Wave Front Sensing and Correction Using Spatial Modulation and Digitally Enhanced Heterodyne Interferometry

David Tarquin Ralph

A thesis submitted for the degree of
Doctor of Philosophy
at the Australian National University

Submitted
20th August, 2018

Supervisors, Spouse and the System.

Declaration

This thesis is an account of research undertaken between August 2013 and July 2018 at the Centre for Gravitational Physics, Department of Quantum Science, in the Research School of Physics and Engineering at the Australian National University, Canberra, Australia.

Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not be submitted in whole or part for a degree in any university.

David Tarquin Ralph
July 2018

Acknowledgements

It is an extraordinary experience for someone of my age and background to have the opportunity to undertake a PhD in Physics. While my first degree was in Physics, I completed it 45 years ago and graduated at a time when the study of Physics was quite different. There was no standard model, quantum mechanics was far less developed, there was no Matlab to solve complex equations and model physical systems and photonics had yet to be invented. The only laser I saw was the one wielded by Goldfinger as he tried to get James Bond to talk. (I know they were bluffing now because they weren't wearing goggles.) This year interferometry is exactly twice as old as I am!

Therefore I have more people than most to thank.

First and foremost is my supervisory team, Professor Daniel Shaddock who invented digital interferometry and patiently watched me grapple with the theory. The essential ideas explored here are his. Dr Paul Altin whose experimental and theoretical help was invaluable and Dr. David Rabeling, who helped me with many of the early set ups and showed me how to align a cavity, which, hopefully, I'll never have to do again. Collectively they brought an impressive amount of intellect, knowledge, patience and generosity and I am extremely grateful for their help and support.

I would also like to thank Professor David McClelland who was a member of my supervisory panel and leads our optics group here at ANU. These were momentous times for gravitational wave research. The detection of gravitational waves was a massive achievement and so was David's sustained contribution to that effort.

I should also thank the other students and post doctoral fellows I worked with – Dr. Bram Slagmolen who intervened at a number of strategically significant moments, Dr. Robert Ward and Dr Jong Chow, whose summer schools helped me fill my knowledge voids, Dr Roland Fleddermann and Dr. David Bowman who helped me in the early days of the first experiment, and the entire gravitational wave group who welcomed me wholeheartedly into their midst. They put up with all my inane questions and dad jokes and their friendship made most of this great fun.

Next I need to recognise the system that allowed someone of my age to return and study the frontiers of science. Few people have this opportunity and I remain amazed that I did. The downside of this is that we reward universities for making PhDs rather than employing them and I have seen several of my brilliant peers look for opportunities outside physics because of the uncertainty that surrounds careers in science. Do we want to be the clever country or not?

Finally, of course, I must thank Leslie, my long suffering wife, who has lived through 40 years of my obsessions and passions. The PhD has stopped us taking some trips, kept me in my study at night and generally got in the way of the life she would like to lead. I thank her for tolerating it.

What fun it has been to learn so much!

Abstract

This thesis is about light. Specifically it explores a new way sensing the spatial distribution of amplitude and phase across the wavefront of a propagating laser. It uses spatial light modulators to tag spatially distinct regions of the beam, a single diode to collect the resulting light and digitally enhanced heterodyne interferometry to decode the phase and amplitude information across the wavefront. It also demonstrates how using these methods can be used to maximise the transmission of light through a cavity and shows how minor aberrations in the beam can be corrected in real time. Finally it demonstrate the preferential transmission of higher order modes.

Wavefront sensing is becoming increasingly important as the demands on modern interferometers increase. Land based systems such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) use it to maximise the amount of power in the arm cavities during operation and reduce noise, while space based missions such as the Laser Interferometer Space Antenna (LISA) will use it to align distant partner satellites and ensure that the maximum amount of signal is exchanged. Conventionally wavefront sensing is accomplished using either Hartmann Sensors or multi-element diodes. These are well proven and very effective techniques but bring with them a number of well understood limitations. Critically, while they can map a wavefront in detail, they are strictly sensors and can do nothing to correct it.

Our new technique is based on a single-element photo-diode and the spatial modulation of the local oscillator beam. We encode orthogonal codes spatially onto this light and use these to separate the phases and amplitudes of different parts of the signal beam in post processing. This technique shifts complexity from the optical hardware into deterministic digital signal processing. Notably, the use of a single analogue channel (photo-diode, connections and analogue to digital converter) avoids some low-frequency error sources. The technique can also sense the wavefront phase at many points, limited only by the number of actuators on the spatial light modulator in contrast to the standard 4 points from a quadrant photo-diode. For ground-based systems, our technique could be used to identify and eliminate higher-order modes, while, for space-based systems, it provides a measure of wavefront tilt which is less susceptible to low frequency noise.

In the future it may be possible to couple the technique with an artificial intelligence engine to automate more of the beam alignment process in arrangements involving multiple cavities, preferentially select (or reject) specific higher order modes and start to reduce the burgeoning requirements for human control of these complex instruments.

Contents

List of Figures

List of Abbreviations

AC	Alternating current
ADC	Analogue-to-Digital Converter
ANU	Australian National University
AOM	Acousto-Optic Modulator
BS	Beamsplitter
CBC	Coherent Beam Combining
CCD	Charge-Coupled Device
CDF	Cumulative Density Function
CDMA	Code-Division Multiple Access
CIC	Cascade Integrator-Comb
CPU	Central Processing Unit
CW	Continuous Wave
DAC	Digital-to-Analogue Converter
DAM	Debris Avoidance Manoeuvre
DC	Direct current
DEHI	Digitally Enhanced Heterodyne Interferometry
DI	Digital Interferometry
DO	Digital Output
DR	Dynamic Range
DSP	Digital Signal Processing
EOM	Electro-Optic Modulator
FC/APC	Fibre-Coupled / Angled Physical Contact
FC/FPC	Fibre-Coupled / Flat Physical Contact
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FP	Forward Path
FPGA	Field-Programmable Gate-Array
FSM	Fast Steering Mirror
GPS	Global Positioning System
GRACE	Gravity Recovery and Climate Experiment
GSPS	Giga-Samples Per Second
GW	Gravitational Wave
IQ	In-Phase and Quadrature

LFSR	Linear Feedback Shift Register
LIGO	Laser Interferometer Gravitational Wave Observatory
LISA	Laser Interferometer Space Antenna
LO	Local Oscillator
LMA	Large Mode Area
LPF	Low-Pass Filter
LSB	Least Significant Bit
LOCSET	Locking of Optical Coherence by Single-detector Electronic frequency Tagging
LUT	Look-Up Table
MSB	Most Significant Bit
NCO	Numerically Controlled Oscillator
Nd:YAG	Neodymium-doped Yttrium Aluminium Garnet
NI	National Instruments
NPRO	Non-Planar Ring Oscillator
OLG	Open-Loop Gain
OPA	Optical Phased Array
PBS	Polarising Beamsplitter
PCIe	Peripheral Component Interface Express
PDF	Probability Density Function
PDH locking	Pound Drever Hall Locking
PLL	Phase-Locked Loop
PI	Proportional plus Integral
PRN	Pseudo-Random Noise
PSD	Power Spectral Density
PZT	Piezoelectric Transducer
RAM	Random Access Memory
RF	Radio Frequency
RMS	Root Mean Square
ROM	Read Only Memory
RPSD	Root Power Spectral Density
SND	Sum-and-Dump
SLM	Spatial Light Modulator
SNR	Signal-to-Noise Ratio
SPGD	Stochastic Parallel Gradient Descent
TCS	Thermal Compensation System
TF	Transfer Function
USB	Universal Serial Bus
VHDL	VHSIC Hardware Description Language
VHSIC	Very High-Speed Integrated Circuit

Introduction and Overview

For centuries light has been the province of artists and poets. During the last two centuries, however, our understanding of light has increased substantially. Since Theodore Maiman demonstrated the first working laser at the Hughes Research Laboratory on 16 May 1960 Ref [1], our knowledge and control has increased to the point were lasers are now a ubiquitous part of our daily lives from removing unwanted hair, levelling brick courses in new buildings and cutting metal, to supporting complex medical procedures, sensing trace amounts of escaping gases or driving the most precise instruments on the planet, the interferometers that have been used to discover gravitational waves.

Recently, the search for gravitational waves has increased the demand for precision and it has become increasingly important to be able to detect and correct the spatial shape of the signal laser beam so that measurements are more accurate and more efficient. Wavefront sensing and correction is the underlying theme of this thesis.

This chapter provides a overview of this thesis. It:

- begins with a very brief outline of how our understanding of light has evolved from a mystery to the well understood tool that is an an integral part of modern life and how, in the process, our concept of the wavefront has emerged,
- outlines the main issues with mapping a wavefront,
- reviews the current approaches to wavefront sensing and introduces a new approach which forms the first half of the work presented in this thesis, and
- concludes with an outline of the thesis and how the chapters fit together to tell a complete story.

1.1 The Mystery of Light - an Historical Review

To understand our current concept of a wavefront it is useful to briefly review how our understanding of light has developed.

The nature of light has been a puzzle since before recorded history and our understanding has evolved over time. As early as fifth century BC, Empedocles postulated that everything was made of four elements - earth air, fire and water and that Aphrodite lit fire in human eyes which shone out and illuminated objects so that we could see them Ref [2]. In 300BC Euclid decided that light travelled in straight lines and created the first formalism for the laws of refection Ref [3]. Around the same time the ancient Egyptians had developed rudimentary lenses. In 55 BC, Lucretius postulated the first particle theory describing light

as atoms which travelled infinitely fast although his ideas were not generally accepted Ref [4]. 250 years later Ptolemy described the refraction of light.

The next major advance was made in Basra in what is now Iraq. Ibn al-Haytham published the *Book of Optics* which attempted to explain both reflection and refraction Ref [5]. Ibn al-Haytham, was born around 965AD in Basra, Iraq and died around 1040AD in Cairo, Egypt. He was part of the golden age of the Muslim civilisation that saw many advances in science, technology and medicine,Ref [6]. He moved to Cairo to develop a regulation system for the Nile but this proved too hard and he spent his time developing theories in optics, geometry and astronomy,Ref [7]. His work in optics drew on and advanced the work of Ptolemy. He realised that the eyes received light rather than transmitting it and devised experiments to prove this. He also developed complete descriptions for reflection and refraction, concluding correctly that light travelled more slowly in denser media.

René Descartes (1596–1650) believed light was a mechanical property of the luminous body, and published a theory of the refraction of light assuming, incorrectly, that light travelled faster in a denser medium than in a less dense medium. Descartes arrived at this conclusion by analogy with the behaviour of sound waves. Although Descartes was wrong about the sign of the speed change, he was correct in assuming that light behaves like a wave and that refraction can be explained as the speed of light changes in different media.

Pierre Gassendi (1592–1655), proposed a particle theory of light which attracted a young Newton (1642-1727), who rejected the wave description of light on the basis that other waves, water etc., could bend around things while light appeared to travel only in straight lines. Newton described light in terms of corpuscles and successfully explained reflection but could not explain refraction. He was one of the first to examine the nature of light and discovered that white light could be broken into the colours of the spectrum using a prism Ref [8]. Newton's work on light was summarised in the highly influential book *Opticks*,Ref [9].

It was Robert Hooke (1635-1703) who first championed the wave based description of light. In his 1665 work *Micrographia*, he explained the origin of colours, developed a "pulse theory" and compared the spreading of light to that of waves in water, suggesting that light's vibrations could be perpendicular to the direction of propagation Ref [10].

Christiaan Huygens (1629-1695) expanded these ideas and developed the first mathematical wave theory of light in 1678, publishing it in 1690 Ref [11]. He proposed that light was emitted in all directions as a series of waves in a medium called the luminiferous ether. As waves did not appear to be affected by gravity, later revisited by Einstein, it was assumed that they slowed down upon entering a denser medium. His ideas have remained very important and his Huygens-Fresnel principle, which states that every point on a wavefront is a source of wavelets and that these wavelets spread out in the forward direction at the same speed as the source wave, remains very important as to understanding diffractions and wavefronts today. The new wavefront is a line tangent to all of the wavelets, Figure 1.1.

As the incident wave, shown in blue, strikes the mirror, it creates small wavelets which radiate out at the same speed as the incident wave as illustrated by the increasing grey half circles. The reflected beam emerges perpendicular to the wavefront formed by drawing a tangent to these half circles. Importantly the direction of propagation is perpendicular to this wavefront.

The wave theory took hold and for the next 200 years was pre-eminent. In around 1800, Thomas Young performed a famous experiment in which he illuminated two adjacent slits with the same light and showed the two resulting two beams interfered with each other

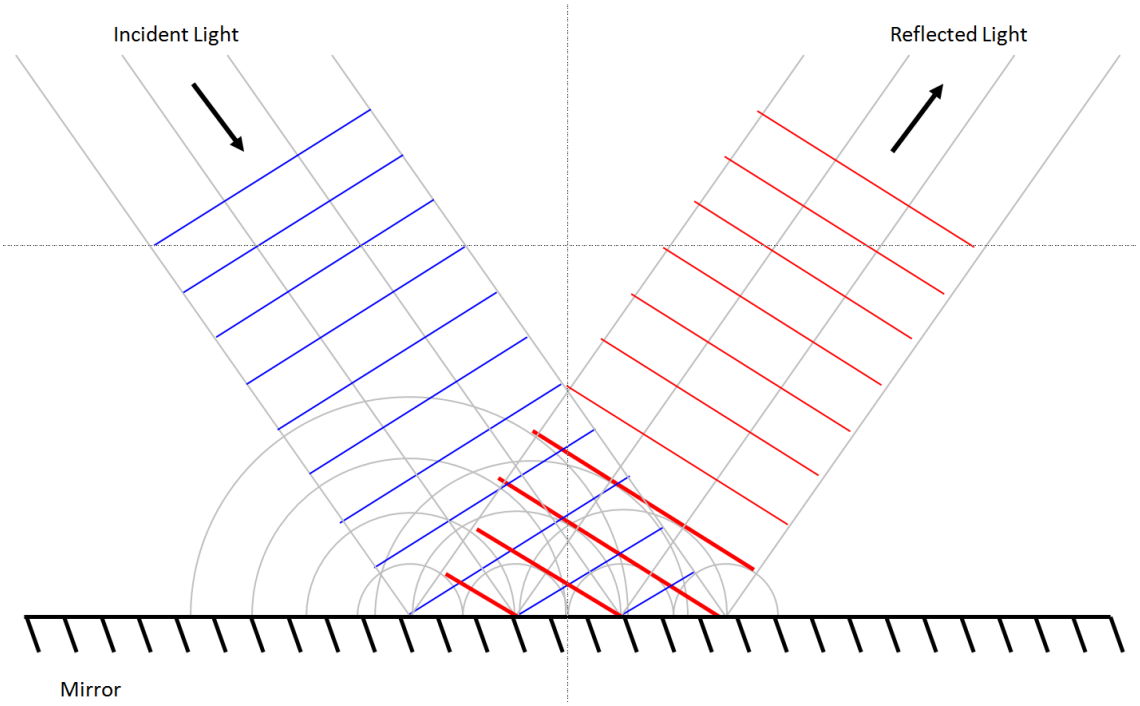

Figure 1.1: The Huygens-Fresnel principle states that every point on a wavefront is a source of wavelets, that these wavelets spread out in the forward direction and that the new wavefront is a line tangent to all of the wavelets

to produce a series of dark and bright fringes, bright fringes where peaks on the screen overlapped, reinforcing each other and dark fringes where peaks and troughs coincided cancelling each other out.

It was James Clerk Maxwell who finally created the formalism for the propagation of electromagnetic radiation in 1860. He realised that light, and indeed all electromagnetic waves, comprised an electric oscillation and a magnetic oscillation that were orthogonal but in phase in the direction of propagation, figure 1.2.

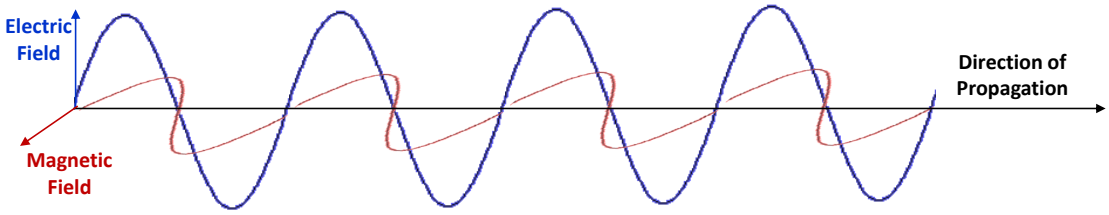

Figure 1.2: Light is a transverse wave comprising an electric oscillation and a magnetic oscillation that are orthogonal but in phase in the direction of propagation.

He summarised the propagation of light in 4 equations, Maxwell's equations, which are shown below in their integral form Ref [12] and which ultimately determine which modes can propagate in free space (2.1) and within a cavity (4.6):

- Gauss' law - $\int_S E.da = \frac{1}{\epsilon_0}\int \rho dV$ - Electric charges produce an electric field. The

electric flux across a closed surface is proportional to the charge enclosed.

- Gauss' law for magnetism - $\int_S B.da = 0$ - There are no magnetic monopoles. The total magnetic flux passing through any closed surface is zero.
- Faraday's law - $\oint E.ds = \frac{d}{dt}\int_S B.da$ - Time-varying magnetic fields produce an electric field.
- Ampère's law - $\oint B.ds = \mu_0 \int_s J.da + \mu_0\epsilon_0 \frac{d}{dt}\int_S E.da$ - Steady currents and time-varying electric fields (the latter due to Maxwell's correction) produce a magnetic field

This formalism is still used today and solving it leads to the paraxial Helmholtz equation that describes the propagation of a Gaussian laser beam,(2.1).

By the late 19 century the wave theory was pre-eminent. The interpretation of Young's experiment seemed irrefutable. A few observations remained unexplained however. Chief amongst these was the photoelectric effect. Physicists noticed that when some substances were illuminated with the right light they produced an electric current. Strangely this current depended on the colour of the light used. Red light might produce no current no matter how intense, while faint blue light might produce a current. Albert Einstein explained this phenomenon. He suggested that light comprised particles and that the energy carried by each was linked to their frequency which in turn determined their colour. Blue light is more energetic than red light so a single particle of blue light is able to displace an electron from a target while red light particles, however plentiful, lacked the energy individually to dislodge any electrons. Einstein's explanation, for which he won the Noble Prize, were presented concurrently with quantum theory and the notion of a wave particle duality emerged.

The strangeness of the wave particle duality is perhaps a reflection of the limitations of human experience and language rather than something intrinsically odd about the quantum state. Quantum theory is the province of atomic scales, which have only been accessible during the last 100 years. The vast majority of human experience has been concerned with the macroscopic world in which things are either waves or particles and our languages and understandings reflect this. We live in a world without direct exposure to quantum effects and it is only once we move to atomic scales that the quantum behaviour of nature becomes evident. It turns out on closer inspection, however, that all subatomic particles exhibit wave like properties.

We now understand that light is a stream of photons, particles with wave like characteristics which move at a set speed of 3x10^8 m/s per second (in vacuum) and which carry an energy that is dependent upon their frequency. As they move collectively, or indeed individually, they exhibit the properties of a transverse wave, that is they diffract, reflect, can be polarised and interfere. Further we now understand that light is only part of a much broader electromagnetic spectrum which extends from radio waves through, infra-red light, visible light, ultraviolet light, x-rays and gamma rays, figure 1.3.

1.2 What is a Wavefront

This thesis is about the detection and correction of wavefronts. This begs the question, what is a wavefront. For a propagating beam, a wavefront is a surface over which the optical wave, or more specifically the electric field, has the same phase (within a single cycle). It is a two dimensional surface in three dimensional space. For a plane wave,

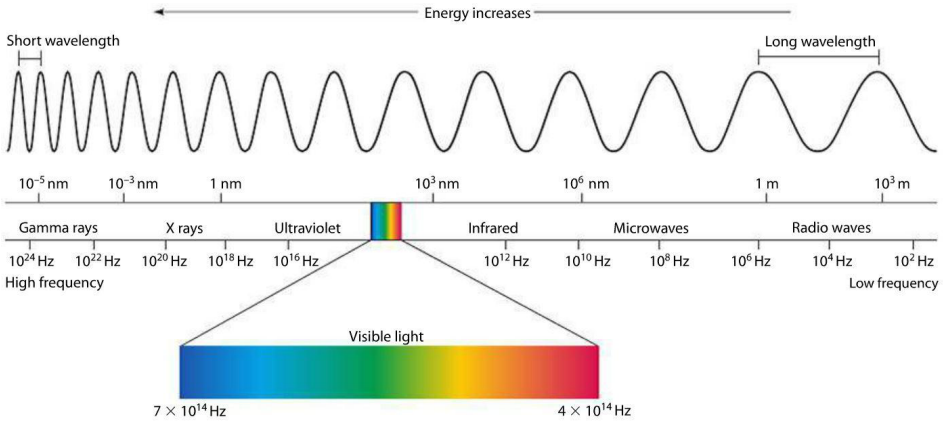

Figure 1.3: The Electromagnetic Spectrum

which only occur at beam waists or in the infinitely far field, the wavefront would be the plane perpendicular to the direction of propagation of a beam of light. For a laser beam, propagating as a Gaussian beam, it changes shape as it propagates from flat at the waist, to increasing curved as the beam approaches the Rayleigh Range and then becoming flatter again as the beam propagates to infinity, see equation 2.1 and figure 2.2. For reflections from an imperfect surface it can become quite jagged mimicking the irregularities in the surface from which it has been reflected.

A well defined optical wavefront typically requires a laser. A torch might provide a unidirectional stream of light but the phases and frequencies of the component photons are random making definition of a wavefront impossible. A laser provides a continuous beam of coherent, monochromatic light from a localised source. The Innolight Promethius Laser, which was used for the experiments outlined in the following chapters, is based on an Nd:YAG (neodymium-doped yttrium aluminium garnet; $Nd:Y_3Al_5O_{12}$) crystal and emits light at 1064nm in the near infra-red. The light emerging from the laser is localised, of an approximately single frequency and coherent in phase.

Of course, no system is perfect and the light from a real laser contains variations across all these attributes:

- Localisation and direction. A laser beam spreads as it propagates due to diffraction (see Chapter 2). Furthermore any minute imperfections in the laser cavity used to create it, inhomogeneous regions in the medium through which it travels and the optical elements it encounters on its journey adds small, additional distortions.
- Frequency. While the Nd:YAG laser nominally radiates at a single frequency there is, in reality, a range of frequencies spread around a peak at 1064 nms.
- Phase. While stimulated, the emission of photons from the lasing crystal remains a random, quantum process constant only at the macro level. This is shot noise and is due to uncertainty in the arrival times of photons at the sensor. It creates a spread in the precise phases of the photons emitted by the laser. This can be exacerbated as the beam propagates by flaws in the optics it meets and local variations in the medium through which it travels.
- Amplitude and Intensity. Quantum variations also affect the total intensity of the beam over time and its amplitude across its extent. In additional classical intensity noise create additional, small fluctuations in power.

1.2.1 The Importance of Wavefront Sensing

For many applications these variations are too small to be of concern. Compact disk players, spirit levels and etching machine can all function perfectly well with an imperfect laser. But as our application of lasers has evolved into a range of precision measuring tools so has the need for improved laser quality. In almost all high precision applications the quality of the beam directly couples into the accuracy of the measurement and the power required to operate the system.

The best examples is perhaps the Laser Interferometer Gravitational-Wave Observatory (LIGO), which has successfully been used to discover gravitational waves. This required measuring a strain of the order of $2\text{x}10^{-22} per \sqrt{Hz}$ in the 100Hz band which in turn required the highest quality lasers and optics available Ref [13]. Every distortion in the wavefront reduces the sensitivity the instrument and a huge amount of effort has gone into mapping the wavefront and correcting distortions. These distortions are a problem because they are indistinguishable from wavefront movements caused by changes in the positions of the test masses and couple into the sensitive measurements of the interference fringes.

LIGO employs mode matching cavities, Hartmann sensors and the very best optics supported by elaborate heating systems to monitor and maintain the quality of its laser beam wavefronts. LISA, the Laser Interferometer Space Antenna, will face a different set of challenges. Planned for launch in 2034 Ref [14], it will comprise three satellites separated by millions of kilometres. While the signals arriving will be plane waves, because of the distances involved, they will be very faint yet accurate measurement of the angle of a plane wavefront will be essential if the alignment between the three satellites is to be maintained.

1.3 The Difficulty with Wavefront Sensing

Mapping a wavefront requires measuring the phase at different points across its surface. Accomplishing this presents us with three interrelated problems.

1. The information required is encoded on a signal which oscillates at $\sim 3\text{x}10^{14}$Hz, far faster than the capabilities of any detection device currently available,
2. Phase cannot be measured directly, a photodiode measures power, and
3. A photodiode responds to the total photons striking it within the measurement cycle and averages the results to provide the power incident on its entire surface, there is no information about spatial distribution.

The first problem has been solved by beating the signal beam with a local oscillator of similar but different frequency, a process called heterodyning. The beat note produced is the difference between these two frequencies and is selected to be in the radio frequency range so can be read by a photodiode. Importantly this beat note carries all the phase information from the original signal beam.

The second problem is addressed using In-phase and Quadrature Demodulation (IQ demodulation) which uses a process of demodulation, integration and some trigonometry to recover the phase information (see Chapter 2). The phase of the carrier beam is encoded onto the phase of the beat note.

The last problem is typically solved by either directing the beat note to a segmented photodiode or a grid of holes (a Hartmann Sensor) to resolve spatial complexity. A Hartmann

sensor works by monitoring how a laser beam propagates through a lattice of very regularly space holes. Light travels perpendicularly to its wavefront so minor distortions in the image of the lattice can be used to map the wavefront incident upon it. This work has been developed to a high level of precision by the University of Adelaide (see Ref [15] and Ref [16]) and is the basis of wavefront assessments made in LIGO. Both these techniques are well established and reviewed in Chapter 2.

1.4 A New Approach to Wavefront Sensing

This thesis explores a new way of sensing the spatial distribution of amplitude and phase across the wavefront of a propagating laser beam. It uses a spatial light modulator (SLM) to tag different spatial regions of the beam, a single diode to collect the resulting light and digitally enhanced heterodyne interferometry (DEHI) to decode the phase and amplitude information across the wavefront in real time. A spatial light modulator is a series of closely packed mirrors, typically arranged in a square, which are individually controllable. Each can be advanced or retracted by precise, small amounts, typically fractions of a wavelength, which means that a laser beam incident across the active surface can have parts of its wavefront advanced of retarded, see5.4.1.

The bulk of the work in this thesis concerns two experiments. The first uses spatial tagging and DEHI, to develop detailed maps of a laser beam wavefront and show that the decoding process can follow and correct wavefront variations over time. It also shows that phase variations can be imprinted onto and then decoded from a wavefront, again in near real time.

The second applies these techniques to optimising the transmission of a laser beam through a cavity by coupling to the fundamental mode of of that cavity. Laser light is reflected from two SLMs and then directed into a fixed spacer cavity. The transmitted beam is measured on a single diode, decoded and then fed into an algorithm which makes simultaneous adjustments to both SLMs to maximise cavity transmission. The experiment is extended to show how higher order modes can be selected and maximised. Finally there is discussion around the role that Artificial Intelligence might play in optimising a system containing several cavities.

