


Figure 8.11: Shown is the RF cross-talk between the probe beam signal and the local oscillator signal.

The interferometer was initially limited in sensitivity by the cross-talk between

the two signal chains leading to the AOMs. The worse case and most likely method of RF cross-talk between the two signal chains is illustrated in figure 8.11. Here the output signal at ' ω_2 ' couples into the input of the ' ω_1 ' signal chain and is thus amplified. This produces a beam with both heterodyne frequency components, which can produce a signal on the photodiode without the presence of the other beam. This was measured initially to be -86 dB below the optimally align interferometer signal level.

To suppress the cross-talk the local oscillator ω_1 signal chain was attenuated by 12 dB on entry to the amplifier. This suppressed both the input signal and the cross-talk entering the amplifier; however by increasing the input signal size from the signal generator the output power of the local oscillator signal was maintained. This reduced the ratio between the local signal and the cross-talk signal of the probe beam by 12 dB. This led to the amplitude sensitivity limit of the interferometer changing from -86 dB to -90 dB. If the interferometer had still been limited by the cross-talk then the sensitivity limit should have dropped by 6 dB.

The reason this is most potent source of cross-talk is that the local oscillator beam is unsuppressed on its route to the photodiode. Thus the small amount of cross-talk is comparable to the probe beam signal from the beam dump. In the other case, where there is a small cross-talk into the probe beam, this beam is then suppressed by the beam dump. The amount of this signal reaching the photodiode is then insignificantly small when compared to the local oscillator signal.

8.6.4 Results

The suppression of each object tested was made by comparing the optimised interference signal between the local oscillator and the probe beam when the probe beam was reflected off a mirror. The measurements shown in table 8.4 are the result of a series of optimised setups. The limit of interferometer sensitivity was approximately -90 dB compared to the reference.

Table 8.4: *Amplitude suppression values for: prototype beam dumps, and flat sections with coating and finish consistent with the prototype beam dumps.*

Test object	Surface/finish	Suppression
Flat test section	Black/rough	-69 dB
Flat test section	Aluminum/rough	-64 dB
Large beam dump	Black/rough	-90 dB
Large beam dump	Aluminum/rough	-86 dB
Small beam dump	Black/rough	-82 dB
Small beam dump	Aluminum/rough	-84 dB

These results show that the black coating on the beam dumps did not improve the suppression of the beam dump design by a significant margin.

8.7 Analysis

The measured performance of the uncoated Aluminum prototype beam dumps was -84 dB for the small prototype design, and -86 dB for the larger beam dumps which are equivalent to coefficients of amplitude suppression 6.3×10^{-5} and 5.0×10^{-5} . To relate this to the performance goal of 1.4×10^{-5} set in section 8.2, the following modification must be discussed.

The design of this beam dump was made under the assumption that it would be used with an absorptive black coating. The Acktar ‘Vacuum BlackTM’ coating selected for the eLISA beam dump, is 2% reflective at 1064 nm [54], compared with Aluminum which is 95% reflective. When applied to one of these beam dumps it is expected that this will be equivalent to a factor of 6.9 in amplitude suppression by a beam dump, see table 8.2. This relaxes further the amplitude suppression goal from 1.4×10^{-5} to 9.8×10^{-5} . The goal of 9.8×10^{-5} is less suppression than both of the Aluminium beam dumps achieved.

8.8 Conclusion

The scatter of stray light from a beam dump on the ‘eLISA’ optical bench could make its way to the science interferometer, where it could possibly pull the phase of the science interferometer signal and thus induce path length noise. To avoid this an upper limit on the amplitude of stray local oscillator light interfering in the science interferometer was calculated as being 1×10^{-6} of the amplitude of the local oscillator beam. This challenging target can be moderated, through assumptions about the affects of the coatings of the components and the beam dump. When considered these factors combine to give a relaxed suppression goal of the beam dump to 1.4×10^{-5} , where with the addition of an absorptive coating on the beam dump further reduces the goal of an uncoated beam dump design – as presented here – to 9.8×10^{-5} . The measurements taken using the prototype beam dumps showing an amplitude suppression of 5×10^{-5} . This suggests that the fully representative beam dumps are capable of suppressing the amplitude of the stray light in the science interferometer.