1.5 The Structure of this Thesis

The structure of this thesis is shown below. It comprises nine chapters and is in five parts, figure 1.4. This is Part 1 - the Overview - which sets the scene and outlines the structure of the document.

Part 2 - the Background - comprises three chapters. The first, Current Wavefront Sensing Techniques, is a review of wavefront sensing as practised today. Specifically it examines wavefront sensing using segmented diodes and Hartmann sensors. As noted above, phase is complicated to measure because it cannot be measured directly. Segmented diodes are used to make measurements at different parts of the wavefront and then local variations are demodulated to reveal both the phase and amplitude for a given segment. The mathematics behind this demodulation is explored in some detail because it is the foundation of Digitally Enhanced Heterodyne Interferometry, the technique employed in these experiments. Finally the chapter recognises that both these techniques only measure the structure of the wavefront and cannot correct it. It includes a brief summary of how the

information from the Hartmann sensors is used to make minute adjustments to the test masses (40kg mirrors at each end of each arm of the interferometers).

Chapter 3 is a detailed explanation of Digitally Enhanced Heterodyne Interferometry (DEHI). This technique was pioneered by Professor Daniel Shaddock and is the basis of much of the work we do in the Gravitational Wave and Space Groups at the Australian National University. In particular, the demodulation techniques involved are the basis for both the experiments presented here.

Chapter 4 contains a brief review of cavity dynamics and control and is the foundation for the experimental work reported in the next Chapter 6. It can be skipped by those familiar with cavities, cavity locking and the creation of higher order modes.

Part 3 - Experimental Work comprises two chapters. Chapter 5, Wavefront Sensing and Correction, outlines the work completed in 2015 and 2016 and published in Ref [17]. We report the experiment in detail, in particular how DEHI can be used in conjunction with a spatial light modulator and a single diode to develop a detailed map of a laser beam wavefront and follow variations in this wavefront as they occur. We review the noise characteristics of this new technique compared with sensing using segmented photodiodes. We then extend extend this discussion to show how the technique has been used to not only map a laser beam wavefront but make corrections to it, offering a new approach to wavefront detection and correction.

Chapter 6, Controlling Wavefronts in Real Time, shows how dithering techniques can be used to reshape a laser beam wavefront so that they are better mode matched to a fixed spacer cavity. It uses two SLMs to show a significant increases in the transmitted power of the fundamental mode. It then extends the experiment to show how other modes can be preferentially selected and transmitted. It also estimates the transfer function of the experimental set up and discusses how this might be improved.

The final part, Chapter 7, Future Work - considers how the work of this thesis might be extended to automate fully the locking of an optical cavity, use artificial intelligence to lock systems involving multiple cavities, and replace the SLMs with optical phased arrays to increase the speed and performance of laser beam reshaping.

The thesis concludes with a list of references and an appendix which contains details of the LabVIEW and Matlab code used to accomplish the results.

1.6 Publications

Below is a list of publications resulting from the work in this thesis.

Published

D.T. Ralph, P.A. Altin, D.S. Rabeling, D.E. McClelland and D.A. Shaddock, *"Interferometric wavefront sensing with a single diode using spatial light modulation"* Appl. Opt. **56**(8), 2353-2358 (2017).

In Preparation:

D.T. Ralph, P.A. Altin, D.S. Rabeling, D.E. McClelland and D.A. Shaddock, *"Automatic mode matching of a Fabry-Pérot cavity with a single photodiode and spatial light modulation"*

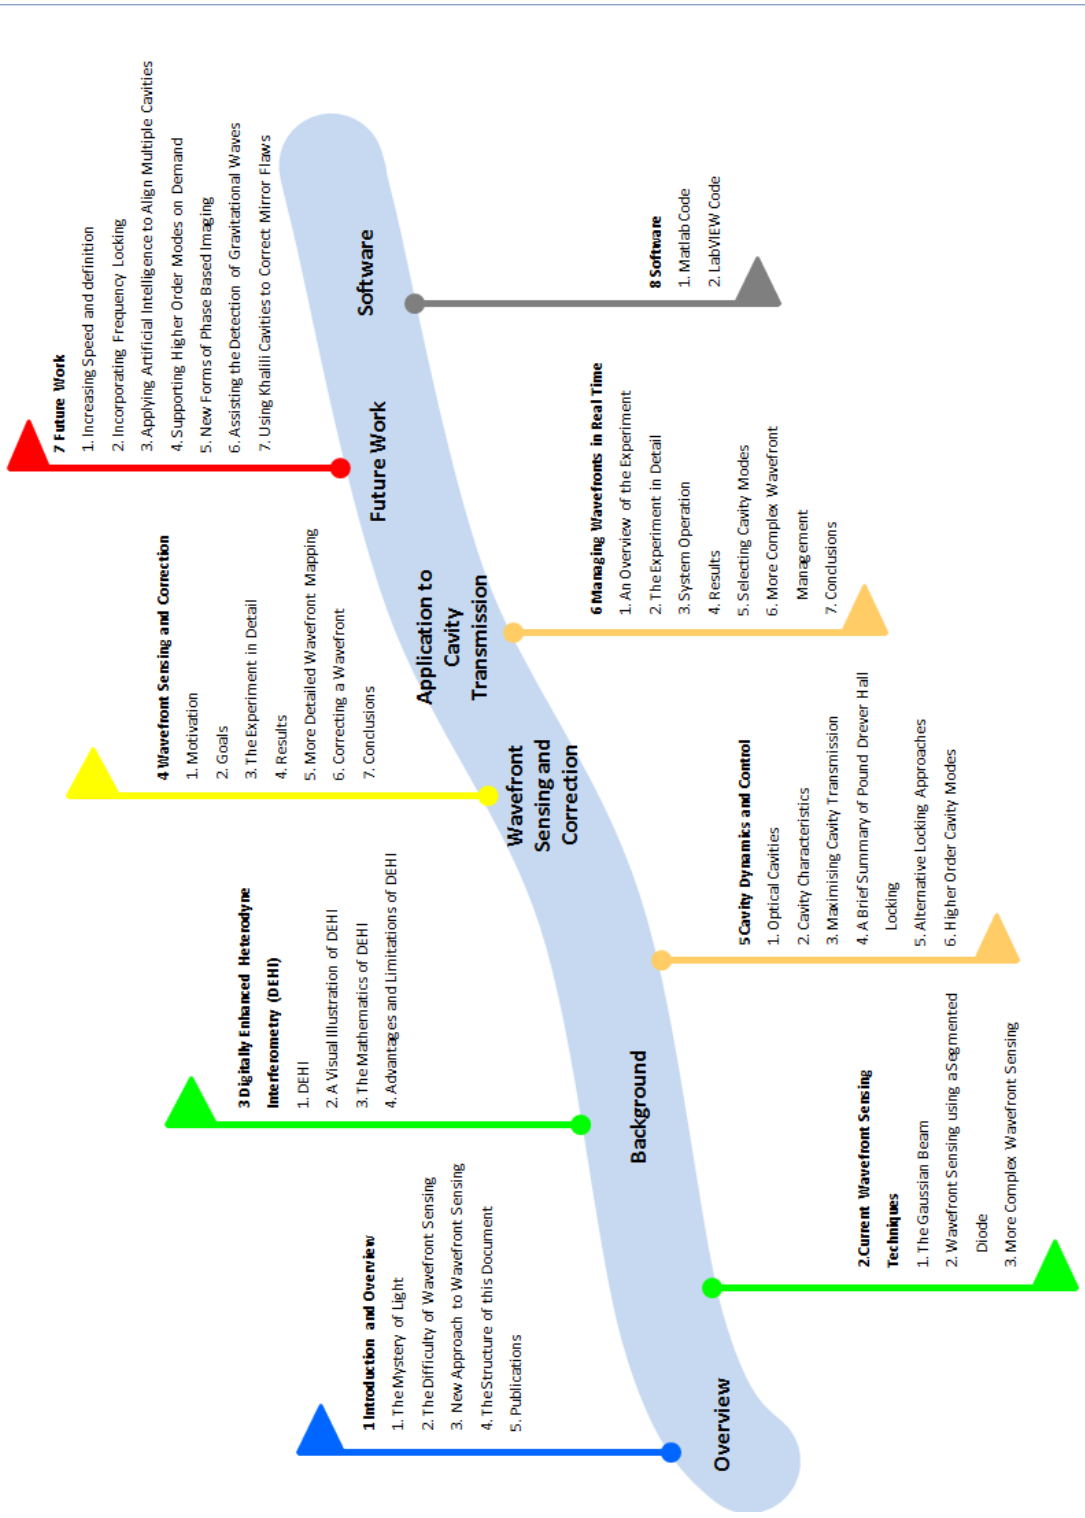

Figure 1.4: The structure of this thesis

Current Wavefront Sensing Techniques

This chapter looks at the traditional methods of wavefront sensing and discusses their advantages and limitations. After a brief review of Gaussian beam dynamics, it explains in detail how sensing is undertaken using multi-segment diodes, the Ward Technique and Hartmann Sensors.

2.1 The Gaussian Beam

A laser beam propagating from a high quality source as a narrow beam can be represented by the paraxial approximation, a solution to the Helmholtz wave equation. This assumes that z, the axis along which the beam propagates, is much larger than x and y, the axes across the beam which carry the transverse electric and magnetic fields, and that the beam radius at a beam waist is much larger than the wavelength. The beam is a Gaussian beam which means that it transverse electric and magnetic fields have amplitude profiles that are defined by the Gaussian function and in turn produce the fundamental Gaussian mode TEM_{00}. The equation which defines the electric field of such a beam is, to a very good approximation, a solution to the paraxial Helmholtz equation and is (for polarisation in the x direction and propagation in the z direction) Ref [18] and Ref [19]:

$$E(r,z) = E_0\hat{x}\frac{w_0}{w(z)}.e^{\left(\frac{-r^2}{w(z)^2}\right)}.e^{\left(-i(kz+k\frac{r^2}{2R(z)}-\psi(z))\right)} \tag{2.1}$$

where

- r is the radial distance from the centre axis of the beam,
- z is the axial distance from the beam's focus (or "waist"),
- i is the imaginary unit,
- $k = 2\pi/\lambda$ the wave number (in radians per meter) for a wavelength λ
- $E_0 = E(0,0)$, the electric field amplitude (and phase) at the origin at time 0,
- $w(z)$ is the radius at which the field amplitudes fall to $1/e$ of their axial values (i.e., where the intensity values fall to $1/e^2$ of their axial values), at the plane z along the beam,

- $w_0 = w(0)$ is the waist radius,
- $R(z)$ is the radius of curvature of the beam's wavefronts at z, and
- $\psi(z)$ is the Gouy phase at z, an extra phase term attributable to the phase velocity of light.

The paraxial approximation can be used to describe the transverse, spatial dependence of the wave as it propagates. It applies to most situations in precision metrology and, specifically, to the work in this thesis. (There are occasions where a more detailed analysis is required such as that provided by Davis Ref [20], Lax et al Ref [21] and Chen et al Ref [22]). It describes a wave which diverges as the beam propagates and that the wavefront evolves from flat at the waist, where the beam is narrowest, to curved and then near flat again far from the waist, figure 2.1. The maximum curvature occurs at the Rayleigh Range which is given by:

$$z_R = \pi . w_0^2 / \lambda \tag{2.2}$$

where z_R is the Rayleigh Range, w_0 is the waist size at $z = 0$ and λ is the wavelength of the light, Ref [23].

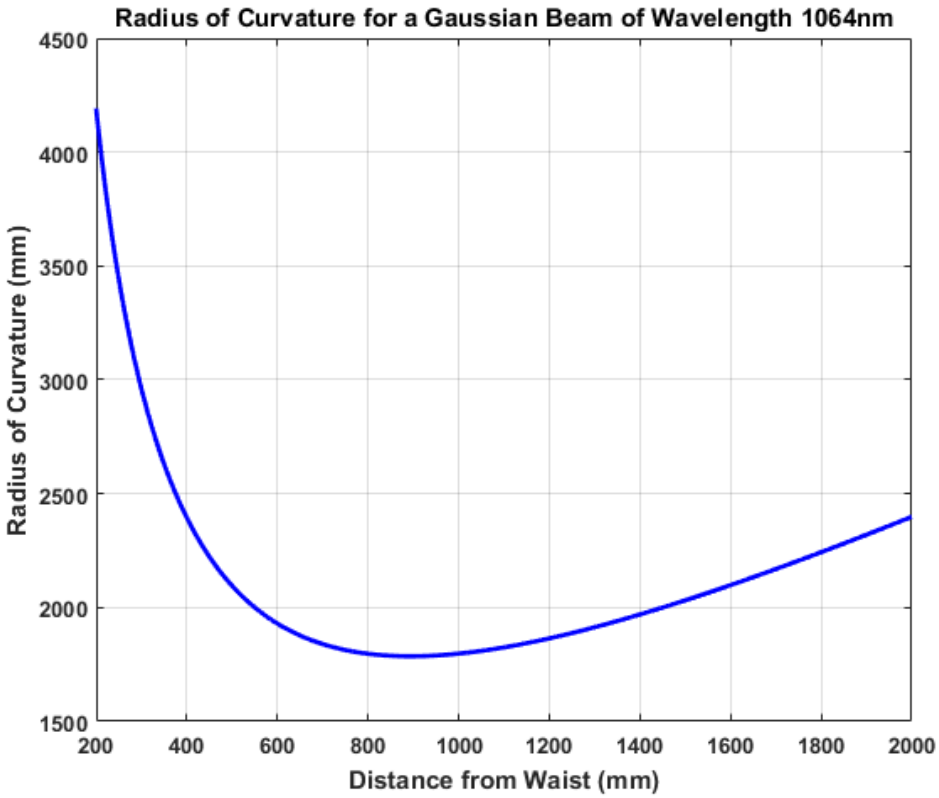

Figure 2.1: The evolving radius of curvature of a Gaussian Beam from a waist of 0.55mm. The finite width of the laser beam at its waist means the the propagating wavefront begins flat, with infinite curvature, and then begins to a become progressively more curved to reach a maximum curvature at the Rayleigh Range (893mm) before flattening out again in the far field. The radius of curvature is symmetrical about the waist.

The radius of curvature of the beam is given by Ref [23]:

$$R_z = z\left(1 + \left(\frac{z_R}{z}\right)^2\right) \tag{2.3}$$

$$= z\left(1 + \left(\frac{\pi w_0^2}{z\lambda}\right)^2\right) \tag{2.4}$$

and the Gouy phase by Ref [23]:

$$\psi(z) = arctan\left(\frac{z}{z_R}\right) \tag{2.5}$$

Gouy phase arises from the limited spatial extent of the beam in the x and y directions and is a measure of how much the actual curvature of a wavefront diverges from a plane wave at a given value of z. It is a useful tool for determining the placement of optics which are required to act primarily one of two orthogonal attributes, for example beam tilt and offset. It changes by $\pi/2$ from the waist to the far field reaching a value of $\pi/4$ (45°) at the Rayleigh Range.

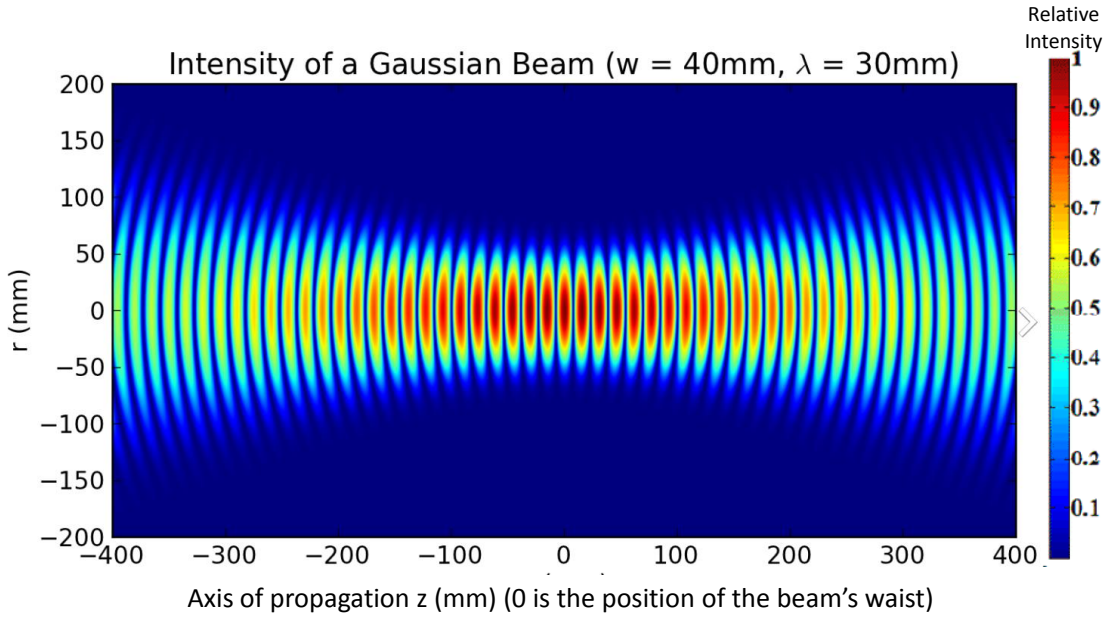

Figure 2.2: The wavefront of a propagating laser beam derived using the the paraxial approximation. The wavefront is flat at the waist but spreads and curves as it progresses resulting in a wavefront that is curved rather than flat. This image was reproduced from Ref [24].

Finally the width of the beam at any point z is given by w_z where:

$$w_z = w_o\sqrt{1 + \left(\frac{z}{z_R}\right)^2} \tag{2.6}$$

These parameters are summarised in figure 2.3.

2.1.1 An example for a 1064nm Laser

As an example consider the 1064 nm laser beam incident upon the cavity used in the second experiment reported in this thesis. This cavity was a fixed spacer, plano-concave cavity,

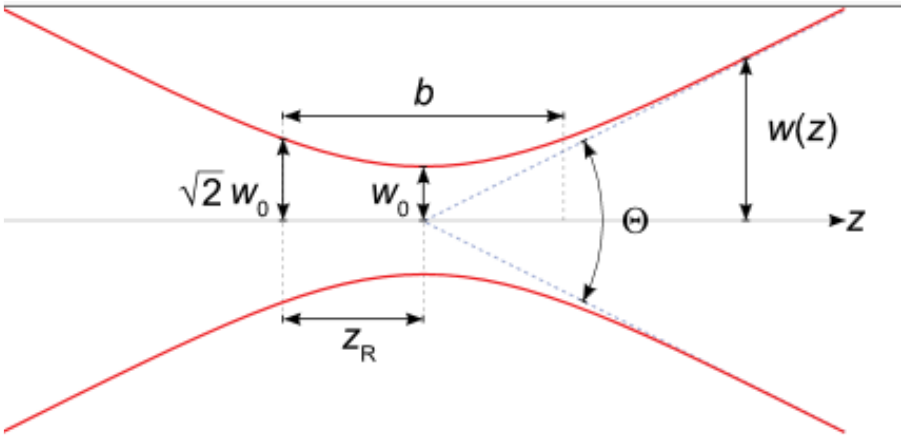

Figure 2.3: The diagram shows the Gaussian beam width w(z) as a function of the distance z along the beam where w_0 is the beam waist and z_R is the Rayleigh range. Θ is the total angular spread and b is the depth of focus and is twice the Rayleigh range, Ref [18]

.

300 mm in length with the concave mirror having a radius of curvature of around 3 m. A 1064nm laser beam resonating perfectly within this cavity would need the characteristics shown in the figure (2.4). In particular it would require a waist of .55 mm positioned perfectly on the plane mirror. A beam mode matched to this cavity would expand within the cavity with its wavefront curving progressively as it progressed. 30 cm from the waist, still well within the Rayleigh range, all parts of the wavefront would meet the concave mirror at the other end of the cavity at normal incidence with the wavefront precisely matching the radius of curvature of that mirror so that the beam is reflected back on itself until it hits the plane mirror, whereupon the whole process would repeat. This would continue indefinitely with more and more laser power collecting in the cavity. In practice, however, the system quickly reaches a steady state where losses equal the new energy being supplied. Losses arise because of mirror reflectivities that are less than 100% and wavefront aberrations from imperfect optics.

The figures in the charts in figure (2.4) show how such a beam would evolve in free space and were derived using beam.m, part of the suite of beam tools developed by Alexei Ourjoumtsev from the Ecole Normale Superieure in Paris in 2003, see Ref [25].

The wavefront is flat at the waist with a Gouy phase of 0°. The beam expands through the Rayleigh Range (893 mm) where the Gouy Phase is 45° and the curvature is at its greatest, actually spherical. In the far field the wavefront approaches flat again and the Gouy phase is 90°. These results are presented graphically in figure 2.5, created using the beam.m model.

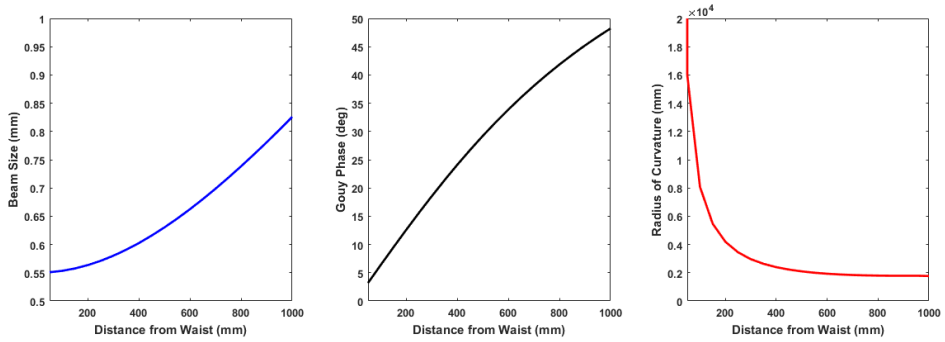

Figure 2.4: Values for the width, Gouy phase and radius of curvature for a 1064 nm laser beam propagating from a waist of 0.55 mm at position 0 mm. The Rayleigh range of the beam is 893 mm where the Gouy Phase is 45.

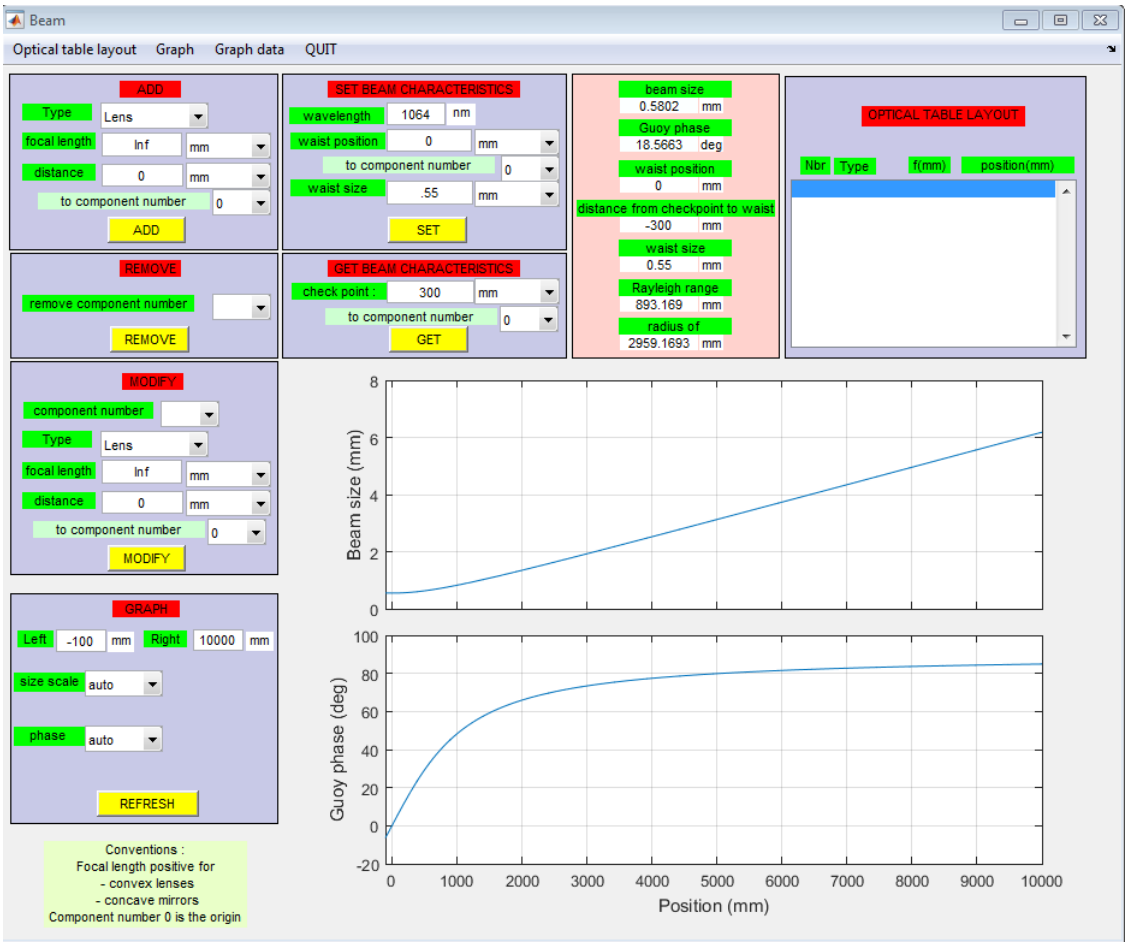

Figure 2.5: An extract from the beam modelling program showing the evolutions of a 1064nm laser beam from a waist of .55 mm.

2.2 Determining the Wavefront Angle Using Segmented Photodiode

2.2.1 Differential Sensing and the Ward Technique

In 1994, a group at the University of Glasgow developed a technique to measure wavefront tilts and beam offsets by using a pair of photodiodes to examine the reflected signal from a plano-concave cavity. The theoretical discussion of the approach is outlined Ref [26], while Ref [27] reports details of the experimental verification. They understood that highly sensitive interferometers would need ongoing alignment but that the use of mechanical dithering to create an error signal might introduce an undesirable level of noise or could be impractical in some system such as satellite based interferometers, Ref [26].

Maintaining alignment requires continual control of four independent errors - vertical tilt and offset and horizontal tilt and offset. The Glasgow group recognised that if two beams were slightly tilted with respect to each other, even though the average phase difference across the combined wavefront (ω) might be constant, there would be a difference in phase between the two sides of the beam ($\omega - \phi$ in one half and $\omega + \phi$ in the other) and that this could be detected, using suitable demodulation, by a segmented diode. They also understood that beam misalignment could be transformed into tilt using a lens and that consequently they could derive the four independent error signals with two quadrant photodiodes and some lenses. The group was lead by Dr. Henry (Harry) Ward and the technique now commonly bears his name. It has become the standard technique for the maintenance of alignment between satellites for space based missions.

2.2.2 Measuring Phase with a Segmented Diode

Traditional wavefront signal processing using a segmented photodiode is a two stage process:

- First the light is interfered with a local oscillator reference light of slightly different frequency to transfer variations in the signal beam to a beat note visible to the photodiode - a process called heterodyning.
- The resulting signal is then split into two channels, multiplied by the sine and cosine of the same frequency as the heterodyne signal to recover the in-phase and quadrature components. These can be divided to recover the phase of the wavefront at a given point - a process called In-phase Quadrature Demodulation or IQ Demodulation.

The phase of the heterodyne beat note wavefront is recovered as an average value for each segment of the photodiode and assumed to be linear between segments, providing an approximation of deviations across the wavefront. This heterodyne beat note carries the phase information of the original signal beam. Note that this phase measurement is differential, relative measurement and made with respect to a reference, in this case the phase of the local oscillator.

2.2.3 Heterodyne Signal Processing

Laser light with a wavelength of 1064 nm has a frequency of $2.82\text{x}10^{14}$ but a photodiode measures power which is the electric field multiplied by its complex conjugate, equation

2.7 and this is a constant.

$$P_d(t) = \tilde{E}.\tilde{E}^* = E.e^{-i\omega_l t}.E.e^{i\omega_l t} = E^2 \tag{2.7}$$

In 1901, a Canadian inventor called Reginald Fessenden invented a signal processing technique which used the fact that, when signals of two frequencies are combined, they produce a beat note in a different frequency range, a process called heterodyning Ref [28]. While first applied in the radio frequency domain, the process works just as well with light. Importantly, the beat frequency captures the differential phase between the two original signals so, if the signal of interest beats with a pure sine wave, called the local oscillator, the variations in the heterodyne signal capture the phase information from the original signal measured with respect to the local oscillator and imposed on a carrier frequency which can be resolved by a photodiode.

Consider a field at a detector E_d comprising a local oscillator E_l and a coded signal E_s. The total field at the detector is given by the sum of the individual fields.

$$\tilde{E_d} = \tilde{E_s} + \tilde{E_l} \tag{2.8}$$

Photodiodes measure power and the total power at the detector, P_d, is this field multiplied by its complex conjugate.

$$\begin{aligned} P_d(t) &= (\tilde{E_s}^* + \tilde{E_l}^*).(\tilde{E_s} + \tilde{E_l}) \\ &= \tilde{E_s}^*.\tilde{E_s} + \tilde{E_l}^*.\tilde{E_l} + \tilde{E_s}^*.\tilde{E_l} + \tilde{E_l}^*.\tilde{E_s} \\ &= \tilde{E_s}^*.\tilde{E_s} + \tilde{E_l}^*.\tilde{E_l} + \tilde{E_s}^*.\tilde{E_l} + (\tilde{E_s}^*.\tilde{E_l})^* \end{aligned} \tag{2.9}$$

But, in general, for a complex number, $(a+ib)+(a+ib)^* = (a+ib)+(a-ib) = 2a = 2Real[a+ib]$ so, if the electric field of the signal beam is given by $\tilde{E_s} = E_s e^{-i\omega t}$, then:

$$\tilde{E_s}.\tilde{E_s}^* = E_s e^{-i\omega t}.E_s e^{i\omega t} = E_s^2(e^0) = E_s^2 \tag{2.10}$$

Combining the two fields

$$P_d(t) = \tilde{E_s}^2 + \tilde{E_l}^2 + 2Re[\tilde{E_s}^*.\tilde{E_l}] \tag{2.11}$$

Now $\tilde{E_s}^2$ and $\tilde{E_l}^2$ are constants so appear as a direct current (DC) signal. The alternating current (AC) signal is:

$$P_{AC}(t) = 2Re[(\tilde{E_s}^* \cdot \tilde{E_l})] \tag{2.12}$$

Now the local oscillator, $\tilde{E_l}$, can be represented by:

$$\tilde{E_l} = E_l.e^{-i\omega_l t} \tag{2.13}$$

where ω_l is the angular frequency of the local oscillator and t is the elapsed time. Similarly the signal, E_s can be represented by:

$$\tilde{E_s} = E_s.e^{-i\omega_s t + \phi_s} \tag{2.14}$$

where ω_s is the angular frequency of the signal and ϕ_s is the phase difference between the signal and the local oscillator at time $t = 0$. The power produced by a photodiode is

proportional to the electric field falling on it multiplied by the complex conjugate of that field so substituting these values we get:

$$\begin{aligned} P_{AC}(t) &\propto 2Re[(\tilde{E_s}^* \cdot \tilde{E_l})] \qquad (2.15)\\ &\propto 2Re[E_s.e^{i\omega_s t+\phi_s}.E_l.e^{-i\omega_l t}] \\ &\propto 2Re[E_s.E_l e^{i(\omega_l-\omega_s)t-\phi_s}] \\ &\propto 2E_sE_lRe[e^{i(\omega t-\phi)}] \\ &\propto 2E_sE_l\cos(\omega t-\phi) \qquad (2.16)\end{aligned}$$

The crucial point here is that combining the signal with a local oscillator of a slightly different frequency has created a beat note which includes a cross term varies in power at a frequency the photodiode can detect. Importantly, this beat note carries all the phase information, ϕ, embedded in the original signal beam. This is heterodyning.

IQ Demodulation

Recovering this phase is typically accomplished using In-phase Quadrature (IQ) Demodulation. The detected beam, AC_d, is passed through and Analogue to Digital Converter (ADC) to the central processing unit and split into two and multiplied by sine and cosine of a wave with the same frequency as the beat frequency, ω. The resulting signals are $2E_sE_l\cos(\omega t-\phi).\cos(\omega t)$ and $2E_sE_l\cos(\omega t-\phi).\sin(\omega)$. These can be solved further using the identities:

$$\begin{aligned}\cos(a).\cos(b) &= 1/2(\cos(a+b)+\cos(a-b)) \\ \sin(a).\cos(b) &= 1/2(\sin(a+b)-\sin(a-b))\end{aligned}$$

The signal multiplied by $\cos(\omega t)$ becomes:

$$\begin{aligned}\cos(\omega t-\phi).\cos(\omega t) &= 1/2((\cos(\omega t-\phi+\omega t)+\cos(\omega t-\phi-\omega t)) \\ &= 1/2(\cos(2\omega t-\phi)+\cos(-\phi)) \qquad (2.17)\end{aligned}$$

while the signal multiplied by $\sin(\omega t)$ becomes:

$$\begin{aligned}\cos(\omega t-\phi).\sin(\omega t) &= 1/2((\sin(\omega t-\phi+\omega t)-\sin(\omega t-\phi-\omega t)) \\ &= 1/2(\sin(2\omega t-\phi)-\sin(-\phi)) \qquad (2.18)\end{aligned}$$

The final stage in the process is to average or sum the signal over a period that is long compared to the angular frequency ω and, ideally, an integer multiple of it. The term $1/2(\cos(2\omega t-\phi))$ and $1/2(\sin(2\omega t-\phi)$ average to near zero (or precisely zero if n covers an integer number of angular cycles). The parts containing ϕ alone compound as follows:

$$\begin{aligned}\sum_{t=1}^{n} 2E_sE_l.1/2(\cos(2\omega t-\phi)+\cos(-\phi)) &= -n.E_sE_l.\cos(-\phi) \\ &= -n.E_sE_l.\cos(\phi) \qquad (2.19)\\ \sum_{t=1}^{n} 2E_sE_l.1/2(\sin(2\omega t-\phi)-\sin(-\phi)) &= -n.E_sE_l.\sin(-\phi) \\ &= n.E_sE_l.\sin(\phi) \qquad (2.20)\end{aligned}$$

ϕ can not be recovered from either one of these results directly because the amplitude of the signal beam may vary over time However, if the two expressions are divided as follows, these errors are common and divide away so that ϕ can be recovered using the arctangent.

$$\begin{aligned}\phi &= -tan^{-1}\left(\frac{n.E_sE_l.\sin(\phi)}{-n.E_sE_l.\cos(\phi)}\right)\\ &= tan^{-1}\left(\frac{n.E_sE_l.\sin(\phi)}{n.E_sE_l.\cos(\phi)}\right)\end{aligned} \tag{2.21}$$

Interestingly the amplitudes and length of summation divide away making the process extremely robust to changes in laser power which are long compared to the angular frequency, ω.

Mapping the Wavefront

If a quadrant photo-detector is used to collect the signal beam and 4 signals are returned for processing, the above approach can be applied to determine the relative average phase of each quadrant of the detector and a basic map of the wavefront can be estimated as a surface linking these 4 points. This is a good approximation for satellite based interferometers where the distances involved are large and the wavefronts are close to plane waves by the time the reach a distance detector (although not exactly, see Ref [29]). The process also provides enough information to effect first order corrections to satellite orientation in the form or tip and tilt.

Advantages and Disadvantages of Conventional Wavefront Sensing

Using IQ demodulation to work out the point phase at different points on a wavefront is a robust and well proven technology. It is relatively easy to implement and independent of short term variations in laser power. It's main limitations are that:

- it can only provide as many data points for mapping the wavefront's surface as the photodiode has segments (typically four) so only provide a linear approximation of the wavefront, and
- signals from each photodiode segment return to the central processing unit through separate channels and ADCs, allowing small, uncorrelated errors which are indistinguishable from phase variations.

Furthermore:

- the technique provides no mechanism to correct any imperfections detected in a wavefront, if and when this is important, and
- the frequency of the measurement is determined by the length of the Hadamard code used rather than the number of individual point measured which limits the frequency at which measurements can be made (although, in practice, this is rarely a limitation). (This disadvantage also applies to our new technique.)

2.3 More Complex Wavefront Sensing Using Hartmann Sensors

Description

The Hartman sensor provides an alternative and now very accurate method of sensing a wavefront Ref [15], Ref [16] and Ref [30]. It was developed in 1900 by Johannes Franz Hartmann Ref [31] and works because a light wave always travels perpendicularly to its wavefront. A Hartmann sensor is a screen of very precisely arrayed holes through which a beam is shone onto a Charge-Coupled Device (CCD) camera figure 2.6).

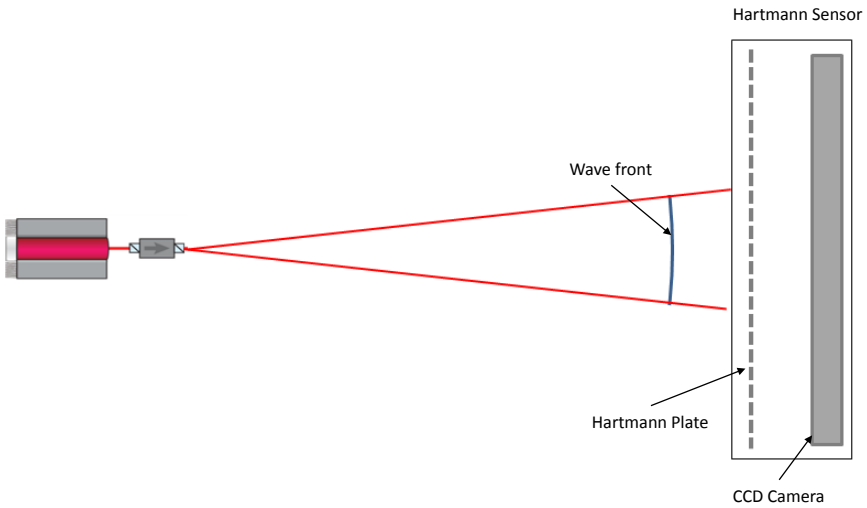

Figure 2.6: The Hartmann sensor allows a user to sense the shape of a wavefront from the distortion caused in the image of a regular array of holes projected onto a CCD camera.

If the wavefront striking the screen is a plane wave, the image of the holes on the CCD camera is a precise recreation of the screen. If the wavefront is curved or irregular, the image is distorted, figure 2.7). Algorithm can then be used to infer the shape of the wavefront from these distortions Ref [15].

Hartmann sensors are being used in the most demanding sensing application on the planet, the Laser Interferometer Gravitational Wave Observatory (LIGO). LIGO comprises two Michelson interferometers with 4 kilometre arms, separated by about 4,300 km. The two interferometers are one instrument as it is only by comparing signals received at both and matching them (with appropriate delays) that observers can be certain that any signal is a gravitational wave and the product of not some local, unsuppressed noise source. Further it allows observers to estimate the direction from which the gravitational wave originated and this can be enhanced by adding data from the Virgo interferometer in Europe.

In each observatory, a high quality, high power laser beam is split, injected into the arm and then recombined to look for tiny changes in relative arm length. The precision of the measurement is breathtaking. LIGO can measure a strain of 10^{-22} which is equivalent to measuring a change the width of a human hair in the distance between Earth and the second nearest star (that is, not the Sun)! Obviously any source of extraneous noise can

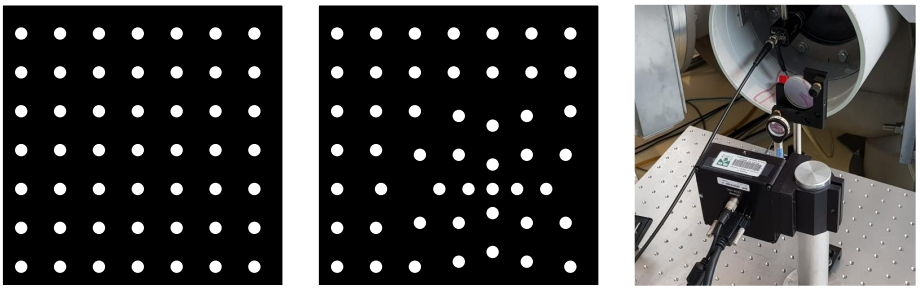

Figure 2.7: The picture on the left represents the regular image produced by a Hartmann sensor sensing a plane wave. The middle picture shows how this pattern changes as the wavefront changes; these changes can be used to infer the wavefront's shape. On the right is a photograph of one of the Hartman sensors being installed in the Gin Gin Gravitational Wave facility in Western Australia.

ruin the measurements and one noise source that must be managed is the distortion of the large mirrors (the test masses) at the end of each cavity by the large amounts of heat that the circulating lasers create. This is accomplished by the Thermal Compensation System (TCS) which uses Hartmann sensors to infer thermal distortion from wavefront distortion and ring heaters to try to correct it.

LIGO has put new demands on the accuracy of Hartmann sensors and the team at the University of Adelaide reported in 2016 that they had achieved a wavefront sensitivity of 1.35 nm RMS Ref [32], within the tolerances required for LIGO.

Currently LIGO requires that the Hartmann sensors use laser light of a different frequency (532nm and 800 to 840 nm) from that used in the interferometer (1064nm) and some scatter loss from defects may be wavefront dependent. Like any sensor, the Hartmann sensors only provide measurements can cannot actively correct any wavefront aberrations.

Nonetheless they represent a remarkable achievement and were able recently to detect wavefront distortions in one of the test masses that were not visible to human observers, see Ref [33].

Great care is taken to ensure that the light injected into the main interferometer arm in LIGO is coherent, monochromatic and Gaussian. Wavefront distortions that arise therefore, tend to come from two sources:

1. Defects in the optics, such as the spot found on one of the test masses, or
2. Distortions of the optics caused by expansion and absorption due to the high optical powers resident within the cavity.

There is little that can be done to correct defects of the first type other than the reworking or replacing the optic. Defects of the second type can be mitigated using the active half of the Thermal Compensation System (TCS) which attempts to compensate for low order distortions in the test masses using ring heaters, Ref [32].

The Shack-Hartmann Sensor

The Shack-Hartmann sensor is a variant of the Hartmann sensor which grew out of a demand for better satellite imaging during the Cold War, see Ref [34] and Ref [31]. The

traditional Hartmann sensor suffered from two disadvantages when it came to satellite imaging, see Ref [35]:

1. The sensor was in line with the imaging beam, blocking it, and
2. The sensor blocked a large proportion of the light falling on on it, allowing only that fraction passing through the holes in the sensor to create an image.

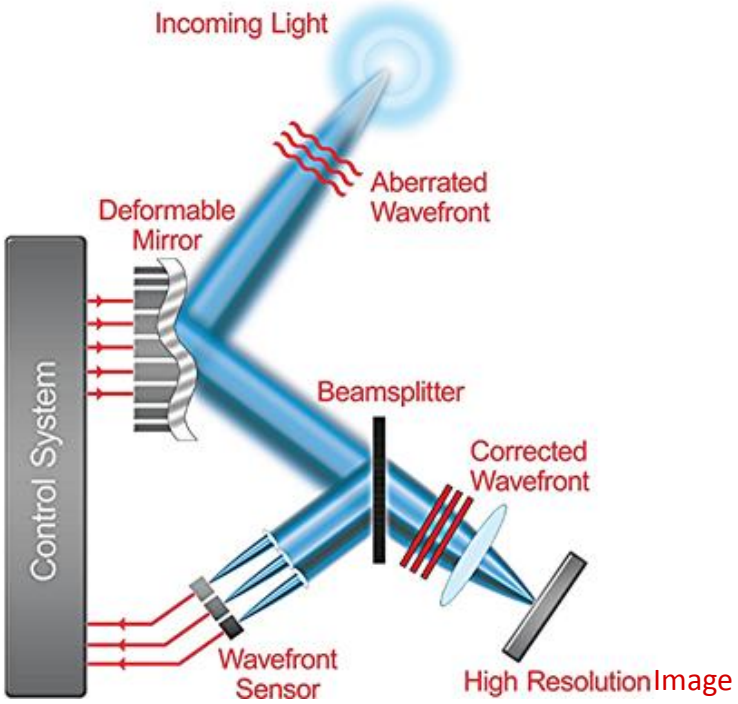

Figure 2.8: The Shack-Hartmann Sensor. A variant of the Hartmann sensor, the Shack-Hartmann sensor uses a beam splitter to divert a small proportion of the incoming light onto a lenslet array which intensifies the pattern of spots it produces. The bulk of the beam is free to pass through to a CCD camera and produce a high resolution image. The control system analyses displacements in the positions of the spots and produces an error signal. This in turn can adjust a deformable mirror which then adjusts the incident wavefront in real time.Of course, the feedback from the control system must be faster than the rate of change aberrations being corrected. This image is taken from Ref [36]

After an approach from the military, Dr Aden Meinel and the Optical Sciences Center (OSC) at the University of Arizona devised an approach which involved tapping off a small part of the image beam, using a beam splitter, and intensifying the images of spots produced using a lenslet array. This left the bulk of the beam free to produce an image while the lenses focused the tapped portion into a series of spots. It was Roland Shack, whom Dr. Meinel had recruited to OSC, who was given the task of determining whether this technique was feasible, see Ref [35].

Figure 2.8 shows a typical adaptive optics system based on a Shack-Hartmann sensor using in astronomy. The incoming light strikes a deformable mirror (initially set flat) and then a beam splitter. A small proportion is diverted to the Shack-Hartmann sensor and produces an image of an array of dots. Their deviation from the expected regular array is used by the control system to create an error signal which is fed back to the deformable mirror, which in turn modifies the next part of the incoming light. Provided the response of the

control system is faster than changes due to turbulence in the atmosphere, the main part of the beam passes to the CCD camera where it produces a corrected, high resolution image, see Ref [36].

The Shack-Hartmann sensor is now a well proven and well used part of adaptive optics being employed for everything from reducing distortions due to atmospheric turbulence for astronomical observations to making detailed assessments in ophthalmology, see Ref [37] and Ref [31]. Figure 2.9 shows the type of correction possible. The figure on the left is the uncorrected image, while that on the right is the corrected image. This figure is reproduced from Ref [37].

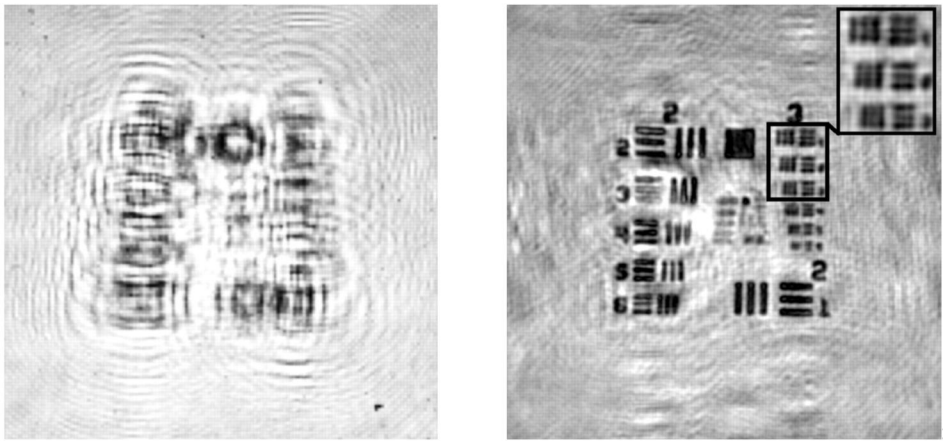

Figure 2.9: A demonstration of image Correction Using a Shack-Hartmann Sensor. The figure on the left is the uncorrected image, while that on the right is the corrected image. This figure is reproduced from Ref [37].

2.4 The Ongoing Demand for Wavefront Correction

There is no doubt that wavefront aberrations will continue to be a problem in future gravitational wave detectors. In a comprehensive review of the challenges facing the next gravitational wave detectors, Ref [38] (and see also Ref [39]), Section 4 recognises that the *'operation and sensitivity of interferometric gravitational waves detectors can be strongly limited by wave-front aberrations in the core optics.'* and notes that *'these aberrations are due to defects in the optics (surface figure errors and refraction index inhomogeneity) and to the rise of thermal effects and can be compensated for by generating proper optical path length corrections.'*

It proposes that the solution is to induce in the mirrors an optical path length increase equal but opposite to the distortions and that suggests that the answer is to use localised heating to reshape optical components. It also recognises that parametric instabilities will need to be better controlled in future detectors if they are not to cause interferometer locking to fail, (see Ref [40], Ref [41] and Ref [42]) and will require wavefront sensing to assist management of this. A Hartmann sensor is currently being installed at the Gingin test facility to assist with this. The need for high power lasers (100 W of optical power) operating at a single frequency and in a single spatial mode (TEM_{00}) is also identified as a future requirement to help reduce the shot noise, see Ref [40], meaning that the identification and sustaining of a single mode will be increasingly important. This relates to the proofs of concept presented in Chapter 6.

2.5 Summary

This chapter has reviewed how laser light propagates, to a close approximation, as a Gaussian beam. It describes the Gaussian beam parameters and how a wavefront can be sensed using either segmented diodes or Hartmann sensors. These techniques are now very accurate and well proven. The definition that segmented diodes can provided is limited by the number of sensors that can be place in a wavefront and any errors that arise in individual analogue channels and Analogue to Digital Converters (ADCs) as the signals are sent back to the central processor for demodulation. Hartmann and Shack-Hartmann sensors can provide a much more detailed picture of a wavefront and are a critical part of the Thermal Compensation System in LIGO and adaptive optics generally. Neither technique can make corrections to the wavefront.

Chapter 3

Digitally Enhanced Heterodyne Interferometry

This chapter provides a detailed review of Digitally Enhanced Heterodyne Interferometry(DEHI). It is a technique which uses pseudo-random codes (PRN codes) to tag either temporal or spatial regions of a laser beam so that they can be isolated and analysed in post processing. It was developed by Professor Daniel Shaddock Ref [43] and pioneered by our group at ANU and is the tool underlying all the experimental work presented in this thesis. It is still relatively new to the photonics community so this chapter provides a detailed explanation of technique and how it can be used to measure the phase and amplitude of a wavefront.

DEHI builds on traditional heterodyne sensing which was outlined in Chapter 2. This chapter begins with a qualitative description of DEHI and a discussion of the differences between temporal and spatial code tagging and how this affects code selection. Hadamard codes have been used for the experiments presented in this thesis. They are explained in detail and compared with m-Sequences. The chapter concludes with a mathematical explanation of how these codes can be used to first tag and then isolate spatial regions to facilitate development of a detailed wavefront map. The chapter concludes with a summary of the advantages and disadvantages of the technique.

3.1 Digitally Enhanced Heterodyne Interferometry

3.1.1 A Qualitative Overview

DEHI includes all the steps involved in traditional wavefront sensing with segmented diodes but adds two more. First it embeds codes onto different regions of the signal beam and second it uses these codes to separate these regions of the heterodyne signal in post processing so that a detailed picture can be recreated of its wavefront's phase and amplitude over time or space. Finally it no longer requires a quadrant photodetector to resolve different spatial regions but accomplish this using only a single photodetector.

The steps in DEHI are listed below:

1. The signal beam is phase modulated with a series of codes using either a spatial light modulator for spatial codes or an Electro-Optical Modulator for temporal codes.
2. The signal beam is then beaten with a local oscillator to create a heterodyne signal that oscillates at a frequency that can be detected by a photodiode (just as was the

case in traditional wavefront sensing - see equation 2.16).

3. The heterodyned signal is collected on a single photodiode and returned through an ADC for processing. Segmented photodiodes or Hartmann sensors are not required.

4. The signal is split into as many channels as there are codes. Each channel is decoded by multiplying it by the relevant code in phase with the original code. This facilitates either determining the delay of the entire signal (for temporal coding with m-Sequences) of isolating components (for spatial coding with Hadamard codes).

5. Each channel is then split into two for IQ demodulation. This was outlined in Chapter 2 and involves multiplication by sine or cosine, integration and division, see equation 2.21, to recover the phase or amplitude for that all or part of the signal.

There are two interesting points about this approach. First, even though the introduction of the codes would seem to add a considerable amount of noise to the signal this is not the case. The noise added by the codes is either completely removed (in the case of Hadamard code or reduced to almost zero (for m-Sequences) by the integration in step 5. Integration is, of course, a part of traditional wavefront sensing and is used to isolate the phase or amplitude of the desired part of the signal. Certainly the integration period required may have to be longer for long code sequences (often the case for m-Sequences) or slow code modulation (compared to the heterodyne frequency).

Second the order in which the demodulations take place does not matter. In the explanation given above the signal was first separated into as many channels as there are codes and demodulated using the codes before being split a second time for demodulation with $\sin(\omega)$ and $\cos(\omega)$. These steps can be reversed and, in practice, usually are as it is more efficient from a programming point of view.

3.2 Spatial and Temporal Encoding

The DEHI technique is useful for addressing a whole range of problems but in general they fall into two types - temporal encoding or spatial encoding. This is the first time it has been used for spatial encoding. Temporal encoding occurs when a code is applied to an entire signal which is later collected and demodulated with the same code. Finding a match involves sliding a copy of the code along the returning signal at different delays until it aligns with the code originally embedded within it. This is indicated by a large spike in the frequency domain amplitude where the frequency has been 'de-spread'. The delay between the embedded code and the demodulating code can be used to determine the time of flight. Tracking space debris, for example, and in particular calculating its range, works by reflecting an encoded signal off a piece of debris and then seeing how far the reflected signal has travelled using this demodulation process Ref [44]. Halving this distance gives the range. Typically one, long m-Sequence is used and the code must be long enough to determine range unambiguously. The technique has been developed to a point where it can can recover distances shorter than the length light travels during the duration of a single code and accommodate the effects of Doppler shifts on the range information.

Spatial tagging involves a different set of challenges. Here the task is to use codes to tag (modulation) and then separate (demodulation) different spatial regions of the same signal. The challenge is to be able to isolate and examine one part of the wavefront without contaminating the information with data from other regions. Codes are applied

to different areas simultaneously and the part of any sequence used need only be as long as the number of spatial segments being examined. The codes applied to different areas remain in phase throughout the process and the best solution is Hadamard codes.

In summary:

1. m-Sequences are best when a whole signal is to be compared to itself and a delay calculated. This is because m-sequences are autocorrelated.
2. Hadamard codes are best when regions of a signal are to be isolated and analysed independently so that a picture of the entire wavefront can be created. This is because Hadamard codes are cross correlated.