Chapter 9

Outlook

As we look ahead to the second half of the decade, it is widely believed that the detection of gravitational waves will come with the second generation ground based detectors. This would provide the opportunity to look at the universe in a new ‘light’ through the gravitational wave window. The window is expected to be filled with a grand variety of exciting sources. A LISA-type mission like eLISA would enable the study of a range of gravitational wave sources inaccessible to the ground based detectors.

LISA Pathfinder is due to launch sometime in the near future around 2015. With the successful launch and operation of this mission, most of the technical risk will have been demonstrated for a LISA-type mission. This would leave a possible LISA-type mission primed for selection in the post 2020 era. To this end much of the evolution of the technology transfers from LISA Pathfinder to a LISA-type mission has been completed.

Bibliography

- [1] Albert Einstein. *"Relativity : the Special and General Theory"*. Methuen & Co Ltd, 1920.
- [2] Bernard F. Schutz. *A First Course in General Relativity*. Cambridge Univ. Press, 1985.
- [3] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Ajpl*, 195:L51–L53, Jan 1975.
- [4] J. Weisberg and J. Taylor. Relativistic binary pulsar b1913+16: Thirty years of observations and analysis. *ASP Conference Series*, 328:25–32, 2004.
- [5] J. Antoniadis et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340(6131), Apr 2013.
- [6] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, and et al. Planck 2013 results. I. Overview of products and scientific results. *ArXiv e-prints*, March 2013.
- [7] Matthew Pitkin, Stuart Reid, Sheila Rowan, and James Hough. Gravitational wave detection by interferometry (ground and space). *Living Reviews in Relativity*, 14(5), 2011.

-
- [8] B. P. Abbott and the LIGO Scientific Collaboration. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901+, July 2009.
- [9] J. Degallaix and the VIRGO Scientific Collaboration. Advanced Virgo Status. In G. Auger, P. Binétruy, and E. Plagnol, editors, *9th LISA Symposium*, volume 467 of *Astronomical Society of the Pacific Conference Series*, page 151, January 2013.
- [10] H Grote and the LIGO Scientific Collaboration. The GEO 600 status. *Classical and Quantum Gravity*, 27(8):084003, 2010.
- [11] Masaki Ando and the TAMA Collaboration. Current status of the TAMA300 gravitational-wave detector. *Classical and Quantum Gravity*, 22(18):S881, 2005.
- [12] Pau Amaro-Seoane et al. Low-frequency gravitational-wave science with eLISA/NGO. *Classical and Quantum Gravity*, 29(12):124016, 2012.
- [13] J R Smith and the LIGO Scientific Collaboration. The path to the enhanced and advanced LIGO gravitational-wave detectors. *Classical and Quantum Gravity*, 26(11):114013, 2009.
- [14] K Kuroda and the LCGT Collaboration. Status of LCGT. *Classical and Quantum Gravity*, 27(8):084004, 2010.
- [15] B Willke et al. The GEO-HF project. *Classical and Quantum Gravity*, 23(8):S207, 2006.
- [16] Jim Hough, Sheila Rowan, and B S Sathyaprakash. The search for gravitational waves. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 38(9):S497, 2005.