3.2.1 m-Sequences

m-Sequences are examples of Pseudo Random Noise codes. A Pseudo Random Noise code (PRN code) is one that has a spectrum similar to a random sequence of bits but is deterministically generated and periodic, see Ref [45] and Ref [46] but most usefully for practical applications Ref [47]. They have the following characteristics:

1. Each code in a sequence can have the value of 1 or -1.
2. Any m-Sequence multiplied element wise by itself out of phase and summed gives a value of -1, that is a discrete autocorrelation of -1 for any non-zero delay.
3. Any m-Sequence multiplied element-wise by itself in phase and summed gives a value equivalent to the length of the sequence.
4. Any m-Sequence multiplied element wise by another of the same length and summed produces low level, random noise, see the middle and right hand graphs in figure 3.1.

So modulating a signal with an m-Sequence and then demodulating with the same m-sequence, will produce a small, negative signal until the two codes are in phase to within plus or minus a chip, at which point the amplitude of the signal will experience a jump in amplitude proportional to the length of the sequence. Figure 3.1 shows this correlation. Codes c1 to c3 are three different m-Sequences each of of 1023 digits (with values of 1 or -1). The difference between these three m-Sequences is simply the order in which the 1s and -1s are arranged. Each sums to +1. The first (left) of the graphs shows c1 (the first m-Sequence) auto-correlated with itself. The value is always -1 (see the expanded region around the bottom of the peak in figure 3.1), except when the m-Sequences are in phase when it jrises to 1023, the length of the sequence Ref [18]. The value of the m-Sequence c1 multiplied by c2 or c3 produces random noise (middle and right hand graphs) regardless the relative phase of the m-Sequences.

So applying code c1 to a signal, reflecting it from something and then demodulating with the same code when it returns will, in general, provide a signal with its current amplitude multiplied by -1. As the demodulation code is slid along the reflected, modulated signal, the amplitude remain at the same level (-1) until the phase of the modulating code and demodulating code align, at which point the amplitude will jump by a factor of 1023. Knowing the degree to which the codes have been shifted (which isolates the underlying heterodyne phase), the modulation frequency and the speed of light allows the time of flight to be calculated.

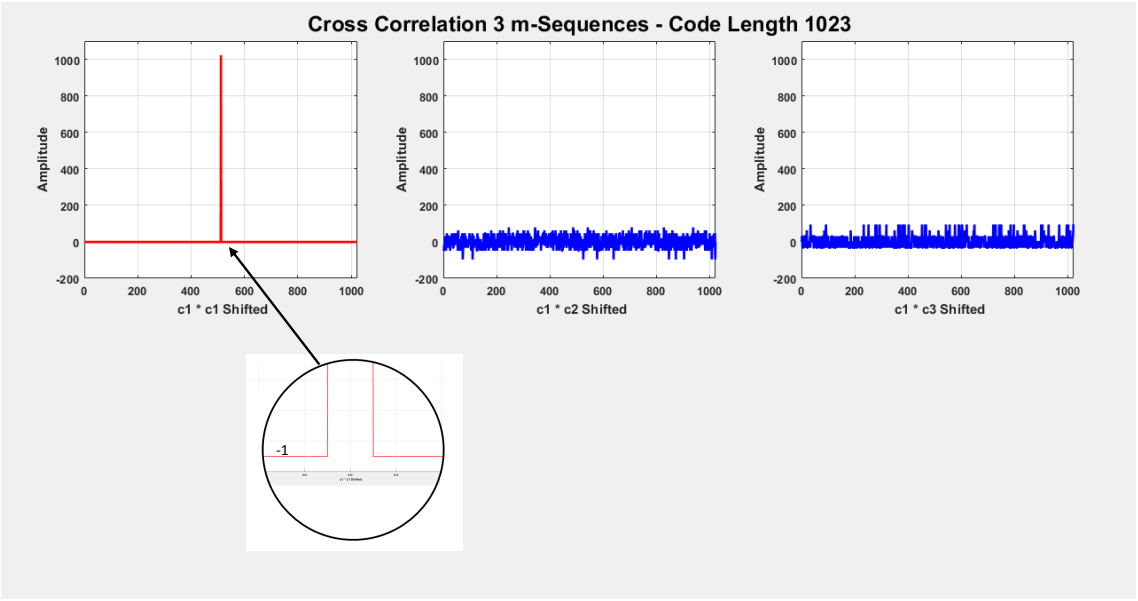

Figure 3.1: These graphs show the effect of multiplying an m-Sequence by itself and with other m-Sequences of the same length in various relative positions. 3 m-Sequences each of 1023 characters are explored.The graphs on top left shows an m-Sequence (c1) auto-correlated with itself. The result is -1 (see insert) in all cases except when the codes align, in which case it is +1023. Cross correlation with any of the other 1023 digit m-Sequences produces low level noise, graphs 2 and 3.

3.2.2 Hadamard Codes

This thesis is concerned with wavefront sensing which presents a different problem. Rather than applying one long code sequence across time, spatial tagging requires the simultaneous application of a number of code sequences to different parts of a wavefront at the same time. These codes remain in phase as the wave front evolves and, if they are orthogonal, provide a way of separating the attributes associated with one region of the wavefront from all the other regions in post processing. Provided there are enough codes in the set to tag each spatial region of interest, the codes can remain quite short, reducing integration times and increasing the rate at which measurements can be taken. Hadamard codes are best suited to this task. (m-Sequences could also be used for this task because their cross correlation produces low level noise but demodulation would be unable to remove this noise making them less attractive than Hadamard codes.)

Hadamard codes (also known as Walsh codes Ref [48]) are a set of codes with the following characteristics.

1. Each code in a sequence can have the value of 1 or -1.
2. Each code set forms a square matrix which is identical to its transpose.
3. The first row (or column) of any set is always all 1s so, in general, cannot be used to tag information usefully.
4. Any row (or column) multiplied by itself and summed gives a value equal to the length of the row.
5. Any row (or column) multiplied by another row (or column) and summed equals 0.

6. Hadamard matrices can be generated by a number of proven algorithms, for example the Matlab command *code* = *hadamard*(4); produces the matrix shown in figure 3.2, which was used in the first experiment.

	1	2	3	4
1	1	1	1	1
2	1	-1	1	-1
3	1	1	-1	-1
4	1	-1	-1	1

Figure 3.2: A Matrix of 4 Hadamard Codes. Rows 2 and 3 were used in the first experiment.

It is the orthogonality of Hadamard codes that makes them so ideally suited to wavefront analysis Ref [49]. (See also Ref [50] and Ref [51]). Physicists have long been familiar with orthogonal bases to deconstruct complex signals. Fourier analysis, for example, decomposes complex waveforms into their component sine functions and any continuous function can be deconstructed this way because the individual sine waves are orthogonal, that is, it is impossible to make one sine wave from any combination of other sine waves.

There are many other frequently used orthogonal functions, including Bessel functions, Hermite Gaussian polynomials and Hermite Lagrange polynomials which are used to analyse complex optical modes.

The digital realm has also developed sets of orthogonal series which can be used to distinguish individual signals delivered through a single channel. This happens every day with mobile phones using Code Division Multiple Access (CDMA) technology. The antenna on a mobile phone sees all the traffic in its area. It is able distinguish signals meant for its particular owner because all the signals are encoded with orthogonal codes and this allows it to reject all but its own signals Ref [52].

Orthogonal codes, then, have the property that a code from the set multiplied by itself sums to the number of elements in the sequence, while codes multiplied a different code from the same set sums to zero. This can be illustrated by using rows 2 and 3 from the 4 by 4 matrix of Hadamard codes, figure 3.2, [+1, -1, +1 ,-1] and [+1, +1, -1, -1].

Multiplying the elements of the first code (row 2) by itself and summing them yields 4 as shown below (here the operator .* is used to indicated element wise multiplication):

$$\sum([+1,-1,+1,-1].*[+1,-1,+1,-1]) = \sum[+1,+1,+1,+1] = 4$$

Similarly the second code (row 3) multiplied by itself and summed across the result equals 4.

$$\sum([+1,+1,-1,-1].*[+1,+1,-1,-1]) = \sum[+1,+1,+1,+1] = 4$$

Cross multiplying codes by each other yields zero:

$$\sum([+1,-1,+1,-1].*[+1,+1,-1,-1]) = \sum[+1,-1,-1,+1] = 0$$

Hadamard Codes are a class of codes with this property. Figure 3.3 shows the Hadamard set for 20 codes (as always, the first line is all ones). Taking any individual code (say row

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1
3	-1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1
4	-1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1
5	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1
6	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1
7	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1
8	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1
9	-1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1	1
10	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1	-1
11	-1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1	1
12	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1	-1
13	-1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1	1
14	-1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1	1
15	-1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1	1
16	-1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1	1
17	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1	-1
18	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1	-1
19	-1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1	1
20	-1	1	-1	-1	1	1	1	1	-1	1	-1	1	-1	-1	-1	-1	1	1	-1	1

Figure 3.3: A Matrix of 20 Hadamard Codes

3) and multiplying it by itself and summing gives a total of 20, which is the number of codes in the set, while taking the same code and multiplying it by any other from the set sums to zero. As an example, consider rows 3 and 7.

$$\begin{aligned}
Row3. * Row3 &= [-1-1+1+1-1-1+1+1+1+1-1+1-1+1-1-1-1+1+1] \\
&.*[-1-1+1+1-1-1+1+1+1+1-1+1-1+1-1-1-1+1+1] \\
&= [+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1] \\
&= 20 \\
Row3. * Row7 &= [-1-1+1+1-1-1+1+1+1+1-1+1-1+1-1-1-1+1+1] \\
&.*[-1-1+1+1-1-1+1+1+1+1-1+1-1+1-1-1-1+1+1] \\
&= [+1+1-1+1-1-1+1+1+1-1-1-1+1-1+1+1-1+1-1-1-1] \\
&= 0
\end{aligned}$$

More generally, the codes are orthogonal if:

$$\sum_{j=1}^{n} c_{i,j} \cdot c_{k,j} = \begin{cases} n \text{ if } i = k \\ 0 \text{ if } i \neq k \end{cases} \tag{3.1}$$

Hadamard codes are used as part of the Code-Division Multiple Access (CDMA) Ref [53] to distinguish mobile phones and are the basis used in this experiment to tag different areas of the wavefront.

3.3 A Visual Illustration of Digitally Enhanced Heterodyne Interferometry

Before presenting a mathematical analysis of DEHI, it is useful to see a visual representation of simple modulation and demodulation of a signal with Hadamard codes. This is shown in figure 3.4. The blue sine wave on the top row of the graph is modulated by the code ([+1 +1 -1 -1]) shown as the black trace on the second row. This results in the modulated signal shown in green in the centre row (three times). Row 4 shows three different demodulation codes; the original ([+1 +1 -1 -1]), a different code from the same Hadamard set ([+1 -1 +1 -1]) and the original, this time out of phase ([+1 -1 -1 +1]). As can be seen from the bottom row. Only demodulating with the original code in phase with itself recovers the signal. The other resulting signals are jumbled. The challenge for

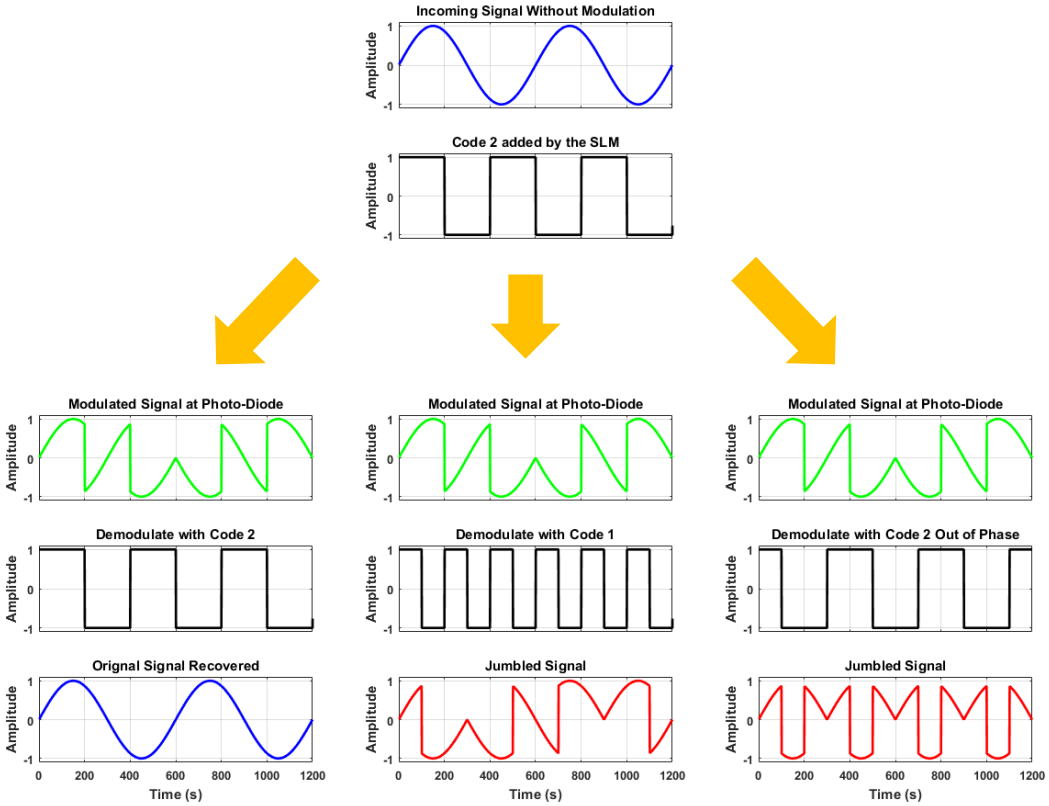

Figure 3.4: An example of modulation and demodulation with Hadamard codes. A sine wave is first modulated by the black Hadamard code resulting in the green modulated signal. Only demodulation with the same code in phase, the left hand column, recovers the original signal. If the same code is used out of phase (right hand column) or a different code is used (centre column) the signal remains unreadable.

DEHI is to reproduce this process, demodulating each channel to recover a specific part of the signal and then remove any jumbled clutter which obscures the desired information.

3.4 The Mathematics of Digitally Enhanced Heterodyne Interferometry

Chapter 2 outlined the standard IQ demodulation used with segmented diodes to make simple maps of a wavefront. Digitally Enhanced Heterodyne Interferometry extends this technique by tagging different regions of the wavefront with orthogonal codes using a Spatial Light Modulator (SLM). This is explained in detail in Chapter 5 but, for now, let us assume that a series of Hadamard codes have been applied successfully and synchronously to different spatial regions. As a starting point, consider rows 2 and 3 of the simplest codes, figure 3.2, $[= 1-1 = 1-1]$ and $[+1+1-1-1]$ and assume these have been applied to the left and right hand sides of the wavefront. For this case, 'synchronous application' means that the codes are applied in phase, at the same frequency and in the same order to both sides of the SLM so:

1. When the left hand side has the code +1, the right hand side has the code +1, then
2. when the left hand side has the code -1, the right hand side has the code +1, then
3. when the left hand side has the code +1, the right hand side has the code -1, then
4. when the left hand side has the code -1 , the right hand side has the code -1, then
5. the sequence repeats

Strangely, as we will see in Chapter 5, the frequency of application does not need to be constant provided it is controlled and changes to each side remain synchronised.

3.4.1 The Effect of Code Modulation

Equation 2.14 in Chapter 2 defined the signal beam as a combination of the angular frequency evolving over time added to a phase offset ϕ_s which reflected its phase difference from a local oscillator.

$$\tilde{E_s} = E_s.e^{-i\omega_s t+\phi_s} \tag{3.2}$$

If a code, $c_{i,j}$ is added to the signal using, for example, a spatial light modulator, the signal beam is changed to:

$$\tilde{E_s} = E_s.e^{-i(\omega_s t+\phi_s+Mc_{(i,j)})} \tag{3.3}$$

where M is the modulation depth of the code and $c_{(i,j)}$ is the i^{th} code in the j^th row of an orthogonal sequence Ref [54], namely one cell from a Hadamard matrix.

Again, the first step in deriving the total incident power on the photodiode is to sum the electric fields (once digitised by the ADC), namely:

$$\tilde{E_d} = \sum_{i=1}^{n} \tilde{E_i} + \tilde{E_l} \tag{3.4}$$

where $\tilde{E_i}$ is the electric field for the region of the wavefront encoded with the i^{th} row of the Hadamard matrix. The total power received is then:

$$\begin{aligned} P_d &\propto \tilde{E_d}.\tilde{E_d}^* \\ &\propto \left(\sum_{i=1}^{n} \tilde{E_i} + \tilde{E_l}\right).\left(\sum_{i=1}^{n} \tilde{E_i} + \tilde{E_l}\right)^* \\ &\propto \mid \tilde{E_l} \mid^2 + \tilde{E_l}^* \sum_{i=1}^{n} \tilde{E_i} + \tilde{E_l} \sum_{i=1}^{n} \tilde{E_i}^* + \sum_{i=1,j}^{n} \tilde{E_i}^* \tilde{E_j} \\ &\propto \mid \tilde{E_l} \mid^2 + \sum_{i=1,j}^{n} \tilde{E_i}^* \tilde{E_j} + \sum_{i=1}^{n} \left(\tilde{E_l}^* \tilde{E_i} + \tilde{E_l} \tilde{E_i}^*\right) \end{aligned} \tag{3.5}$$

Equation 3.5 contains three terms:

1. The first, $\mid \tilde{E_l} \mid^2$, which is the DC power of the carrier,
2. the second, $\sum_{i=1,j}^{n} \tilde{E_i}^* \tilde{E_j}$, which represents, when $i = j$, the DC power of the modulated electric field and when $i \neq j$ a series of orthogonal terms that average to 0, and
3. the third, $\sum_{i=1}^{n} \left(\tilde{E_l}^* \tilde{E_i} + \tilde{E_l} \tilde{E_i}^*\right)$, which is the signal of interest.

As the photodiode is AC coupled, the DC terms are removed so that the alternating power on the photodiode is proportional to the electric field multiplied by its complex conjugate and becomes:

$$\begin{aligned} P_{AC}(t) &\propto \sum_{i=1}^{n} \left(\tilde{E_l}^* \tilde{E_i} + \tilde{E_l} \tilde{E_i}^*\right) \\ &\propto 2 \sum_{i=1}^{n} Re\left(\tilde{E_l}^* \tilde{E_i}\right) \\ &\propto 2\tilde{E_l} \sum_{i=1}^{n} E_i \cos(\omega t - \phi_i - Mc_i) \end{aligned} \tag{3.6}$$

as explained in Chapter 2 , equation 2.16

3.4.2 IQ Demodulation

Splitting the signal into two channels and performing IQ demodulation gives:

$$\begin{aligned} P_{AC}(t)\cos(\omega t) &\propto 2\tilde{E_l} \sum_{i=1}^{n} E_i[\cos(-\phi_i - Mc_i) + \cos(2\omega t - \phi_i - Mc_i] \\ P_{AC}(t)\sin(\omega t) &\propto -2\tilde{E_l} \sum_{i=1}^{n} E_i[\sin(-\phi_i - Mc_i) + \sin(2\omega t - \phi_i - Mc_i] \end{aligned} \tag{3.7}$$

A low pass filter is used to eliminate second harmonics, specifically the terms containing 2ω, leaving:

$$\begin{aligned} P_{AC}(t)\cos(\omega t) &\propto 2\tilde{E}_l \sum_{i=1}^{n} E_i[\cos(-\phi_i - Mc_i) \\ P_{AC}(t)\sin(\omega t) &\propto -2\tilde{E}_l \sum_{i=1}^{n} E_i[\sin(-\phi_i - Mc_i) \end{aligned} \tag{3.8}$$

Now, applying standard trigonometrical identities, namely:

- $\cos(a+b) = \cos(a)\cos(b) - \sin(a)\sin(b)$
- $\sin(a+b) = \sin(a)\cos(b) + \cos(a)\sin(b)$
- $\sin(-x) = -\sin(x)$
- $\cos(-x) = \cos(x)$

and remembering that c_i is always 1 or -1, leads to:

$$\begin{aligned} P_{AC}(t)\cos(\omega t) &\propto \tilde{E}_l \sum_{i=1}^{n} E_i[\cos(\phi_i)\cos(M) - c_i\sin(\phi_i)\sin(M)) \\ P_{AC}(t)\sin(\omega t) &\propto -\tilde{E}_l \sum_{i=1}^{n} E_i[c_i\cos(\phi_i)\sin(M) + \sin(\phi_i)\cos(M)) \end{aligned}$$

3.4.3 Code Demodulation

These signals are now demodulated further by multiplying by one of the original codes, c_j:

$$\begin{aligned} c_j \cdot P_{AC}(t)\cos(\omega t) &\propto \tilde{E}_l \sum_{i=1}^{n} E_i[c_j\cos(\phi_i)\cos(M) - c_jc_i\sin(\phi_i)\sin(M)) \\ c_j \cdot P_{AC}(t)\sin(\omega t) &\propto \tilde{E}_l \sum_{i=1}^{n} E_i[c_jc_i\cos(\phi_i)\sin(M) + c_j\sin(\phi_i)\cos(M)) \end{aligned}$$

But the codes satisfy the following conditions:

- Each code is normalised, that is $\langle c_i^2 \rangle = 1$;
- Each code averages to zero over a complete cycle, that is $\langle c_i \rangle = 0$; and
- Each code is orthogonal to every other code when averaged over a complete cycle, that is $\langle c_i c_j \rangle = \delta_{ij}$

So integrating (or averaging) over an integer number of repetitions of the code we get, Ref [54]

$$\langle c_j \cdot P_{AC}(t)\cos(\omega t) \rangle \propto \tilde{E}_l E_j \sin(\phi_i)\sin(M) \tag{3.9}$$

$$\langle c_j \cdot P_{AC}(t)\sin(\omega t) \rangle \propto \tilde{E}_l E_j \cos(\phi_i)\sin(M) \tag{3.10}$$

and dividing the two expressions gives:

$$\frac{\langle c_j \cdot P_{AC}(t)\cos(\omega t)\rangle}{\langle c_j \cdot P_{AC}(t)\sin(\omega t)\rangle} = tan(\phi_j) \tag{3.11}$$

The phase of the j^{th} region of the wavefront (i.e. that encoded with the j^{th} row of the Hadamard matrix) can then be recovered from the arctangent:

$$\phi_j = tan^{-1}\frac{\langle c_j \cdot P_{AC}(t)\cos(\omega t)\rangle}{\langle c_j \cdot P_{AC}(t)\sin(\omega t)\rangle} \tag{3.12}$$

Notice that recovering the phase ϕ_j is not dependent upon the modulation depth M or the amount of light in the j^{th} region of the wavefront as both these quantities disappear during the final division. (There are some signal to noise ratio effects however.)

3.5 The Advantages DEHI Enhanced Wavefront Sensing

It is true that the signal used for Digitally Enhanced Heterodyne Interferometry is instantaneously noisier than the unmodulated signal and that the codes only disappear when the signal is integrated over the code length, however some form of integration is also required for traditional wavefront sensing so, properly managed, this is less of a disadvantage than it first appears. The technique does offer a number of advantages.

- Like wavefront measurements made with a segmented diode but unlike measurements made with a Hartmann sensor, the technique is a coherent measurement, that is it is made with the signal beam.
- Signals are returned through a single channel so all the noise sources introduced by the ADC and capacitance of the connections are common and divide away. This can be very important for satellite based interferometers as they move in and out of the Earth's shadow, for example.
- The signal is collected in a single diode and is less dependent upon precise alignment which may make the optical set up simpler providing the code modulation device is straightforward.
- The detail of wavefront map is limited only by coarseness of code modulation device and availability of back end processing power and more detailed maps can be created than those created using segmented diodes. (If course, it is possible to create segmented diodes with more than 4 segments but these are not readily available and unlikely to be able to compete with the definition that an SLM can provide.)
- The technique is not dependent upon the modulation depth (or its consistency) allowing some latitude in the specifications of the code modulator. (The modulation depth is a factor which divides out during the final calculation of phase so any short term variation in the accuracy of SLM displacements are not important if they are slow compared to a typical measurement cycle.)
- The technique does not require a highly accurate frequency of modulation providing the system can determine when codes have changed and that codes change across each wavefront region together.
- The technique can accommodate different light intensities across the wavefront and over time as this quantity is also removed by the final division step. (This is, of course, true of conventional wavefront sensing as well.)

- The code sequence need only be as long as the number of regions being mapped which means that changes in a wavefront can be followed and adjusted in near real time, as illustrated in the second experiment presented in this thesis. (Temporal demodulation with m-Sequences often requires very long codes for an unambiguous result because the speed of light is so fast.)
- Finally, from equations 3.9 and 3.10, the process also allows measurement of the relative amplitude of different parts of the wavefront, both over time and with respects to other spatial areas. This property is harnessed in the second experiment which aims to maximise the transmission through a cavity.

It is also worth noting that, beside proving a wavefront sensing capability, the SLM activators also make it possible to real time corrections to the wavefront on the basis of the information collected. (This is illustrated in the first experiment, presented in Chapter 5.)

Cavity Dynamics and Control

The second experiment, presented in chapter 6, shows how the DEHI techniques developed in the first experiment, chapter 5, can be used to maximise the transmission of light through an optical cavity. It relies on the fact that the technique can measure both the phase and amplitude both across a wavefront and over time. As an introduction this chapter examines the fundamentals of cavity dynamics and control, with a particular emphasis on how these are impacted by spatial variations in the wavefront. It starts with a brief review of cavities and their uses. It then goes on to examine cavity modes and how the resonance of the fundamental mode can be sustained using an error signal created using Pound Drever Hall locking. Finally it considers higher order cavity modes and how these may be identified, rejected or, indeed, selected.

Those familiar with the dynamics of cavities can skip it on move onto the experimental work presented in Chapters 5 and 6.

4.1 Optical Cavities

The first optical cavity most people encounter is the mirrored elevator in an expensive hotel. People standing in the lift can look to their right or left and see multiple reflections of themselves heading off into infinity. These reflections are misaligned and get fainter and fainter because the mirrors do not reflect all of the light striking their surfaces. Nonetheless, ten or so reflections are often visible. This is a simple but particularly poorly designed optical cavity.

A linear optical cavity or Fabry-Pérot cavity (or etalon) is formed by two reflecting surfaces that face each other. These surfaces may be flat or curved so that, once the light has entered the cavity, usually through the back of one of the mirrors, it resonates (reflects back and forth) between them many times before emerging from the other side of the cavity as a high quality beam. The beam is of high quality (namely coherent and centred around a single frequency) because light of other phases and frequencies does not generally resonate, that is reflect back and forth, within the cavity and gets rejected by (reflected off) the incident mirror rather than circulating within the cavity and being transmitted through the non-incident mirror, figure 4.3.

4.1.1 Cavity Uses

Cavities have a multitude of uses. Linear cavities can act as mode cleaners transmitting only those frequencies which match the resonances of the cavity. They can also store power and act as an optical delay. In the LIGO interferometers the arm cavities act to increase

the effective interaction length of the interferometers arms from 8 kms (a single round trip), to many hundreds of times that length Ref [55], Ref [56], Ref [32] and Ref [13].

Ring cavities, typically behave as travelling wave cavities in which the light passes any given point in the cavity in only one direction and, as such, is more immune to back scatter than a liner cavity. These cavities have spatially separated input and output ports and have a range of uses in sensing and, when light is sent in both directions simultaneously while the cavity is rotated (Sagnac interferometry), as gyroscopes on aircraft and other position sensitive devices.

4.1.2 Cavity Dynamics

Consider a cavity in which two mirrors are perfectly aligned along the optical axis to face each other, figure 4.1. Further imagine a laser emitting a single frequency at a constant amplitude, E_i, incident along the optical axis from the left hand side. The amount of light transmitted from the cavity is entirely dependent upon E_i and the reflection coefficients and transmission coefficients of the two mirrors comprising the cavity. The reflection coefficient of the left hand mirror is r_1 and is defined as the fraction of the electric field which is reflected each time a beam of light is incident upon the mirror. Similarly t_1 is the transmission coefficient of the left hand mirror and is defined as the fraction of the electric field which is transmitted through it each time a beam is incident upon it. R_1 is the percentage of power that is transmitted when a beam is incident upon the left hand mirror and is r_1^2. Similarly T_1 is the percentage transmitted is and is t_1^2. If the cavity is lossless, all the electric field is either transmitted or reflected and $R_1 + T_1 = r_1^2 + t_1^2 = 1$, see page 15 of Ref [57].

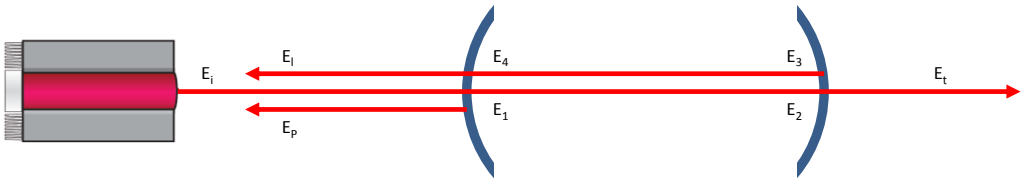

Figure 4.1: This diagram shows incident power is transmitted through a cavity after accounting for internal reflections. In reality all the light rays would lie along the optical axis but they are shown side by side here for clarity.

The electric field emanating from the laser is given by Ref [58]:

$$\begin{aligned} \tilde{E}(x,t) &= E_i \cdot e^{i\cdot(kx-\omega t)} \\ &= E_i \cdot e^{ikx} \cdot e^{-i\omega t} \end{aligned} \tag{4.1}$$

The $e^{-i\omega t}$ term represents the angular frequency of the light and is always present but can be ignored in subsequent analysis. The e^{ikx} term represents the evolving phase of the light as it travels a distance x back and forth along the optical axis of the cavity. This is dependent only upon this distance x and the wavelength, λ, and is captured in k, the wave number, or $2\pi/\lambda$. Thus every time x increases by λ the wave passes through a full cycle, the phase has advanced by 2π and is back to where it was. Consequently if the

cavity is an exact number of half wavelengths long, the return trip within it is an exact number of whole wavelengths and the circulating beam's phases align perfectly each cycle and it reinforces itself through constructive interference.

Using the convention that the electric field of any transmitted beam must be multiplied by i to preserve conservation of energy (see Ref [59]) and referring to figure 4.1 we have at steady state:

$$\begin{aligned} E_i &= E_i i t_1 + r_1 E_4 \\ &= E_i i t_1 + E_3 r_1 e^{ikL} \\ &= E_i i t_1 + E_2 r_1 r_2 e^{ikL} \\ &= E_i i t_1 + E_1 r_1 r_2 e^{ik2L} \end{aligned} \tag{4.2}$$

Where:

- r_1 is the reflection coefficient of the left hand mirror;
- t_1 is the transmission coefficient of the left hand mirror;
- r_2 is the reflection coefficient of the right hand mirror;
- E_i is the amplitude of the electric field of the incident beam;
- E_1 is the amplitude of the electric field just after passing through the left hand mirror;
- E_2 is the amplitude of the electric field just before reflection from the right hand mirror;
- E_3 is the amplitude of the electric field just after reflection from the right hand mirror;
- E_4 is the amplitude of the electric field just before transmission through and reflection from the left hand mirror;
- L is the length of the cavity; and
- k is the wave number, the number of waves per unit of length.

Solving for E_1 we get:

$$E_1 = \frac{E_i i t}{1 - r_1 r_2 e^{2ikL}} \tag{4.3}$$

The amplitude of the transmitted electric field, E_t, is given by:

$$\begin{aligned} E_t &= E_2 i t_2 \\ &= \frac{E_i i t_1 e^{ikL} i t_2}{1 - r_1 r_2 e^{2ikL}} \\ &= -\frac{E_i t_1 t_2 e^{ikL}}{1 - r_1 r_2 e^{2ikL}} \end{aligned} \tag{4.4}$$

where t_2 is the transmission coefficient of the right hand mirror. This expression must be squared to provide the magnitude of the transmitted power.

The amplitude of the reflected electric field comprises the combination of that light reflected directly off the cavity without entering, the prompt reflection or E_p, and the light that leaks out of the cavity after it strikes the left hand mirror from the right hand side, the leakage field or E_l. Similarly we have:

$$\begin{aligned} E_r &= E_p + E_l \\ &= E_i r_1 + E_4 i t_1 \\ &= E_i \left(r_1 + \frac{(it_1 \cdot it_1 r_2 e^{2ikL})}{(1 - r_1 r_2 e^{2ikL})} \right) \\ &= E_i \left(r_1 + \frac{(-t_1^2 \cdot r_2 e^{2ikL})}{(1 - r_1 r_2 e^{2ikL})} \right) \\ &= E_i \left(\frac{r_1 \cdot (1 - r_1 r_2 e^{2ikL}) - (t_1^2 \cdot r_2 e^{2ikL})}{(1 - r_1 r_2 e^{2ikL})} \right) \\ &= E_i \left(\frac{r_1 - (r_1^2 + t_1^2) \cdot r_2 e^{2ikL})}{(1 - r_1 r_2 e^{2ikL})} \right) \end{aligned} \tag{4.5}$$

This is an interesting result. The prompt reflection, E_p, has the opposite phase on resonance to the leakage field, E_l, because the the leakage field has been transmitted as it returns through the left hand mirror and so has had its phase inverted. It is therefore possible that these fields can be of equal magnitude but opposite phase and completely cancel each other out. In this case all the light enters the cavity and is either lost or emerges as the transmitted beam.

In the idealised case where the system is lossless (that is everything is perfectly aligned and both $r_1^2 + t_1^2 = 1$ and $r_2^2 + t_2^2 = 1$ all the light is either transmitted or reflected each time it strikes a mirror and none is lost. When $r_1 = r_2$ and the cavity it lossless, the system is said to be impedance matched or critically coupled. Further if the length of the cavity L is tuned to a precise number of half wavelengths, N, then each round trip within the cavity is a complete number of wavelengths and:

$$e^{2ikL} = 1$$

In this case equation 4.5 becomes:

$$E_r = E_i \left(\frac{r - 1 \cdot 1 \cdot r \cdot 1}{(1 - r^2 \cdot 1)} \right) = E_i \left(\frac{0}{(1 - r^2)} \right) = 0 \tag{4.6}$$

If r_2 is greater than r_1, the leakage field can be up to twice as large as the prompt reflection and the cavity is said to be over coupled. If r_2 is less than r_1, the leakage field is less than the prompt reflection and the cavity is said to be under-coupled.

So, in summary, in a lossless, impedance matched cavity all the light circulates until it is transmitted and none is lost or reflected back towards the laser. Of course, no mirror is perfectly reflective so whenever a beam strikes its surface a portion of the light is reflected, a portion is transmitted and a portion is lost through scattering or absorption. Light striking one of the mirrors of a cavity is no different. Some is reflected away from the cavity, some is lost but some enters the cavity and propagates towards the second mirror. Here again most of the light is reflected but some is transmitted (that is passes through the second mirror and out of the cavity) and some is lost.

If the cavity is well constructed (see section 4.3) and it's length is an integral number of half wavelengths of the incident light, the light reflects back and forth between the mirrors and a standing wave is created. In figure 4.3, for example, the light that enters the cavity, is reflected from the right hand mirror that then again off the left hand mirror, is exactly in phase with new light entering the cavity, which it re-enforces. Meanwhile light of a different wavelength, that is not exactly divisible into twice the length of the cavity, does not interfere constructively as it circulates and begins to interfere with itself destructively. If the reflection coefficients of the mirrors are high and the light circulates many times before emerging as transmitted light, these frequencies disappear (or more accurately become part of the leaked reflection heading back towards the laser) and only those frequencies for which $N\lambda = 2L$ are transmitted. This is shown in figure 4.2 and Ref [60].

Real lasers are not monochromatic and emit a range of frequencies which are distributed as a Gaussian curve around the nominal frequency, the blue curve in figure 4.2 and see Ref [60]. Not all of these are frequencies are resonant in a Fabry-Pérot cavity, however, which transmits only those frequencies which can resonate within it, the vertical red lines. The cavity effectively becomes a filter which passes only a small selection of frequencies from the original broader band of frequencies incident upon it.

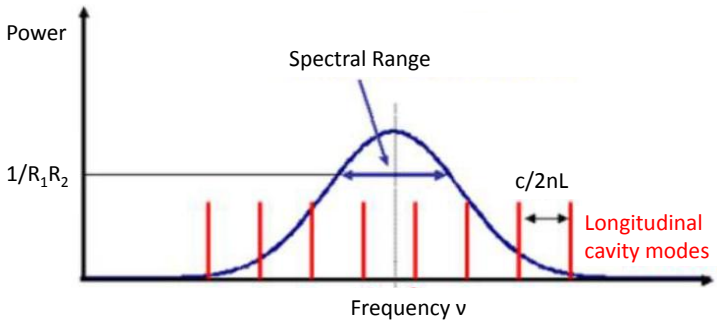

Figure 4.2: Fields emerging from a Fabry-Pérot cavity, see Ref [60]

See also Ref [61] and Ref [62] for more on resonance with the LIGO cavities.

4.1.3 Free Spectral Range

The gap between these transmitted frequencies is called the Free Spectral Range (FSR) and can be calculated from:

$$FSR = c/2nL \tag{4.7}$$

where c is the speed of light, n is the group refractive index of the medium through which the light is travelling and L is the length of the cavity. For the fixed spacer, plano-convex cavity used in the second experiment, chapter 6, L was 0.3m, n was 1 and $c = 3.10^8 m/s$ so the FSR was:

$$FSR = \frac{3.10^8}{(1 \cdot 2 \cdot 0.3)} = 500 \cdot 10^6 = 500\text{MHz}$$

This cavity has the effect therefore of filtering out all those frequencies within the incident beam which are not separated by 500MHz, that is, do not match the mode of the cavity.

(The degree to which this happens is explored below.) In other words, it cleans the emerging beam and narrows its transmission bandwidth. It only transmits light which is coherent and of a very narrow bandwidth, Ref [63].

4.1.4 Quality Factor and Finesse

Two other useful measures for an optical cavity are its Quality Factor (Q) and finesse ($\mathcal{F}$). As we have seen, a laser beam can only resonate within a cavity if the cavity's length is a precise number of half wavelengths, figure 4.3. This occurs every time the length of the round trip increases by a whole wavelength and the gap between these occurrences is the free spectral range. The Q of a cavity is inversely related to the loss during a round trip and measures how long it continues to resonate and can be written as (see Ref [64] and Ref [65]):

$$Q = \frac{\text{Energy stored in the cavity}}{\text{Energy lost per cycle}} \tag{4.8}$$

For high Q cavities, this approximates to (see Ref [66]).

$$Q = \nu_0/\Delta\nu \tag{4.9}$$

High Q cavities store more energy as the light is reflected more times and the transmitted beam has a narrower line width, a desirable quality in high precision applications like LIGO.

Closely related to the quality factor is finesse, $\mathcal{F}$, which is FSR divided by the Line Width, specifically the Full Width Half Maximum (FWHM) line width or spectral range, $\Delta\nu$, in figure 4.2. The finesse is just defined as the free spectral range divided by the line width of the resonance:

$$\mathcal{F} = \frac{c/2nL}{\Delta\nu} \tag{4.10}$$

4.2 The Characteristics of the Cavity Used in the Experiment

The specific cavity used in the experiment reported in the next chapter was an aluminium fixer spacer, plano-concave cavity, 30cm long. The incident mirror was flat with a reflection coefficient, $r_1 = 99\%$. The rear mirror was concave with a radius of curvature of approximately 3m and a reflection coefficient, $r_2 = 99\%$. The finesse of the cavity was calculated to be 314. (This was confirmed by using the oscilloscope to estimate band width which resulted in a value of approximately 300.)

4.3 Maximising Cavity Transmission

Maximising transmission through a Fabry-Pérot cavity requires control of a number of factors including:

1. The structure of the cavity, specifically whether the mirrors are well aligned.

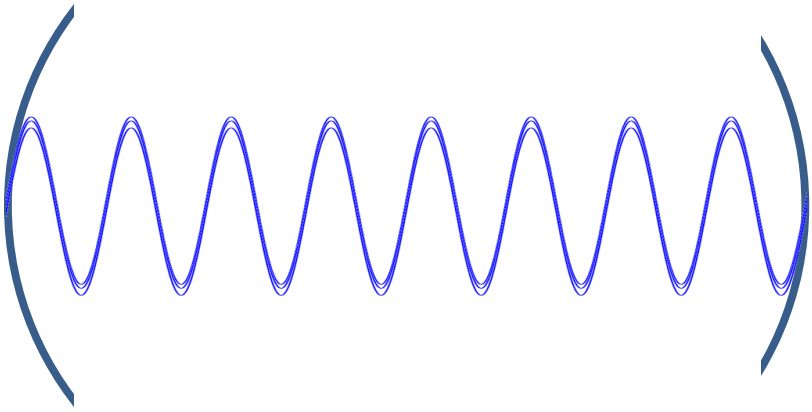

Figure 4.3: A Standing Wave Resonating in a Cavity. When the cavity length is an exact number of half wavelengths the light circulates and each reflection constructively reinforces new light entering the cavity. The energy builds to many times the incident energy with a small portion of the circulating power being transmitted each time. (Again the beams have been shown as slightly separated for clarity but in fact they lie on top of each other.) Beams of similar but slightly different wavelengths begin to destructively interfere with themselves as they circulate, preventing the build up of power of that wavelength in the cavity. This light is reflected back from the cavity rather than resonating within it.

2. The degree to which the cavity is impedance matched. (Cavities which are not impedance matched reflect more light leaving less for transmission.)
3. The match of the incident beam's wavelength to the cavity's length.
4. The alignment and offset of the beam with respect to the cavity's optical axis.
5. The position and size of the waist of the beam (which uniquely determines its divergence).
6. The quality of the beam itself, specifically it's wavefront.

Other factors which may have an effect are the line width and polarisation of the beam.

4.3.1 Mechanical Attributes

The first three of the attributes discussed above are mechanical characteristics of the cavity and cannot be controlled by manipulating the shape of the wavefront. For the experiments reported in the following chapter, the cavity was accepted as delivered (it was a good quality cavity) and the frequency of the laser was locked to the cavity length using Pound Drever Hall Locking (see section 4.4).

Cavity Structure

Clearly the physical characteristics of a cavity affect the fraction of incident power it can transmit. In particular, if the mirrors are of low quality or are misaligned, losses will be large, transmission will be minimal or, in extreme case, non-existent.

Degree of Impedance Matching

Equation 4.4 showed that the reflection coefficients of the two mirrors comprising the cavity place an upper limit of the amount of incident light it is able to transmit. If these are equal (and the cavity is lossless and its length is an integer number of half wavelengths) 100% can be transmitted. If not, this amount is reduced by an amount determined by equation 4.4. For the cavity used in the experiment the reflection coefficients close to equal at $r_1 = r_2 = 99\%$ so the maximum power of a mode matched wavelength that the cavity can transmit is close to 100%.

Locking Cavity Length to Laser Frequency

In the real world the frequency of the laser and the length of a cavity are continually changing, see page 254, Ref [67]. For example, the cavity used in the alignment optimisation experiment (Chapter 6) was 30 cm long and made of aluminium. Aluminium has a coefficient of expansion of $2.5.10^{-5}$ C^{-1} which meant that every change in temperature of 0.1°C resulted in a change of cavity length of 0.75 microns. For a round trip within the cavity, this is approximated 1.5λ for 1064 laser light. Similarly the frequency emitted by the laser is dependent upon the incident current, the ambient temperature and the behaviour of the laser's gain medium, Ref [67].

When transmitted power is maximised, reflected power is minimised. Figure 4.4 presents the results of a simulation which shows how the power reflected from a cavity depends upon how closely the cavity length is an exact number of half wavelengths. Here, because the reflection coefficients are not equal ($R_1 = 98.5\%$ and $R_2 = 95\%$), the cavity is not impedance matched and there will always be some reflected power. The reflected power is used to generate an error signal to lock the laser frequency to the cavity length and this is explained in section 4.4 below.

4.3.2 Higher Order Modes

Figure 4.5 shows the power transmitted by a cavity. It is plotted in the frequency domain as the frequency of the incident light is swept across multiple resonances. The power transmitted when the cavity resonates in the fundamental modes are shown in red and are separated by an FSR.

Other modes lie between the fundamental modes and are also separated by an FSR, (see section 4.6 and Ref [68]). In general, the odd modes are due to misalignment while the even modes are due to errors in waist size and position. Mode matching a cavity involves minimising these higher order modes and coupling as much of the power as possible into the fundamental or TEM_{00} mode.

4.3.3 Mode Matching

Three of the requirements for maximising transmission are concerned with matching the position and shape of the beam to the cavity and are the factors optimised in the following experiments. This optimisation is termed "mode matching" and involves:

1. Ensuring the beam is aligned to the optical axis of the cavity (tilt and offset).

2. Ensuring the beam's waist is the correct size and in the correct position.

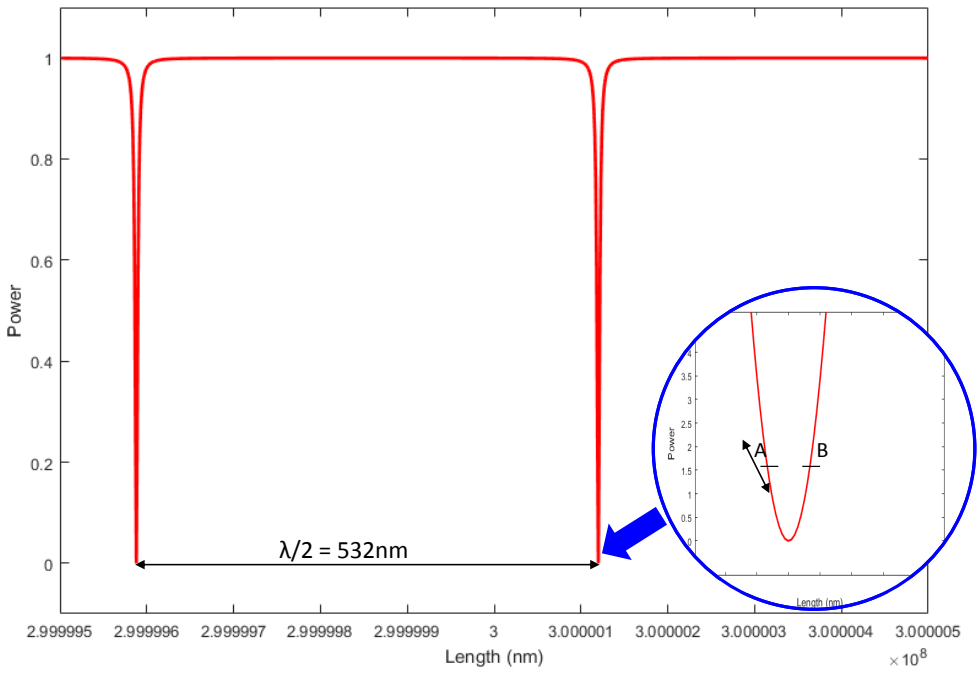

Figure 4.4: The power reflected from a cavity changes as the length changes, being minimal when the cavity length is an exact number of half wavelengths. For a lossless cavity, shown here, it falls to zero. The blue circle shows a magnified picture of the bottom of the trough.

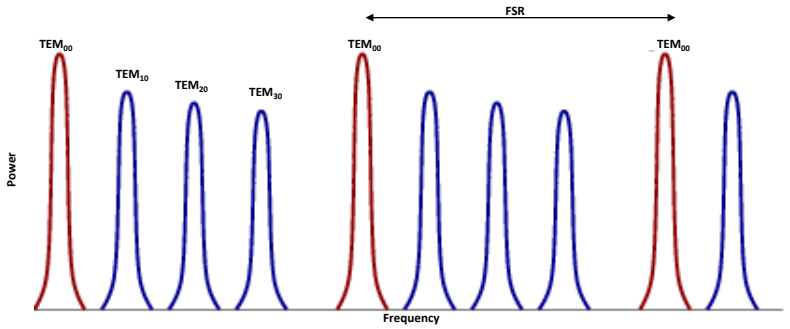

Figure 4.5: An illustration of the power transmitted by a poorly mode matched cavity in the frequency domain as the frequency of the incident light is swept across multiple resonances. The fundamental modes are shown in red and are separated by an FSR. Other modes lie between the fundamental modes.

3. Ensuring the beam is Gaussian in shape.

At least two actuators are required to manage these attributes.

Tilt and Offset

Managing tilt and offset is a common experience for anyone constructing an optical layout and typically involves adjusting two mirrors in the optical path, one far from the target (in this case a cavity) to manage offset and one close to the target to manage tilt, figure 4.6. Off course, as neither mirror is on top of the target or infinitely far from it, adjustments couple into each other but, in general, they are relatively distinct. The distance to the far mirror can be reduced by inserting a lens. Essentially it is a question of leverage. If a long stick is held at one end, (A), and pivoted the other end, (B), which is pointing at a target, B moves around considerably, producing a large offset from the original position. If the same stick is held near B and pivoted, the end B hardly moves at all and only the angle it presents to the target changes.

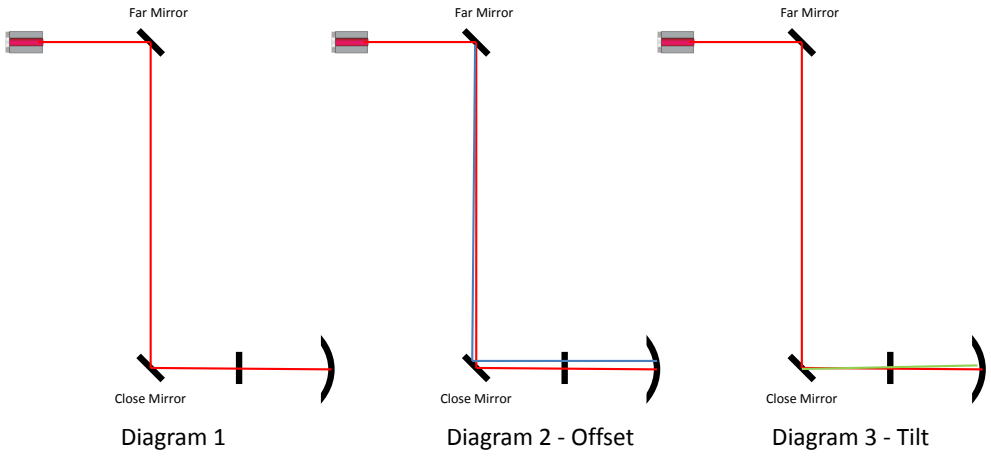

Figure 4.6: Tilt and Offset can be managed reasonably independently using two adjustable mirrors (Diagram 1), one far from the target and one close to it. When the far mirror is adjusted the beam moves as shown by the blue trace (Diagram 2), affecting the offset of the beam. The mirror close to the target is to close to create much offset and only tilts the wavefront as in the green trace (Diagram 3).

Waist Size and Position

The connection between waist size and position is more complex and is shown in figure 4.8. Finding a particular waist size at a particular position involves very careful placement of lenses and this is turn is impacted, in very practical terms, by the availability of the specific lenses needed to create a particular optical layout, and the extreme dependence of waist position on lens position. This is most easily understood by considering that small changes in lens position produce very large changes in waist position. For example, a lens will position the waist:

- at its focal length for a collimated beam which is typically of the order of a tens or hundreds of millimetres (or, if fact, for a Gaussian beam just in front of its focus, Ref [68]);
- at infinity for a beam expanding from the focus of a lens place a focal length away;
- somewhere between the focal length and infinity for a beam emerging from somewhere between these two points.

Figure 4.8 shows how waist size and position are affected by relatively small changes in the curvature of the two SLMs. The curves generates were developed using Alexei Ourjoumtsev's modelling tool once again using the following set up:

- A beam with a waist of 0.55mm at a point 1600mm before the first SLM was established and turned into a roughly collimated beam of diameter 3.5mm using three lenses, figure 4.7, - f1479.6 at -1300mm, f50 at -640mm and f300 at -288.7mm. (The f1470.6 lens is not, of course, available commercially and is included to model characteristics equivalent to those of the beam created in the experiment.) ;
- Two SLMs are inserted, initially flat (i.e. $f_{infinite}$) at the 0mm point and the 1200mm points;
- The beam is reduced down using a telescope to generate a waist of 0.55mm at 1602mm, 602mm from the second SLM.

This approximates the dimensions of the experiments presented in the following chapter. The model shows how sensitive the waist position, in particular, to the the placement of the optical elements. The curves were generated by varying the focal lengths of the SLMs around the zone of interest as follows:

- SLM1 at 0mm focal lengths used -1000, -2000, -4000, -5000, -5500, -6000, -6500, -7000, -7500, -8000, -9000, -10000, -11000, -12500, -15000, -17500, -20000, -30000, -40000, 40000, 30000, 25000, 20000, 17500, 15000, 12500, 10000, 8000, 6000, 4000, 2000, 1000.
- SLM2 at 1200mm focal lengths used -8500, -12500, -17000, 17000, 125000 and 8500.

giving a total of 32 data point for each of the six curves.

One of the goals of the method used in these experiments is to be able to make fine adjustments to a beam's divergence without having to move optical components with high precision and these curves illustrate just how sensitive these adjustments are. They also show why two SLMs are required. A single actuator is insufficient to optimise both waist size and waist position.

A Gaussian Beam

The shape of the wavefront is also important. Each distortion caused by the quality of the laser or imperfections in the optical chain, potentially translates into a loss of transmitted power. Generally, traditional optical components cannot repair flaws in the beam shape. The Hartmann sensor, for example, was able to find an almost invisible flaw in a LIGO test mass Ref [32] but could do nothing to rectify the problem. This method may be able to not only identify some of these flaws but compensate for them.