- [17] Alan J Weinstein. Astronomy and astrophysics with gravitational waves in the advanced detector era. *Classical and Quantum Gravity*, 29(12):124012, 2012.
- [18] K Danzmann for the LISA Study Team. LISA - an ESA cornerstone mission for a gravitational wave observatory. *Classical and Quantum Gravity*, 14(6):1399, 1997.
- [19] K Danzmann. LISA — an ESA cornerstone mission for the detection and observation of gravitational waves. *Adv. Space Res.*, 32:1233–1242, 2003.
- [20] D A Shaddock. Space-based gravitational wave detection with LISA. *Classical and Quantum Gravity*, 25(11):114012, 2008.
- [21] O Jennrich. LISA technology and instrumentation. *Classical and Quantum Gravity*, 26(15):153001, 2009.
- [22] M Armano and the LISA Pathfinder Science Team. LISA Pathfinder: the experiment and the route to LISA. *Classical and Quantum Gravity*, 26(9):094001, 2009.
- [23] F Antonucci and the LISA Pathfinder Science Team. The LISA Pathfinder mission. *Classical and Quantum Gravity*, 29(12):124014, 2012.
- [24] D I Robertson, E D Fitzsimons, C J Killow, M Perreux-Lloyd, H Ward, J Bryant, A M Cruise, G Dixon, D Hoyland, D Smith, and J Bogenschahl. Construction and testing of the optical bench for LISA Pathfinder. *Classical and Quantum Gravity*, 30(8):085006, 2013.
- [25] Pau Amaro-Seoane et al. eLISA: Astrophysics and cosmology in the millihertz regime. 2012.
- [26] NASA. The Gravitational Wave Spectrum, 2012. <http://science.gsfc.nasa.gov/663/research/>.

-
- [27] B. Abbott and the LIGO Scientific Collaboration. Beating the spin-down limit on gravitational wave emission from the crab pulsar. *The Astrophysical Journal Letters*, 706(1):L203, 2009.
- [28] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, B. Allen, and et al. Beating the Spin-down Limit on Gravitational Wave Emission from the Vela Pulsar. *apj*, 737:93, aug 2011.
- [29] B F Schutz. Gravitational wave astronomy. *Classical and Quantum Gravity*, 16(12A):A131, 1999.
- [30] Nils Andersson et al. The transient gravitational-wave sky. *Classical and Quantum Gravity*, 30(19):193002, 2013.
- [31] Steve Drasco and Scott A. Hughes. Gravitational wave snapshots of generic extreme mass ratio inspirals. *Phys. Rev. D*, 73:024027, Jan 2006.
- [32] Leor Barack and Curt Cutler. LISA capture sources: Approximate waveforms, signal-to-noise ratios, and parameter estimation accuracy. *Phys. Rev. D*, 69:082005, Apr 2004.
- [33] Steve Drasco and Scott A. Hughes. Rotating black hole orbit functionals in the frequency domain. *Phys. Rev. D*, 69:044015, Feb 2004.
- [34] Carlos F. Sopuerta. Probing the strong gravity regime with eLISA: Progress on EMRIs. 2012.
- [35] Jonathan R. Gair, Christopher Tang, and Marta Volonteri. LISA extreme-mass-ratio inspiral events as probes of the black hole mass function. *Phys.Rev.*, D81:104014, 2010.
- [36] Dz-Hung Gwo. Ultra precision and reliable bonding method, September 4 2001. US Patent 6,284,085.