4.4 A Brief Summary of Pound Drever Hall Locking

In practice both the frequency of the laser and the length of the cavity vary constantly. Changes in laser frequency occur due to a range of factors including temperature variations, mechanical imperfections and the laser's gain dynamics which change the length of the laser gain medium, the driver current and voltage and even the width of the atomic transition which create the photons within the gain medium Ref [67]. Similarly temperature changes

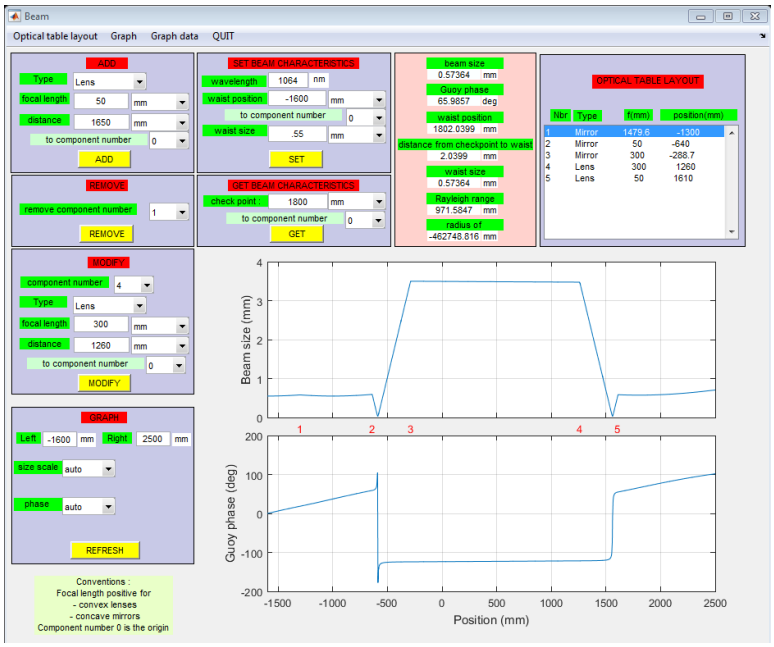

Figure 4.7: Modelling the relationship between waist size and position. A beam with of waist 0.55 mm at the -1600 mm point is established using Alexei Ourjoumtsev's modelling tool and expanded to become a collimated beam of approximated 3.5 mm in diameter at the 0 point. SLMs are place at the 0mm and 1200 mm point before the beam is reduced to a waist of 0.55 mm again for entry into the cavity. Iterating the model to change the curvatures of the two SLMs shows how the waist size and position changes as the wavefront is adjusted, figure 4.8.

and mechanical vibrations can affect the length of the cavity. Figure 4.4 shows how the reflected power varies as the these attributes drift in and out of alignment. Consequently, to maintain maximum transmitted power, the length of the cavity must be locked to the frequency of the laser so that the cavity length remains an integer number of half wavelengths. This requires a control signal that constantly drives these two attributes together. Either the cavity length must be continually adjusted to follow changes in laser frequency, which is accomplished using a piezo attached to one of the cavity mirrors to make continual, small adjustments to cavity length, or small changes are made to the frequency of the laser by changing the injected current or the temperature of the gain medium.

Transmitted power increases as the reflected power diminishes and reaches peaks at every FSR which correspond to points when the length of the cavity has has increased by a half wavelength, figure 4.9 and Ref [69] and Ref [70].

Figure 4.4 suggests that creating this control signal is relatively straight forward. All the control system needs to do is lock onto the reflected power and constantly drive that to a minimum. This simple schema suffers from two drawbacks however:

1. The reflected power is symmetrical about the minimum value so if the control signal measures the reflected power at point A in the figure, it is unable to distinguish this from point B so, even though it can recognise that it is not at a minimum, it cannot provide any information about which way to drive the correcting mechanism.

2. The measurement is a measure of absolute power so intensity variations, which are continual, couple into the measurement and create false indications.

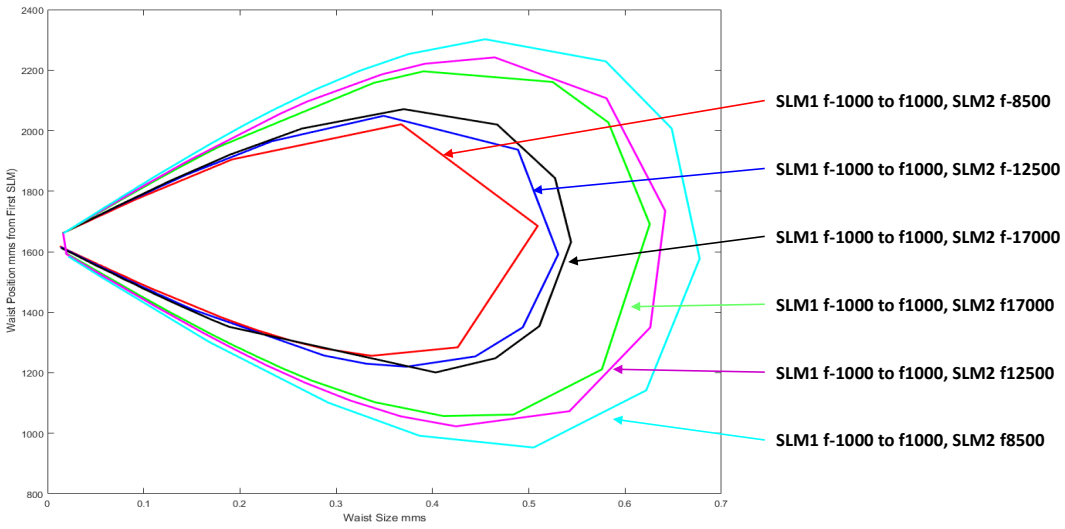

Figure 4.8: The relationship between waist size and position is more complex than simple orthogonality. The curves show contours for various combinations of curvature on the two SLMS. The perfect beam has a waist of 0.55mm at the centre of the plane mirror. This is one specific spot on the contours and uniquely defines the combination of curvatures required.

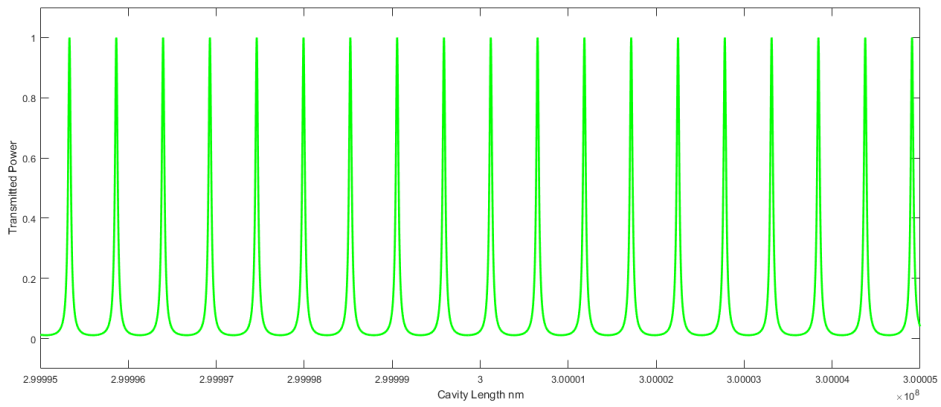

Figure 4.9: The frequency dependent power transmitted by a linear Fabry–Pérot cavity. In this example, the cavity is lossless and the reflection coefficients of both mirrors are 90%. The cavity is impedance matched and perfect mode matching is assumed so that no higher order modes are present.

What is required is a control signal which:

1. Is asymmetric, that is provides a clear indication of which way to make an adjustment, and

2. Is independent of received power (to the first order).

Pound Drever Hall Locking (PDH Locking) does just this. PDH Locking is the standard way of locking the frequency of the laser to the length of the cavity (see Ref [71], Ref [72] and Ref [73]). The technique produces an error signal which has two characteristics. It crosses zero when the frequency and cavity length match and is asymmetric, that is has a different sign before and after this crossing point making it possible to tell which side of that point it is.

The PDH signal is created using a combination of modulation, demodulation and filtering to produce this signal. The signal beam is passed through an Electro-Optic modulator which creates two side-bands before it is directed into the cavity. Now components of the reflected beam and the the side-bands reflect from the cavity and interfere with each other creating an electric field with nine components in total. This is mixed with a electronic local oscillator, which matches the frequency used in the EOM, and then filtered to remove both the up shifted DC terms and the high frequency terms. The resulting signal has the form shown in figure 4.10. The area of interest is around the resonant frequency where the signal crosses zero at the resonant frequency and has a large slope meaning it is a very sensitive indicator of the zero crossing point. Further, as the paragraphs below show, the signal is independent of the power being emitted at the time so is uncontaminated by small power fluctuations.

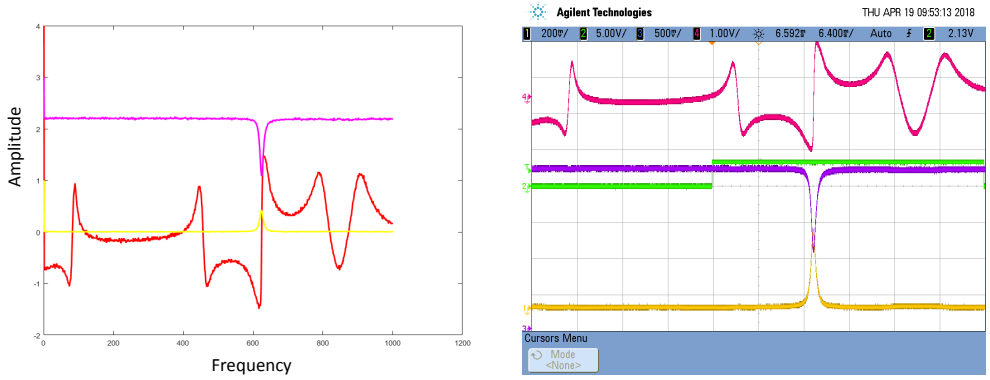

Figure 4.10: The PDH signal overlaid on the transmitted power. The figure shows the plotted values of the signal on the left and a screen shot from the experiment (displayed on an oscilloscope) on the right. In both cases, the red curve is the PDH signal, the yellow curve the transmitted power and the magenta curve the reflected power. (The green curve on the oscilloscope trace is a timing pulse not reproduced on the left.) The red curve, the PDH signal, crosses zero at resonance, has a steep slope and is independent of the reflected power making it an ideal control signal for locking the cavity to the laser.

Briefly the mathematics of PDH locking is as follows. The electric field of the signal beam, $\tilde{E_i}$, can be represented by $E_i e^{i\omega t}$ where E_i is its amplitude, ω is its angular frequency and t is time. E_i is the amplitude term and $e^{i\omega t}$ is the phase term. If this is phase-modulated by an EOM by $\beta sin(\omega_m t)$, two side-bands are added and the electric field becomes:

$$\tilde{E_i} = E_i e^{i(\omega t + \beta sin(\omega_m t))} \tag{4.11}$$

where ω_m is the modulation frequency. Using the Jacobi-Anger expansion allows this to

become, to a close approximation Ref [58]:

$$\begin{aligned}\tilde{E}_i &= E_i e^{i\omega t}[1 + i\beta sin(\omega_m t] \\ &= E_i e^{i\omega t}\left[1 + \frac{\beta}{2}e^{i\omega_m t} - \frac{\beta}{2}e^{-i\omega_m t}\right]\end{aligned} \tag{4.12}$$

It is this field that is a superposition of the original signal and the two side-bands. Now the reflected power is given by equation 4.5 which can be divided by $\tilde{E}_i$ to give a reflected transfer function $R_(\omega)$ of

$$R_{(\omega)} = \frac{E_r}{E_i} = \left(\frac{r_1 - (r_1^2 + t_1^2) \cdot r_2 e^{2ikL})}{(1 - r_1 r_2 e^{2ikL})}\right) \tag{4.13}$$

so the reflected beam. The reflected modulated beam becomes:

$$\tilde{E}_r = E_i\left[R_{(\omega)}e^{i\omega t} + R_{(\omega+\omega_m)}\frac{\beta}{2}e^{i(\omega+\omega_m)t} - R_{(\omega-\omega_m)}\frac{\beta}{2}e^{-i(\omega-\omega_m)t}\right] \tag{4.14}$$

Each of these terms interferes with itself and the two other terms to create an expression for the reflected power comprising 9 terms, some at DC, some at ω and some at 2ω. If P_i is the incident power and P_r is the reflected power, the power registered on the photodiode measuring the reflected light is $\tilde{E}_r * \tilde{E}_r^*$:

$$\begin{aligned}P_r = {} & P_i|R_{(\omega)}|^2 + P_i\frac{\beta^2}{4}\left\{|R_{(\omega+\omega_m)}|^2 + |R_{(\omega-\omega_m)}|^2\right\} \\ & + P_i\beta\left\{\mathrm{Re}[\chi_{(\omega)}]cos\omega_m t + \mathrm{Im}[\chi_{(\omega)}]sin\omega_m t\right\} + (\text{other terms in } 2\omega_m)\end{aligned} \tag{4.15}$$

where

$$\chi(\omega) = R_{(\omega)}R^*_{(\omega+\omega_m)} - R^*_{(\omega)}R_{(\omega-\omega_m)}. \tag{4.16}$$

It is this χ term which provides the asymmetric function that is useful as a control signal. To isolate it the voltage, V_r recorded by the photodiode (which is, of course, proportional to the reflected power) is mixed with a phase-delayed version of ω_m to produce a new voltage V_r' which is given by:

$$V_r' = V_r cos(\omega_m t + \varphi) \tag{4.17}$$

This demodulated signal is then passed through a low pass filter to remove the fast oscillating terms and leave an expression which depends only on χ:

$$V(\omega) = Re[\chi(\omega)]cos\varphi + Im[\chi(\omega)]sin\varphi \tag{4.18}$$

φ becomes the demodulation signal and, with careful selection of ω_m can be extracted from one demodulation path (although, strictly, two are necessary).

4.5 Alternative Locking Approaches

Pound Drever Hall locking is not the only method used to lock a laser to a cavity.

Dither Locking

An alternative approach to PDH locking is Dither Locking. This is illustrated in figure 4.4 by the small arrow near point A. The signal is dithered around the measurement point and variations in the power are used to indicate which direction to move in to reach the minimum. Dither locking uses and absolute measure of power so is not immune from laser fluctuations although the signal is dithered so fast they these do not usually matter. Dither locking differs from PDH locking from in two ways:

1. The modulation frequency used in PDH locking is much higher.
2. The side-bands in PDH are outside the locking peak.

As an example, the width of peak in the experiment reported in the following chapter is given the FSR divided by the finesse, namely 500 MHz/300 = 1.6667 MHz. A PDH locking approach might use a beat signal of 20MHz so the resulting side-bands are well outside peak. Dither locking, on the other hand, might run at 230 kHz so the side bands would be inside the peak bandwidth of the incident laser.

Tilt Locking

A third approach to locking a laser beam to a cavity is tilt locking, see Ref [74], which uses small misalignments of the incident laser beam with respect to the cavity's optical axis to create a non-resonant, non-symmetrical spatial mode which is then interfered with the incident laser beam to create an error signal.

Hansch-Couillard Locking

Hansch-Couillard locking is an approach which has the advantage of not requiring any modulation, see Ref [75] and the original paper Ref [76], but adds a requirement that a linear polariser be inserted within the cavity. The linear polariser or Brewster plate gives the reflected light acquires an elliptical polarization which depends upon frequency. The reflected beam is then split by a polarising beam splitter and the two halves compared to create an error signal which is used to adjust the piezo on the laser.

4.6 Higher Order Cavity Modes

So far we have only considered the fundamental mode resonating within a cavity, that is a beam travelling along the optical axis whose wavelength divides precisely into twice the cavity's length. Other modes are possible, however, because they are equally valid solutions to Maxwell's equations. The fundamental is just a special case.

These modes appear as either multiple spots, concentric rings or combination of both in the transmitted beam. Examples of a number of modes are shown in figure 4.11. These were recorded on a CCD camera during the set up for the experiments reported in the following chapters.

In deriving the solution for a Gaussian beam we started with the equation 4.1:

$$\tilde{E}(x,t) = E_i \cdot e^{ikx} \cdot e^{-i\omega t} \tag{4.19}$$

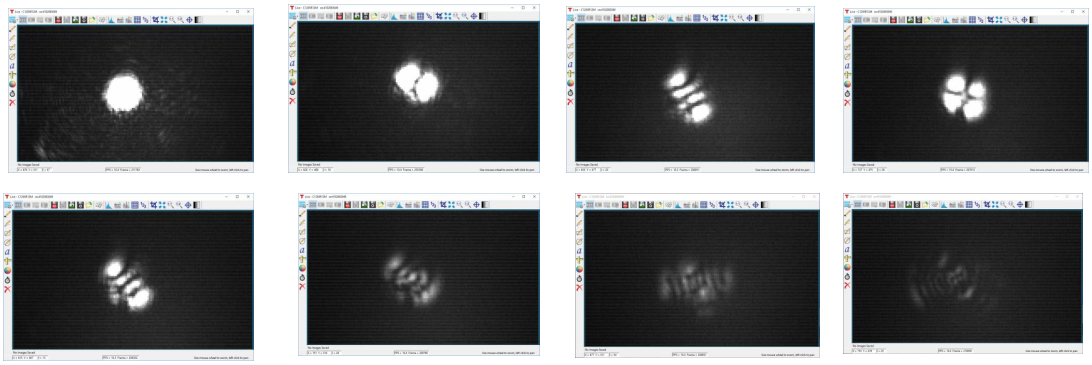

Figure 4.11: A selection of higher order Hermite-Gaussian and Laguerre-Gaussian modes captured during the alignment of the optics used for the experiments reported in chapter 6.

where $E(x,t)$ was the electric field of the incident beam, E_i was its amplitude, $e^{-i\omega t}$ defined a wave travelling in the z direction and e^{ikx} represented its spatial distribution across the direction of propagation and would be 1 for a plane wave. This is allowable as a solution because it satisfies Maxwell's equation for wave propagation, see Ref [68]:

$$\bigtriangledown^2\tilde{E} + \left(\frac{\omega n}{c}\right)^2 \tilde{E} = 0$$

After some manipulation the solution for the Gaussian form can be shown to be, see Ref [68]:

$$\tilde{E} = E_i \left\{ \frac{\omega_0}{\omega_{(z)}} e^{-\frac{r^2}{\omega^2_{(z)}}} \right\} \left\{ e^{-i\left[kz - tan^{-1}\left(\frac{z}{z_0}\right)\right]} \right\} \left\{ e^{\frac{ikr^2}{2R(z)}} \right\} \tag{4.20}$$

where $R(z)$ is the radius of curvature and noting that:

- E_i is the amplitude;
- $\left\{ \frac{\omega_0}{\omega_{(z)}} e^{-\frac{r^2}{\omega^2_{(z)}}} \right\}$ is the shape of the field as it moves away from the optical axis;
- $\left\{ e^{-i\left[kz - tan^{-1}\left(\frac{z}{z_0}\right)\right]} \right\}$ is the longitudinal phase factor; and
- $\left\{ e^{\frac{ikr^2}{2R(z)}} \right\}$ describes the curved phase front.

However, this is not the only form that satisfies this equation. Indeed it is a special case of some more generalised wave structures, in particular, the Hermite-Gaussian and Laguerre-Gaussian wave forms.

4.6.1 Hermite-Gaussian modes

The Hermite-Gaussian modes arise when a beam is slightly misaligned or offset. If a screen were placed in the transmitted beam just after the second mirror, they would appear as they do in figure 4.12, a series of bright spots in both the x and y directions, where the phases of adjacent spots in either direction are precisely π out of phase with their immediate neighbours. They are termed transverse electromagnetic modes and given two subscripts, TEM_{mn}, where the subscripts indicate the number of points of zero illumination (spaces) between bright spots in each direction. So TEM_{00} has no spaces, is a single bright spot and is the fundamental mode. TEM_{32} has 3 spaces between 4 spots in the x direction and 2 spaces between 3 spots in the y direction so comprises a grid of 12 spot in all. Figure 4.12 shows examples of this nomenclature (see Ref [77] and also Ref [78]).

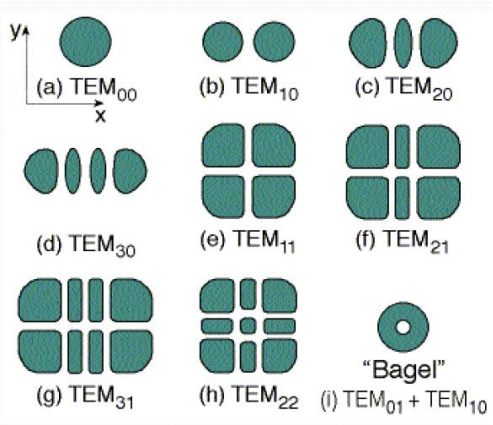

Figure 4.12: Transverse Electromagnetic Modes are classified by counting the number of spaces or points of zero illumination between spots in first the x and then the y directions. Adjacent spots are of opposite phase to each other. The 'Bagel" mode is the results of the TEM_{10} and TEM_{01} modes oscillating together. The diagram is drawn from Ref [77]

These modes arise because there are more forms of equation 4.20 which are also solutions to Maxwell's wave equation. One commonly found form involves including the Hermite polynomials into the answer, giving the more general result:

$$\frac{\tilde{E}}{E_i} = H_m\left(\frac{\sqrt{2}x}{\omega_z}\right) H_p\left(\frac{\sqrt{2}y}{\omega_z}\right) \left[\frac{\omega_0}{w_{(z)}} e^{-\left(\frac{x^2+y^2}{\omega^2_{(z)}}\right)}\right] \left[e^{-i\left[kz-(1+m+p)tan^{-1}\left(\frac{z}{z_0}\right)\right]}\right] \left\{e^{\frac{ikr^2}{2R(z)}}\right\} \tag{4.21}$$

Here the changes that have been made are:

- the conversions of the radial coordinate r^2 in the third term into $x^2 + y^2$ so that changes in x and y can be explored;
- The addition of $H_m\left(\frac{\sqrt{2}x}{\omega_z}\right)$ and $H_p\left(\frac{\sqrt{2}y}{\omega_z}\right)$, the two Hermite polynomials; and
- The addition of a $(1+m+p)$ factor in the fourth term where m and p, specify which polynomials to to be applied in the x and y directions respectively.

The Hermite polynomials are of the form:

$$H_m(u) = (-1)^m e^{u^2} \frac{d^m e^{-u^2}}{d^m u} \tag{4.22}$$

where $\frac{d^m}{d^m u}$ is the m^{th} derivative with respect to u. So when:

- $m = 0$, $H_0(u) = 1$;
- $m = 1$, $H_1(u) \sim u$;
- $m = 2$, $H_2(u) \sim 2u^2 - 1$;
- and so on.

and m becomes the number of nulls along a given axis. Figure 4.13 shows the intensities of the first three Hermite Gaussian modes across the x-axis with the direction of beam propagation being vertical. The figures show how adjacent spots are regions of opposite phase. The Hermite solutions are orthogonal and more complex modes can be represented by weighted sums of them. The black curve in the fourth figure shows the superimposition of the three modes and illustrates how quickly the emerging beam can appear chaotic is multiple modes are allowed to become resident.

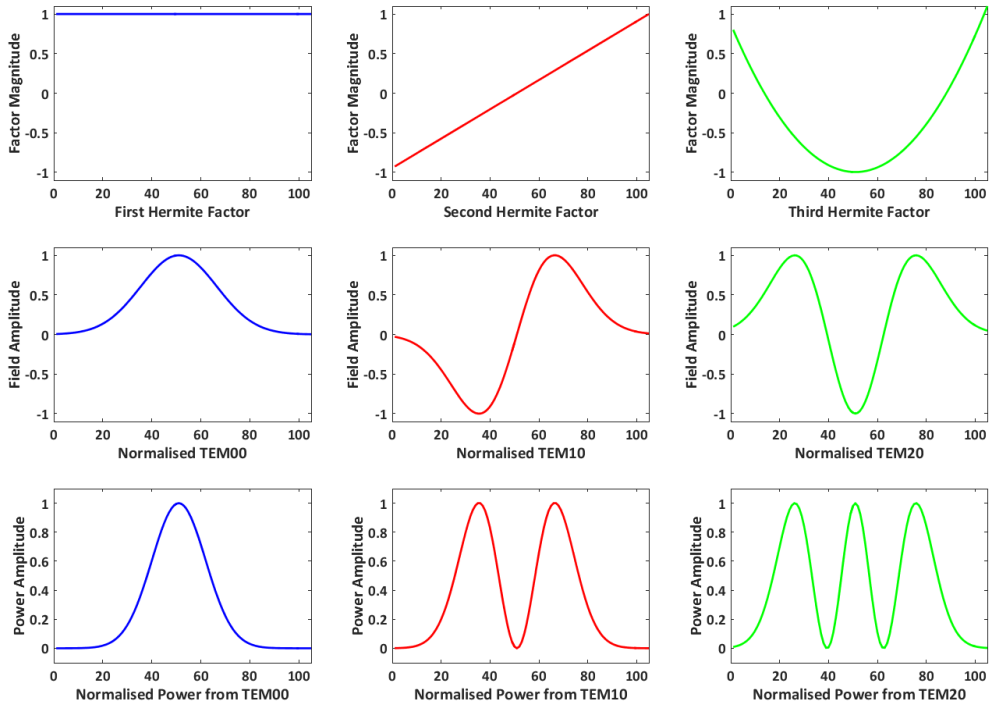

Figure 4.13: The Transverse Electromagnetic Modes TEM_{00} mode (blue) , TEM_{10} mode (red) and TEM_{20} mode (green) showing how adjacent bright spots are of opposite phase. The first row plots $H_0(u) = 1$, $H_1(u) \sim u$ and $H_2(u) \sim 2u^2 - 1$. The second shows how these affect a Gaussian electric field while the third shows the power arising from this field. Plots have been normalised. TEM_{00} produces one spot, TEM_{10} two spots of opposite phase and TEM_{20} three spots of alternating phase. When modes coexist, the emerging beam can quickly take on a chaotic appearance as the modes mix.

4.6.2 Laguerre-Gaussian modes

There is another set of solutions which give rise to mode solutions which appear as concentric circles. These are the Laguerre-Gaussian modes and represent a different class of orthogonal solutions to Maxwell's equations, see Ref [79] and Ref [80]. Figure 4.14 shows the intensity distribution for some of these modes which arise when the resonator has mostly circular symmetry, see Ref [81] Typically they are created when the waist is too large or small, or is in the wrong position (see page 2946 of Ref [82]).

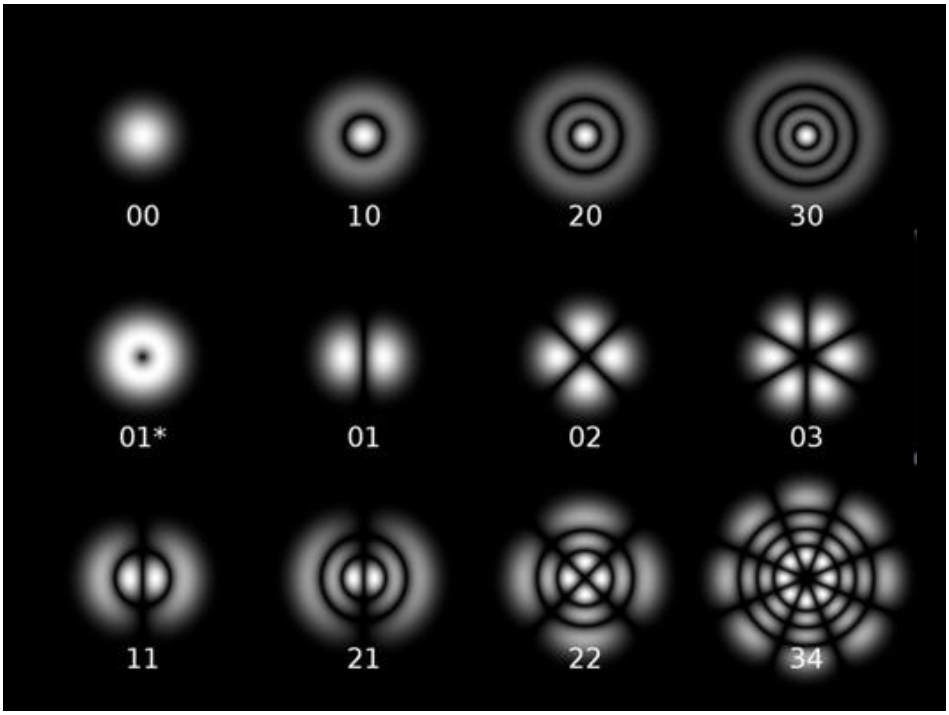

Figure 4.14: The Laguerre-Gaussian Modes. If a resonator has mostly circular symmetry, the modes exhibit a cylindrical symmetry described by the Laguerre polynomials. This figure shows the intensity distribution for modes of this type and is drawn from Ref [60].

4.7 Summary

This chapter is another background chapter and reviews those aspects of cavity dynamics that are important for the experiments presented in Chapter 6. It focuses on how the power transmitted though a cavity can be maximised by fine tuning the position (tilt and offset) and shape (waist position, waist size and wavefront shape) of the incident beam. Using a simulation, it demonstrates the sensitivity of waist position with respect to wavefront curvature.

It goes on to review PDH locking, because this is used in the experiments, and other locking approaches because it may be possible to employ one of these, in particular dither locking, to automate cavity optimisation further.

Finally it looks at the derivation of higher order modes, in particular Hermite-Gaussian modes, because the final experiment, presented in Chapter 6, selects and reinforces one of these modes for preferential transmission.

Wavefront Sensing and Correction

This chapter describes an experiment which creates a detailed map of a laser beam wavefront in near real time. It was reported in Ref [17]. The chapter starts by explaining why this is desirable and goes on to describe the optical layout, electronic arrangements and data processing, paying particular attention to the challenges involved in operating the Spatial Light Modulator (SLM). It concludes by showing how this technique can be used to not only detect but correct a wavefront, an approach applied in the second experiment reported in Chapter 6. Samples of the LabVIEW and Matlab code employed are included in the appendices.

The experimental work reported below was published in Applied Optics early in 2016, see Ref [17]. I designed, set-up and ran the experiment, including the software, on my own. I had assistance from Dr Paul Altin, debugging the software and determining noise contributions.

5.1 Motivation

The search for gravitational waves has made unprecedented demands on interferometry, as experimenters have strived to reduce noise and increase sensitivity. Auto-alignment systems are a key part of the laser interferometers used for gravitational wave detection. For ground-based systems, such as the Laser Interferometer Gravitational-Wave Observatory (LIGO), the challenge is to maximise power within the arm cavities by aligning the mirrors so that coupling into higher-order modes is suppressed and only the fundamental mode resonates (see Chapter 4). Space-based interferometers face a different problem. They operate across very long distances (arm lengths of 2.5 million kilometres are proposed for the Laser Interferometer Space Antenna, LISA) and the signals are so faint that small misalignments can lead to a loss of connectivity between satellites. The Gravity Recovery and Climate Experiment (GRACE) Follow-On planned for 2018 and the full LISA mission, currently planned for 2034, will both need to combat satellite jitter to maintain signal lock. All these problems are addressed by sensing and then correcting for wavefront tilt Ref [17]. Dr. Danzmann and his colleagues have investigated these issues in detail (see Ref [83], Ref [84], Ref [85] and Ref [86]) and concluded that even at the 2.5 million km separation proposed for the LISA satellites (see Ref [86]), the precision of today's metrology means that tiny aberrations in the wavefront can still combine with satellite jitter to couple into the tilt measurement, Ref [87] and Ref [88]. A number of increasingly elaborate models have been created to try to address this, Ref [89], but systems still are still limited by the linear approximations that apply to measurements with segmented diodes.

Wavefront sensing in LIGO is accomplished using the technique (The Ward Technique)

developed by Morrison et al, Ref [27] and Ref [26] which uses multi-segment diodes to determine the radio frequency (RF) beat-note amplitude difference between the two halves of the wavefront and develops a control signal to adjust optical components and remove this difference Ref [55] and see section 2.2.1. The technique is well established and very accurate, but is limited by the number of segments on the diodes used to map wavefront tilt and therefore assumes that tilt to be linear. As explained in Chapter 3, LIGO also uses Hartmann sensors which measure higher-order distortions in the wavefront caused by thermal distortions of the mirror surfaces Ref [32]. These use a fine array of precisely positioned holes to project a small sample of a test laser beam onto a charged coupled device (CCD) and infer the shape of the wavefront by measuring the displacements of the individual image spots from their expected positions Ref [15] and figure 2.6. This technique can achieve an RMS wavefront error of $\lambda/700$ in a single-shot, but relies on having a pre-recorded reference wavefront for comparison, which must be updated offline. It has also been used to identify flaws in the test masses which were not otherwise visible, Ref [90], Ref [91] and Ref [92] for the general approach and Ref [33] for specifics.

In space, the current approach is to sense minute changes in wavefront tilt using segmented photodiodes (see Ref [93], Ref [94], Ref [82] and Ref [95]) and realign the receiving satellite so that it can reflect the signal back to the originating satellite successfully. This approach is also accurate and effective but suffers because signals are recorded by separate channels and then compared. Differential phase shifts in these channels produce errors. These may be due to temperature fluctuations as the instruments move in and out of the Earth's shadow (GRACE) or other, internally generated temperature variations (LISA Pathfinder and LISA) Ref [96]. For example, the analogue-to-digital converters (ADCs) exhibit 1/f noise as their capacitances changes with temperature, decreasing the accuracy of the realignment signal. The capacitance of the photodiode segments themselves may also change with temperature, giving rise to spurious phase shifts between the segments which are indistinguishable from wavefront tilt. There have been attempts to reduce these errors using common pilot signals, but these must be separated in frequency from the heterodyne signal Ref [14] and so react differently to the analogue characteristics of the equipment, reducing their effectiveness.

This experiment presents a new technique based on a single-element photodiode and the spatial modulation of the local oscillator light. It encodes spatial information onto the light and uses this to separate the phases of different parts of the beam in post processing as outlined in Chapter 3. This technique shifts the complexity from the detection hardware into deterministic digital signal processing. Notably, the use of a single analogue channel (single photodiode and ADC) avoids the low-frequency error sources identified above. The technique can also sense the wavefront phase at many points, limited only by the number of pixels on the spatial light modulator in contrast to the standard 4 points from a quadrant diode. For ground-based systems, the new technique could be used to identify and eliminate higher-order modes, while, for space-based systems, it provides a measure of wavefront tilt which is less susceptible to low frequency noise. The technique involves only minimal changes to the physical set up of the alignment system of a satellite interferometer, Figure 5.1.

The experiment uses 1064 nm laser light which is in the near infra-red and is invisible to the human eye.

5.2 Experiment Goals

The goals of the experiment were three fold:

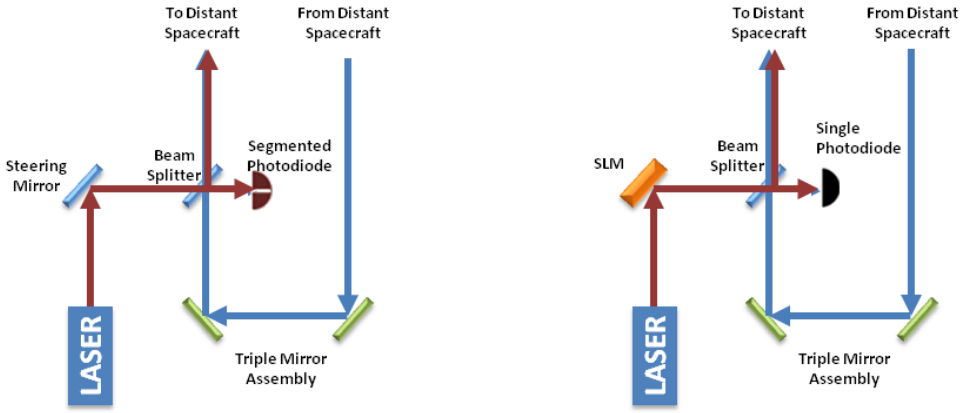

Figure 5.1: A comparison of the current (left) and proposed (right) arrangements for space based interferometers. In the proposed set up, the steering mirror and segmented photodiode have been replaced by an SLM and single photodiode respectively. The separation of the wavefront phases is done in post processing.

1. Demonstrate that digitally enhanced heterodyne interferometry (DEHI) can be used to develop a detailed map of a laser beam wavefront with a single element photodiode.

2. Compare the performance of DEHI with the performance of the tradition, segmented photodiode approach and, in particular, to compare the noise performances of the two approaches.

3. Demonstrate the capacity to not only to detect the shape of the wavefront but, using these measurements (or indeed measurements made with a Hartmann sensor), make corrections to it as desired.

Conceptually the approach was straightforward, see 5.2. A laser beam was split and sent through the two arms of a Mach-Zehnder interferometer. One beam, the local oscillator, was passed through an Acousto-Optic Modulator (AOM) and directed onto a segmented spatial light modulator (SLM) where it was phase modulated with spatially-dependent codes while the other, the signal, was passed through a different AOM and then directed onto a fast steering mirror (FSM) which could apply wavefront tilts of known amplitude and frequency. They were recombined at a beam splitter where the difference between the frequencies of the two AOMs produced a radio frequency beat note which was sent to two segmented photodiodes and a single photodiode. The SLM was driven by a Multi-driver controller which in turn was driven by a Matlab program running on the host PC. Data from the photodiodes and the sync pulse were collected in a LabVIEW ADC and passed to the FPGA which did some initial processing (channel splitting, demodulation and integration) before passing the results to the host PC for analysis and data presentation and capture.

It is worth noting that encoding the local oscillator is mathematically the same as encoding the signal path.

At the highest level, the reason for undertaking this experiment was to see if DEHI can provide an alternative approach to wavefront sensing that has comparable or better noise

performance to the conventional approach but offers a simpler optical set up, in particular, less sensitivity with respect to the alignment of the signal beam onto the photodetector.

The experimental set up is explained in detail below.

5.3 Optical Set Up

The experimental layout is shown schematically in Figure 5.2. The beam from a 1064nm laser (Innolight Prometheus) was passed into a Mach-Zehnder interferometer where it was split and passed through two acoustic-optic modulators (AOMs) driven by modulation signals of around 80 MHz and separated by the heterodyne frequency of $\omega = 2\pi \cdot 160$ kHz. The heterodyne frequency was selected to be as far as possible from the code modulation frequency (around 2 kHz) while remaining well within the 500 kHz Nyquist frequency of the ADCs.

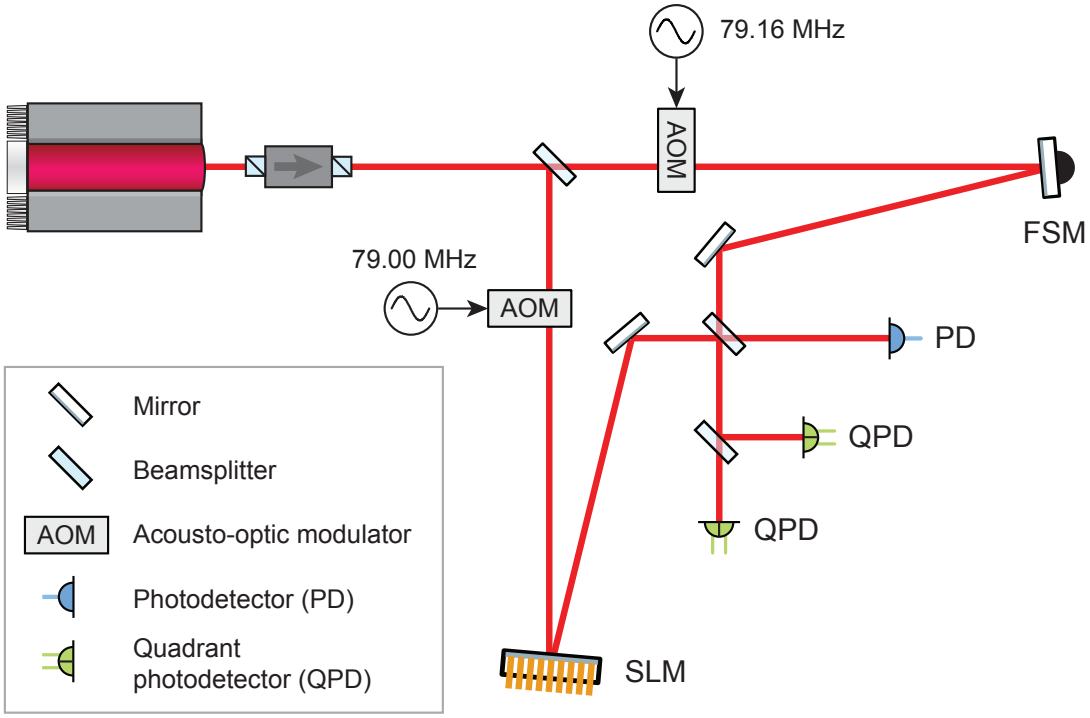

Figure 5.2: A schematic of the experimental layout. A 1064nm laser beam is split and passed through two AOMs to create a heterodyne beat note of 160 kHz. The signal beam (on the left hand side) reflects off a spatial light modulator which tags each half of the beam with orthogonal codes at 2 kHz. The local oscillator strikes a fast steering mirror which can be used to apply a controlled wavefront tilt. The beams are then interfered and directed to both single and quadrant diodes. The diode signals are decoded on an FPGA to recover the wavefront phase.

One beam, denoted the signal beam, struck the segmented SLM (see below) which modulated the left and right halves of the beam with orthogonal Hadamard codes [+1 -1 +1 -1] and [+1 +1 -1 -1] at around 2 kHz. (The Matlab code for this program is included as Appendix 1) The second beam, the local oscillator, was incident upon a fast steering mirror (FSM) which could be used to tilt the wavefront by a controlled amount and at a

controlled frequency to test the system's capacity to recover beam tilt. The beams recombined at the beam splitter and were then sent to two quadrant photodiodes and a single photodiode for comparison. The signals from each quadrant photodiode segment could be used independently, or added before processing to simulate a single photodiode. This allowed direct comparison of the new technique with the conventional segmented photodiode approach. The quadrant photodiodes were arranged perpendicularly to each other on the optical bench because the SLM, being a mechanical device, caused small vibrations on the bench and there was a chance that these might be more pronounced along either the long or short axis.

The signals from the photodiodes passed through ADCs and into a Field Programmable Gate Array (FPGA) (LabVIEW 7856-R) where the digital signal processing occurred. The size of the FPGA determined the maximum number of channels which could be processed concurrently. This particular FPGA reached maximum capacity at 64 simultaneous channels, but this could easily be increased by using a larger FPGA.

5.4 Electronic Set Up

The electronic set up was quite straight forward and involved five separate systems:

1. SLM Control - A USB cable was connected to the Multi-driver controller which, in turn, was connected by 4 ribbon cables which sent data in parallel to the SLM. (The SLM controller was initially mounted on the instrument rack as its cooling fan made it too noisy to be position on the optical bench. Even so the mechanical action of the SLM introduce enough noise to have an effect below 10Hz, see below.)

2. AOM Control - Each AOM was driven by a signal from an Agilent 33250A signal generator which was passed through a PA-4 RF power amplifier with a minimum gain of 40dB.

3. Fast Steering Mirror Control - The FSM was driven by its controller which in turn was driven by a signal from a third Agilent signal generator.

4. LabVIEW ADCs - signals from segments of the photodiodes and the sync pulse from the SLM controller were passed to a LabVIEW junction box which in turn passed them to the FPGA (8channels, 16 bit resolution and a 40MHz clock rate enclosed in a National Instruments USB-7856R LabVIEW chassis).

5. System Clock - the LabVIEW chassis used did not have a clock input so a 10MHz clock was developed in LabVIEW, passed out through a digital channel and then daisy-chained to synchronise the other equipment. The digital channel was power with a 5V power supply.

This information is summarised below in figure 5.3

During operation it became clear that there was some cross talk between these systems, particularly the leads from the three photodiodes to the LabVIEW control box. photodiodes were disconnected while not in use to minimise this feedback.

5.4.1 Managing the Spatial Light Modulators

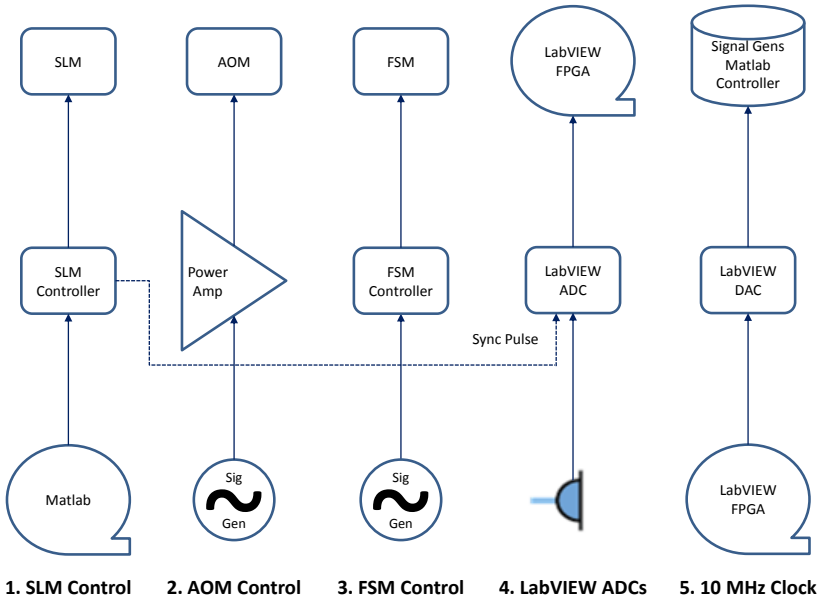

Figure 5.3: A schematic showing the 5 control systems used in the first experiment.

The Spatial Light Modulator

The spatial light modulator was the centrepiece of the experiment. It was a segmented spatial light modulator (specifically a Boston Micromachines Multi-DM 1.5μm) and comprised a 12 by 12 array of gold coated mirrors each of which could be advanced independently by up to 1.5 microns. (The mirror's in the corners were fixed in the 0% or flat position making the actual number of active elements 140.) It is a micro electromechanical systems (MEMS) device, packaged and wire bonded to a ceramic chip carrier and sealed in a window enclosure which has a slightly angled face to help deflect unwanted reflections away from the optic line. The sealed chip carrier is integrated into a compact aluminium enclosure which can be mounted in most standard 2-inch optical mounts (See Section 2.2.3 and Ref [97]) Elements were controlled by creating an array of 140, 14-bit numbers in Matlab and sending them to the SLM controller (see below) which converted them to voltages and passed them in parallel to the SLM where they initiated a change of state. Numbers in the array were percentages with 0% leaving an actuator in the flat position, 100% generating the full 1.5 μm and values in between generating a proportionate response. The controller recognised 14-bit numbers suggesting a minimum actuator step of 1.5μm/16384 or 0.0916nm. (The technical specifications guarantee only movement of less than 1nm but even this was more than adequate for the experiments reported here because of the low dependency on modulation depth.) The attributes of the SLM used in this experiment are shown in Table 5.4.

The SLM Controllers

Each SLM is managed by a controller, figure 5.5 and Ref [97], which has the job of converting the array of numbers created by Matlab or LabVIEW into the voltages which drive the actuators and passing them to the SLM through a parallel cable. Each SLM is matched individually to its controller by a driver which, besides facilitating data transfer

Attribute	Segmented SLM
Boston Micromachines Designation	Multi-DM 1.5 µm SLM
Maximum Stroke	1.5 µm
Surface	Segmented
Actuator Pitch	300 µm
Active Aperture	3.6 mm x 3.6 mm
Step Size	<1nm
Step Granularity	0.0916nm
Surface Covering	Gold
Surface Quality	< 30 nm RMS
Inter-actuator coupling	0%
Power limits (Unpublished)	100W/cm^2
Frame Rate	8kHz with 34kHz burst
Tested Continuous Max Frame Rate	10.374Khz

Figure 5.4: The performance specifications of the Segmented SLM.

also adjusts the output voltages to match the specific characteristics of the SLM so that it performs within the tolerances promised. Drivers, SLM and controllers must be used as a matched set for accurate performance. The calibration for the SLM used in this experiment is provided at Ref [98].

Controllers can be either Multi Drivers or X-Drivers. The Multi Driver is linked to the controlling PC using a USB connection. It is cheaper but introduces a degree of uncertainty about the arrival times of new arrays because the USB channel competes with other resources, being managed by the operating system (Windows 10), and executes with non-deterministic timing. In practical terms, this limited operational frequencies to approximately 2 kHz. X-Drivers are connected to the PC through a PCIe bus and are supposed to be able operate at up to 400 kHz although we have been unable to test this because, as yet, no suitable drivers are available, see Ref [99]. Given that the errors on modern PC clocks are of the order of 11 PPM, see Ref [100], and the PC driving the SLM controller used a 2.7 GHz clock, the effect of PC clock jitter on the X-Driver would be minimal. As always budget was a consideration and so a Multi-driver was used in the proof of concept and the timing uncertainties had to be managed (see Section 5.4.1).

Determining the Sampling Window

The SLM is a mechanical device and, when it is connected through a Multi-driver to a host PC, especially when it is running other software (typically, the LabVIEW Host and the Matlab mirror controller), its operating frequency is variable. It can run reliably at around 2kHz but this remains a variable operating frequency. The heterodyne frequency of the beat note between the two AOMs was selected at 160kHz and, consequently, there were around 80 cycles of the heterodyne frequency for each code change. (The examples in Chapter 3 were based on multiple code changes per heterodyne cycle but either arrangement works equally well.)

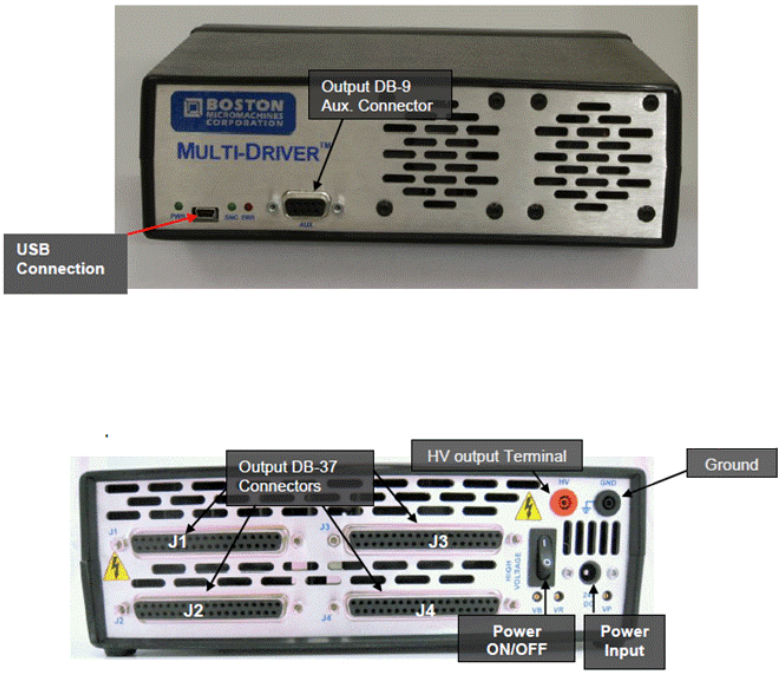

Figure 5.5: The Segmented SLM USB Controller.

A new mirror shape was initiated as soon as the controller received a new array. Samples could only be taken when the system had settled into a stable state, that was after the SLM had changed shape and settled, and before the USB interface sent the next array. This proved more difficult than first thought.

The complications were threefold:

1. Accurately determining the true finish of the sync pulse.
2. Accommodating a settling time for the mirror.
3. Managing the uncertainty around the USB interface and the arrival time of the next sync pulse.

Every time the SLM changed state it sent a synchronisation pulse (sync pulse) to the LabVIEW ADC. The sync pulse manifested as a voltage increase on a signal line, which rose from 0V to 3.3 volts for $31\mu s$. The mirror had now transitioned to a new state. Experiment showed that this was followed by a settling time of up to 40 micro-seconds which depended on the magnitude of the actuator movement. Large movements oscillated for longer before they settled. Once the settling time had finished a window opened which allowed the collection of reliable data (that is, data not corrupted by transients in the actuators) and this continued until the controller provided the voltages to initiate the next change. This, of course, was also somewhat unpredictable so the data collection challenge was to identify a suitable data collection window after the SLM had changed state and settled and before a new transition started.

A typical cycle was as follows and allowed for reliable sampling for 75% of the total cycle time:

1. Sync Pulse - mirror has been sent an array and is changing shape - $31\mu s$.

2. Settle Time - mirror settled into new configuration - up to $40\mu s$ depending on the degree of actuator movement.

3. Sample Time - period of time during which reliable data could be collected.

4. Safety Buffer - safety margin which allowed for the shortest, likely time before a new sync pulse is received. (The occasional sample which overlaps the start of arrival of the next USB pulse is tolerable but it cannot be allowed to extend pass the end of the next sync pulse or the codes will get out of sync.)

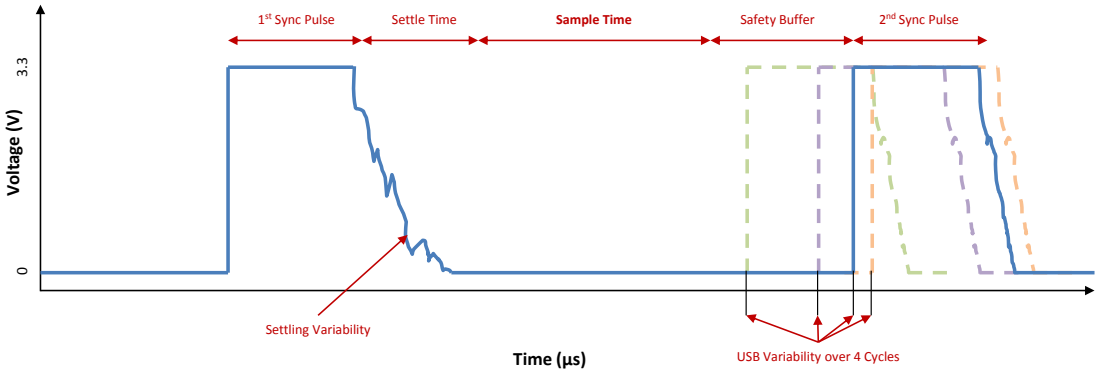

Figure 5.6: A typical SLM cycle showing how the variability of the settling time, and arrival of the next sync pulse reduce the safe sampling time to about 75% of the total cycling.

This is illustrated in figure 5.6 and summarised in table in figure 5.7. Ticks refers to FPGA cycles which occurred at 40MHz, so $1\mu s$ was 40 ticks. It was vitally important to keep track of all of these changes because, as figure 3.4 shows, the demodulation only works if the code used to demodulate remains in phase with the code used to do the initial tagging.