- [37] E J Elliffe, J Bogenstahl, A Deshpande, J Hough, C Killow, S Reid, D Robertson, S Rowan, H Ward, and G Cagnoli. Hydroxide-catalysis bonding for stable optical systems for space. *Classical and Quantum Gravity*, 22(10):S257, 2005.
- [38] A.A. van Veggel, D. van den Ende, J. Bogenstahl, S. Rowan, W. Cunningham, G.H.M. Gubbels, and H. Nijmeijer. Hydroxide catalysis bonding of silicon carbide. *Journal of the European Ceramic Society*, 28(28):303–310, 2008.
- [39] S. Rowan, S.M. Twyford, J. Hough, D.-H. Gwo, and R. Route. Mechanical losses associated with the technique of hydroxide-catalysis bonding of fused silica. *Physics Letters A*, 246(6):471 – 478, 1998.
- [40] P.H. Sneddon, S. Bull, G. Cagnoli, D.R.M. Crooks, E.J. Elliffe, J.E. Faller, M.M. Fejer, J. Hough, and S. Rowan. The intrinsic mechanical loss factor of hydroxy-catalysis bonds for use in the mirror suspensions of gravitational wave detectors. *Classical and Quantum Gravity*, 20(23):5025–5037, 2003.
- [41] J Bogenstahl, L Cunningham, E D Fitzsimons, J Hough, C J Killow, M Perreur-Lloyd, D Robertson, S Rowan, and H Ward. LTP fibre injector qualification and status. *Journal of Physics: Conference Series*, 154(1):012011, 2009.
- [42] Kai Su, Jon V DeGroot Jr, Ann W Norris, and Peter Y Lo. Siloxane materials for optical applications. In *ICO20: Materials and Nanostructures*, pages 60291C–60291C. International Society for Optics and Photonics, 2006.
- [43] A A van Veggel, J Scott, D A Skinner, B Bezensek, W Cunningham, J Hough, I Martin, P Murray, S Reid, and S Rowan. Strength testing and

SEM imaging of hydroxide-catalysis bonds between silicon. *Classical and Quantum Gravity*, 26(17):175007, 2009.

- [44] L. Cunningham, P.G. Murray, A. Cumming, E.J. Elliffe, G.D. Hammond, K. Haughian, J. Hough, M. Hendry, R. Jones, I.W. Martin, S. Reid, S. Rowan, J. Scott, K.A. Strain, K. Tokmakov, C. Torrie, and A.A. van Veggel. Re-evaluation of the mechanical loss factor of hydroxide-catalysis bonds and its significance for the next generation of gravitational wave detectors. *Physics Letters A*, 374(39):3993 – 3998, 2010.
- [45] Nicola Louise Beveridge. *Characterisation of silicon-silicon hydroxide catalysis bonds for future gravitational wave detectors*. PhD thesis, University of Glasgow, 2012.
- [46] Christian J. Killow, Ewan D. Fitzsimons, James Hough, Michael Perreur-Lloyd, David I. Robertson, Sheila Rowan, and Henry Ward. Construction of rugged, ultrastable optical assemblies with optical component alignment at the few microradian level. *Appl. Opt.*, 52(2):177–181, Jan 2013.
- [47] E. D. Fitzsimons, C. J. Killow, M. Perreur-Lloyd, D. I. Robertson, A. Taylor, and H. Ward. LOB-UGL-TN-008.2-14 OB EBB Design Definition File - Opto-Mechanical Design, May 2011.
- [48] LightPath. Lightpath 352220, 2012. <http://www.lightpath.com/352220.html>.
- [49] PI. M-824 compact 6-axis-positioning system, 2013. <http://www.physikinstrumente.com/en/products/prdetail.php?sortnr=700815>.
- [50] Euan Morrison, Brian J Meers, David I Robertson, and Henry Ward. Automatic alignment of optical interferometers. *Applied Optics*, 33(22):5041–5049, 1994.

-
- [51] J. Bryant for the LTP OBI Team. S2-ugb-rp-3026 lewicki photo diodes incoming uniformity tests, December 2012.
- [52] H. Kögel, M. Gohlke, R. Samani, A. Peters, U. Johann, C. Braxmaier, and D. Weise. Comprehensive Experimental Study of the Elements in the LISA Interferometric Noise Chain. In G. Auger, P. Binétruy, and E. Plagnol, editors, *9th LISA Symposium*, volume 467 of *Astronomical Society of the Pacific Conference Series*, page 293, January 2013.
- [53] PerkinElmer Optoelectronics. Data sheet - Large Area InGaAs Photodiodes, 2013. <http://pdf1.alldatasheet.com/datasheet-pdf/view/14856/PERKINELMER/C30619.html>.
- [54] Acktar Advanced Coatings. Black Optical Coating - Vacuum Black, 2013. <http://www.acktar.com/category/VacuumBlack>.