Stage	µs	Ticks	% of Total
Sync Pulse	31	1240	1.55%
Settle Time	219	8760	10.95%
Sample Time	**1500**	**60,000**	**75.00%**
Safety Buffer	250	10000	12.50%
Total	**2000**	**80,000**	**100.00%**

Figure 5.7: Typical sampling windows available for reliable data collection during the cycling of the SLM. The reliable sampling time was 75% of the total cycle. Ticks refers to LabVIEW cycles which ran at 40mHz. $1\mu s$ is 40 ticks.

The settle time and minimum time between sync pulses were determined by observation and rechecked regularly during data collection using an oscilloscope. Determining that a

sync pulse had been completed required the development of some more complex programming rules. Theoretically the sync pulse presented LabVIEW with a perfect square wave, dropping from 3.3V to 0V instantly. In practice, of course, this was not a perfect edge (even after digitisation) as is evident from figure 5.8 which shows a typical sync pulse and, in the circle, details of the falling edge. Unless managed properly, the bounces shown in the trailing edge of the sync pulse could trigger additional false sync pulses which would lead to the demodulation codes applied in the FPGA getting out of step with the original modulation codes which had been applied to the signal by the SLMs. To identify a real edge, we looked for a drop of at least 1.7V that was sustained for around 100ticks ($2.4\mu s$).

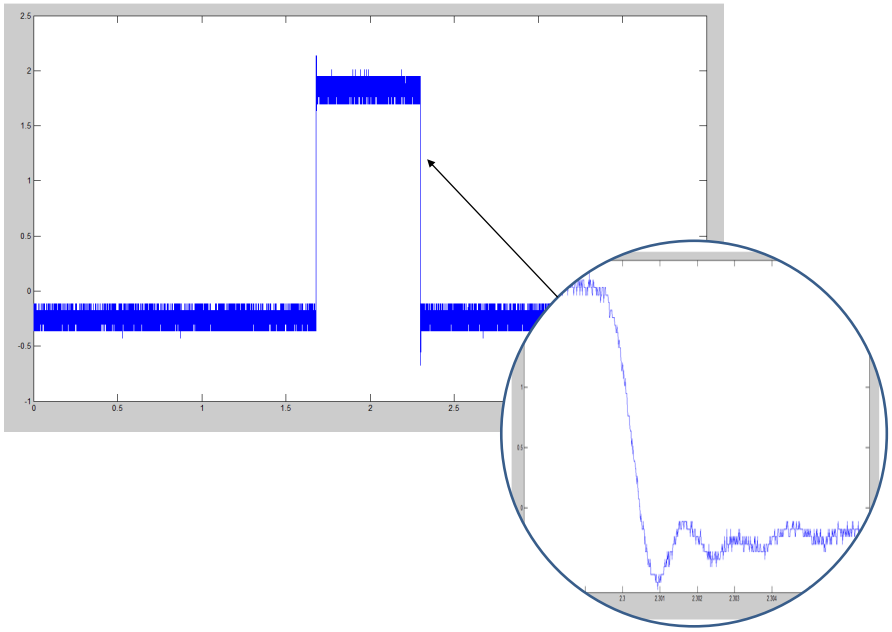

Figure 5.8: The typical shape of a sync pulse shows that it was not a perfect square wave. Unless managed properly these fluctuations could generate additional, false sync pulses which caused the modulating and demodulating codes to get out of step.

5.5 Data Processing

Both Matlab and LabVIEW were employed to manage the experiment - Matlab to drive the SLM and LabVIEW to collect and analyse data and provide the system clock.

5.5.1 Matlab

A Matlab program, see 8.1, was used to drive the SLM which, for the first experiment, required alternate halves of the mirror to operate independently with different orthogonal codes. The detailed code is presented in Appendix 8.1 but figure 5.9 provides a high level overview.

Two variables were entered at the start of the program - modulation depth and time delay. Modulation depth defined how far halves of the the SLM were advanced when they were

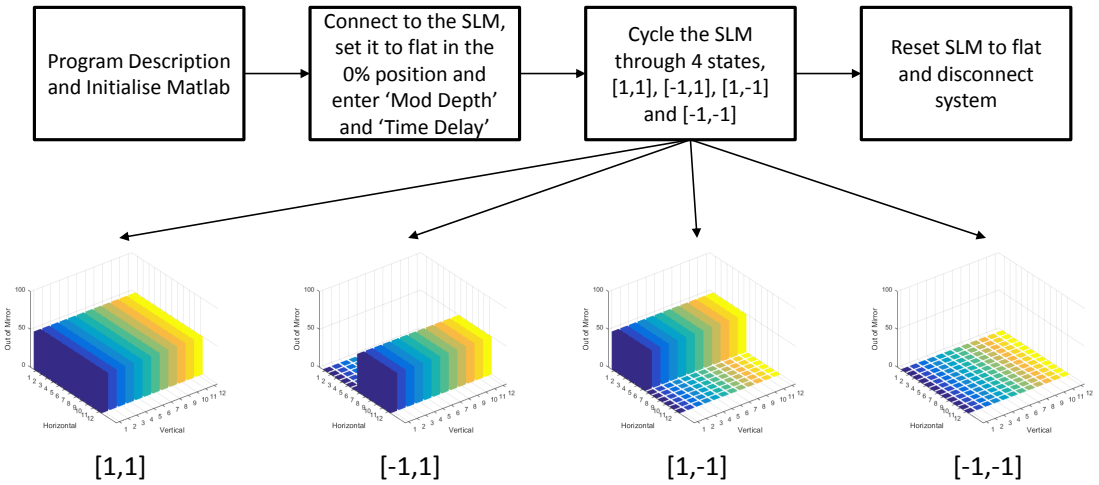

Figure 5.9: A schematic of the Matlab Code used in the first experiment.

in the [+1] state. 0%, left a segment flat (the [-1] state), 100% advanced a segment the full $1.5\mu m$ and a value in between advanced a segment proportionally. Time delay ran a counter after each change of state and the length of this count determined the approximate frequency at which the SLM changed. After opening the SLM and ensuring it was flat, the code cycled through the four states illustrated in figure 5.9 above.

5.5.2 LabVIEW

The LabVIEW system was used to collect and analyse the data. It also provided a 10MHz clock used to synchronise the equipment used in the experiment. It is outlined in some detail in Appendix 8.3 and explained in summary here. The LabVIEW code was split across a LabVIEW chassis (40MHz LabVIEW 7856-R) and the host PC. The LabVIEW chassis contained an FPGA (a Kintex-7 160T FPGA, 1 MS/s Multifunction Reconfigurable I/O Device) and was used for all the code that was time dependent. It:

1. Provided a 10MHz clock.
2. Collected all the data from the photodiodes.
3. Detected the true edge of the sync pulse, advanced the demodulation code and began data processing by:
   (a) Splitting the signal into two halves and executing the IQ demodulation
   (b) Splitting the two demodulated signals into halves again and executing the code demodulation.
   (c) Integrating the signals over a nominated multiple of the code length.
   (d) Creating an array of values for each channel and inserting -2^{13} into the first position to identify a new data set.
   (e) Passing the result into the FIFO for transfer to the host.

Figure 5.10 shows the basic processing loop for each pair of signal channels. The decoding process involves two demodulations which take place in the LabVIEW digital processing

engine, one using sine and cosine to separate the I and Q components and one using the codes to separate the different spatial regions. The process is illustrated in figure 5.11.

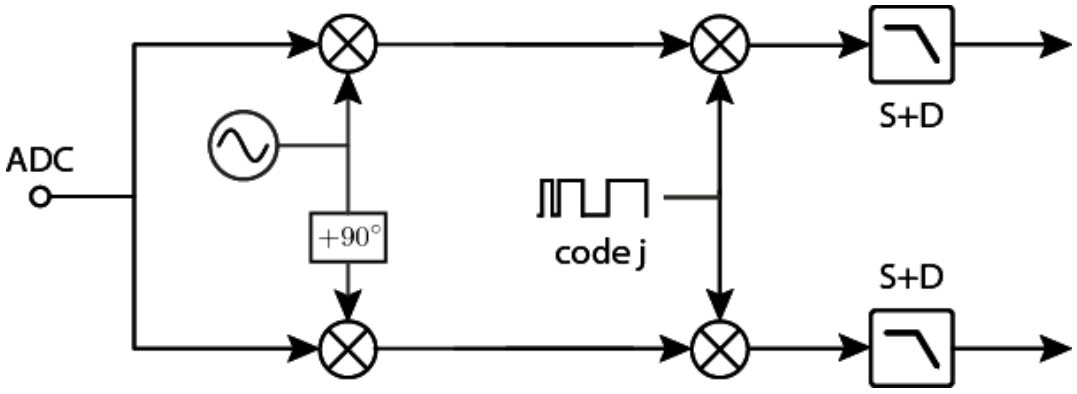

Figure 5.10: A block diagram illustrating signal processing for one pair of channels after collection in the single diode. This is repeated for each orthogonal code in the set. The input signal is first split into two and demodulated with a sine and cosine of the heterodyne frequency. These signals are then split again for each of the orthogonal codes used and each pair is multiplied by one of the codes. Finally all the channels are integrated across the full code length and the resulting outputs processed using arc tangents to determine relative phase of each spatial zone.

The blue traces in the top row of figure 5.11 represent a modulation free, heterodyne signal and the black traces on the second row the codes, c1= [+1 -1 +1 -1] and c2= [+1 +1 -1 -1]. Row 3 shows the result of the SLM imprinting the different codes onto the heterodyne signal and represents the components of the signal arriving at the single diode. Post processing occurs from this point. Row 4 shows the effect of demodulating these signals with a cosine of the same frequency as the heterodyne frequency and Row 5 shows the final effect of decoding these new signals by their respective codes and then averaging them across multiples of the code length (the blue line). Notice that when the code used in demodulation matches the code imprinted by the SLM (left hand column), the average is non-zero and can be used with its pair (the sine demodulation) to recover the phase while other, orthogonal codes (right column) which do not match the imprinted code average to zero and their signals disappear.

The application of and demodulation with the codes has allowed separation of the components of the single signal captured in the single photo diode even though its components were seemingly inextricably mixed together.

Figure 5.12 shows a screen shot of the front panel of the FPGA LabVIEW program used for the first experiment. It is designed to facilitate comparison of single element and multi-element wave front sensing.

While samples of code are presented in detail in the Appendices 8.2 and 8.3 it is useful to review the inputs and outputs here. The inputs, many of which were uploaded from the LabVIEW Host program, can be grouped as follows:

1. Stop and Resets - Controls which either stop all operational processes or reset Sync Pulse count or FIFO after overflow errors.

2. Edge Detection - *Sync Threshold* and *Sync Pulse Width* set the parameters that determine when a change in the Sync Pulse Voltage will be recognised as the conclusion of a new sync pulse, the code is advanced and a new round of signal processing

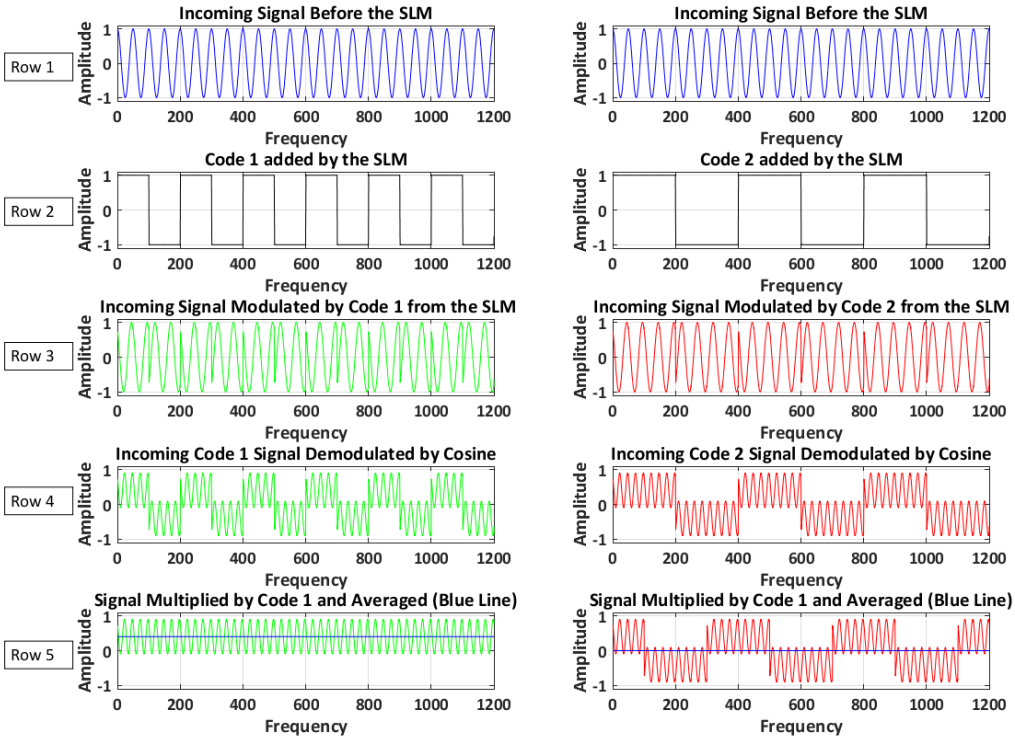

Figure 5.11: Simulation data showing a progressive view of modulation and then demodulation of a heterodyne signal using orthogonal codes. The left hand column shows a signal modulated and then demodulated with the same code and averaging to a non-zero value. The right hand column shows a signal modulated with one code, demodulated with another and averaging to zero. Row 1 shows the modulation free, heterodyne signal and Row 2 the codes, in this simplest case codes c1= [+1 -1 +1 -1] and c2= [+1 +1 -1 -1]. Row 3 shows the result of the SLM imprinting the different codes onto the heterodyne signal and represents the components of the signal arriving at the single diode. Post processing occurs from this point. Row 4 shows the effect of demodulating this signal with a cosine of the same frequency and Row 5 shows the final effect of decoding these new signals by their respective codes and then averaging them across multiples of the code length. Notice that when the code used in demodulation matches the code imprinted by the SLM (left column), the average is non-zero and can be used with its pair (the sine demodulation) to recover the phase while other, orthogonal codes (right column) which do not match the imprinted code average to zero and their signals disappear.

begins. *Code offset* facilitates the alignment of the code used for modulation by the SLM with the code used for demodulation by LabVIEW.

3. Sampling Control - *Sum Delay* and *Sum Duration* determine when the sampling period starts and how long it lasts. *S&D Scaling Factor* ensures that the signal produced by the integration does not overflow when it is transformed back into a 16 bit number for transfer to the FIFO.

4. ADC Voltage Levels - Selects the sensitivity of the ADC for each segment of each diode and the sync pulse

5. Live or Simulated Operation - The first of these toggle switches determine whether the system is run with live or simulated data. The second elects whether artificial noise is added to the signal before processing.

6. Signal Amplification - Allows digital amplification of signals from each photodiode in the segment before they are presented for processing.

7. Analysis Selection - Allows selection of signals from different parts of the processing train for direct pass through to an oscilloscope. This proved invaluable for system diagnosis.

8. Mode of Diode Operation - Selects whether the photodiodes are treated as segmented diodes, in which case the input from their segments were processed separately) or a single diode, in which case the signals from their segments are added before being processed.

9. IQ Demodulation Control - *LUT Step Size* determines the frequency of the Sine and Cosine used for IQ demodulation while *Scale Sine* and *Scale Cosine* adjust their amplitudes. The look up tables which define the sines and cosines used for demodulation in LabVIEW have 10,000 entries. The LabVIEW clock runs at 40MHz so advancing one position per tick means the system cycles through the look up tables 40000000/10000 times a second generating waves of 4,000Hz. The heterodyne frequency is 160kHz so to match it, we need to cycle through the look up tables 40 times faster. Hence the LUT value of 40.

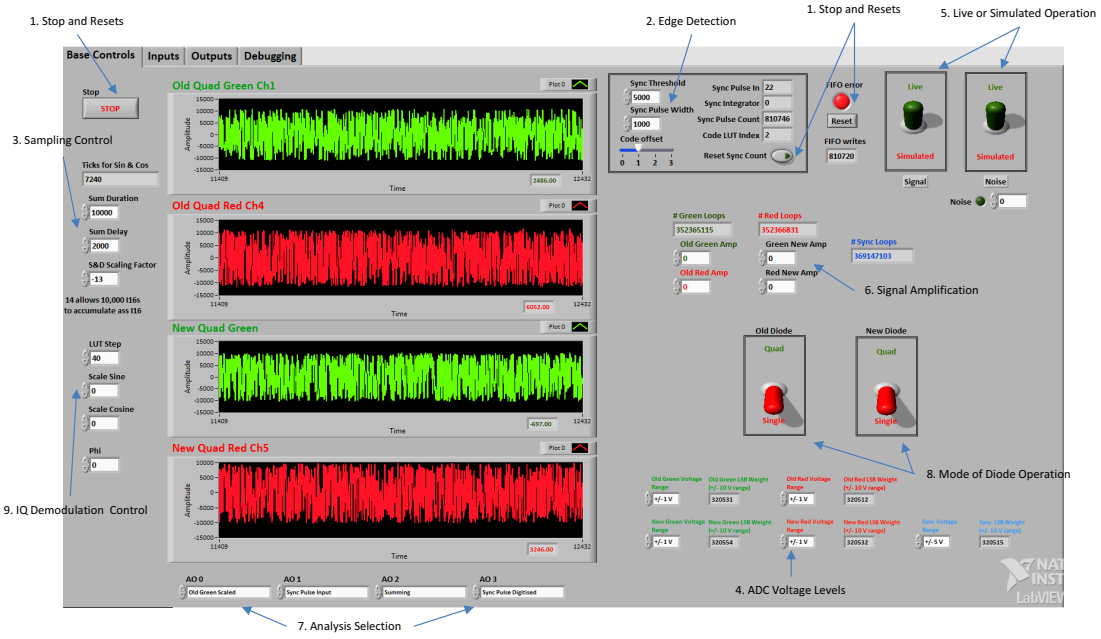

Figure 5.12: The front panel of the FPGA LabVIEW program used for the first experiment.

The panel also showed the values selected and traces from segments of the two segmented photodiodes termed Old Quad and New Quad (for no particularly good reason except the order in which they were introduced to the experiment.) Similarly the green signals comprised the sum of two left hand segments of each photodiode while the red signals comprised the sum of the two right hand segments. When the behaviour of the single diode was required the signals from all for segments were added together before processing.

Panels on other tabs provided a variety of diagnostic traces which were, frankly, of limited use given the speed with which the FPGA executed and the limited capacity of the host to present these traces in real time.

The system was controlled by a LabVIEW instance running on the host PC. The front panel is shown in figure 5.13. Most data input was controlled from the host and the values were passed through to the FPGA so do not need to be re-described here. The *Nominal SLM Frequency* input deserves special mention. It was used to define the scale for the Timeseries plot, recognising that slight variations were not issue provided they were common across all the traces.

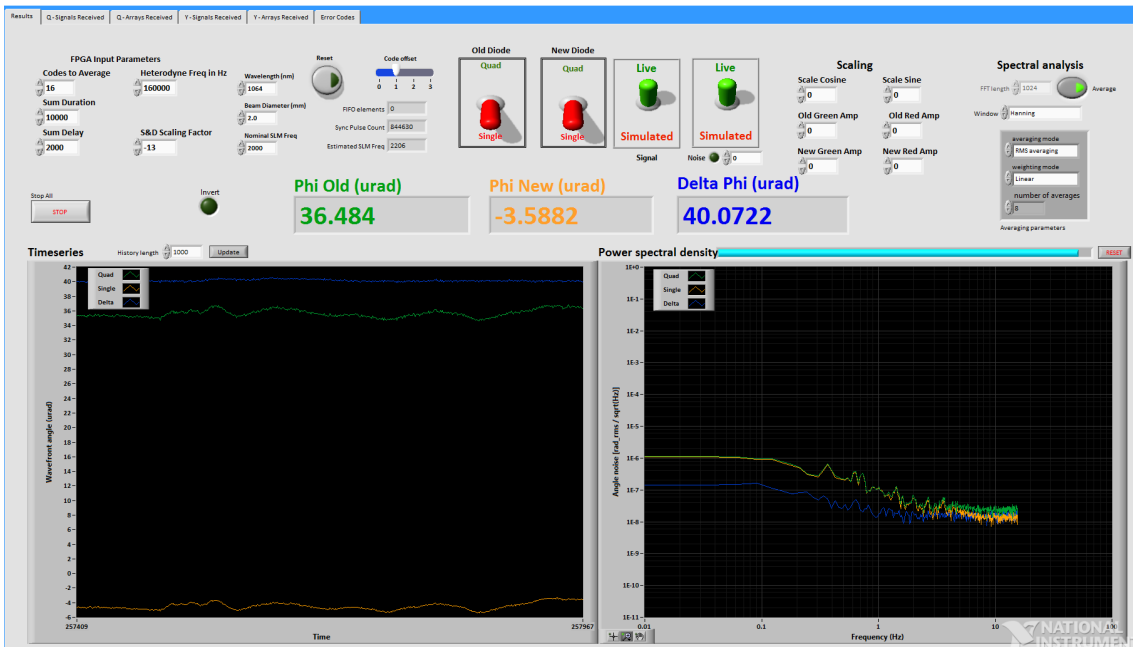

Figure 5.13: The front panel of the FPGA LabVIEW program used for the first experiment. This program controlled the LabVIEW FPGA program and captured and displayed the data collected during experimental runs.

The screen presented the processed data as a time series of the wavefront tilt (Phi) measured by the segmented photodiode acting as both a segmented photodiode and a single photodiode. It also provided the Power Spectral density of the outputs. The Fast Fourier Transform supporting this was controlled by the controls in the top right hand corner which defined its length and the window used to truncate it (A Hanning window). Further tabs provided diagnostic signals.

Other Operational Considerations

Code Alignment

Maintaining the alignment between the codes used to modulate the mirror with the SLM and demodulate the signal with LabVIEW was essential if the individual signals were to be recovered successfully. As figure 3.4 illustrates, demodulation ceases to work once the codes get out of step. Importantly there was no direct connection between the Matlab program controlling the SLM (which ran at a variable frequency) and LabVIEW (which ran at a fixed rate). The only synchronisation available was provided by the sync pulse

which informed LabVIEW that the mirror had changed state and a new code was required. This was accomplished using a variety of techniques:

- The codes were initially aligned by counting the set up events that caused the SLM code to advance and adjusting the Code Offset slide to match.
- Ensuring that only true sync pulses were counted by developing the edge detection routines outlined above.
- Ensuring that the first element of each demodulation array passed to the host through the FIFO began with the number -2^{13} as a reference point.
- Providing general reset capabilities for occasions when the FIFO overflowed or the system lost synchronisation.

It is important to recognise that this is a purely operational issue and note a problem inherent in the technique. If, for example, an X-Driver controller had been available, the operation of the system would have been far more deterministic and the code alignment problem would have been much simpler to manage.

Mechanical Vibrations

SLMs are mechanical devices and their vibrations impacted the optical bench. Two segmented photodiodes were used set at right angles to each other along the long and short axes of the bench. This allowed their signals to be subtracted to remove much of the common noise due to these vibrations. When measurements were made for the segmented photodiodes (rather than when they were used to simulated a single photodiode) the SLMS were turned off so that all the vibration was removed. Any final system would have to address this issue by decoupling mechanical components from the optical bench.

Simulating the Single Diode for Comparison

In order to keep the comparison between the two techniques as fair as possible, the quadrant photodiodes were used to simulated the single photodiodes. This was accomplished by flipping a switch on the LabVIEW Host screen which added the signals from all segments before they were delivered for processing. This understated the performance of the single diode slightly in two ways:

- It ignored the gaps between the segments of the multi-element diode, reducing the detecting area where the signal was at its strongest, and
- The performance of the particular multi-element photodiode employed was smaller and less sensitive that a single photodiode.

Modulation Depth

At full extension, an actuator in the SLM could move forward by $1.5\mu m$ (100%) which would advance the phase of a beam striking it almost normally by $3\mu m$ or nearly 3λ (6π) for 1064 nm laser light. Of course, advancing the phase by a multiple of 2π was pointless as it simply moved the wave forward by an integer number of cycles. The ideal depth was around $\pi/4$ or 4%. Given that experimental constraints did not allow beam to strike the SLM normally, the apparatus was arranged so that it struck the reflecting surface of

the SLM at around 10° from normal. As figure 5.14 shows, the movement required in the mirror to increase the path length by $\pi/4$ for $\theta = 10°$ is x where:

$$\begin{aligned} x &= ycos(\theta)\text{nm} \\ &= 0.9848y\text{nm} \\ &= 0.9848.1064/16 = 65.4898\text{nm} \end{aligned}$$

$$x \equiv \frac{65.4898 \cdot 100\%}{1500}\% = 4.36\% \tag{5.1}$$

In practice the system proved remarkably insensitive to modulation depth as it normalises during the final demodulation, see equation 3.12. For this experiment, we were only interested in the result after demodulation and integration so a modulation depth of 8% or $\lambda/4$ was used, offering a clear and unambiguous change in phase modulation. As the experiment was designed to compare the performance of the conventional segmented photodiode approach to the DEHI based, single photodiode approach and any modulation was simply noise as far as the conventional approach was concerned, all modulations were turned off when results for the conventional approach were collected.

In experiment 2, see chapter 6, there was no comparison between techniques and the modulation depth was reduced to 0.2% to minimise noise while still providing enough depth to facilitate phase modulation.

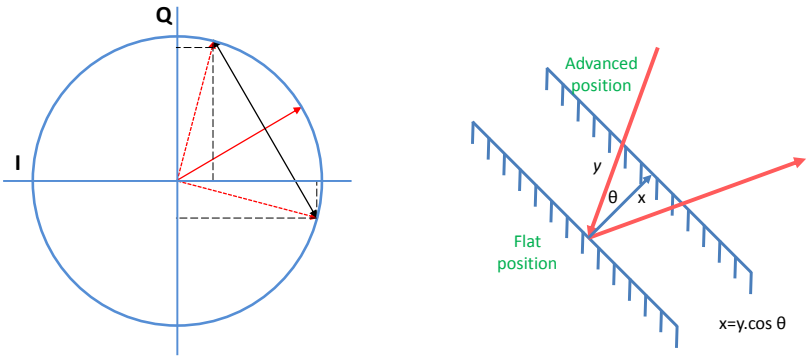

Figure 5.14: The ideal modulation depth is around $\pi/4$. After adjusting for the 10° angle of incidence or 4.36%. In practice the system was not particularly dependent on modulation depth as it divides away during the final demodulation.

The depth of modulation was zero when the segmented photodiodes were used in the conventional manner as the SLM was turned off.

5.6 Results

As mentioned above, the aim of this first experiment was to compare the wavefront sensing performance of the new technique to the conventional technique based on segmented photodiodes. Figure 5.15 shows the horizontal tilt angle of the wavefront calculated using each approach. The fast steering mirror was used to periodically apply a sinusoidal tilt of approximately $100\mu rad$ to the local oscillator beam. This tilt corresponds to a phase

difference between the two halves of the beam of nearly 1 radian. Both techniques are able to recover this signal, with a difference of less than $5\mu rad$.

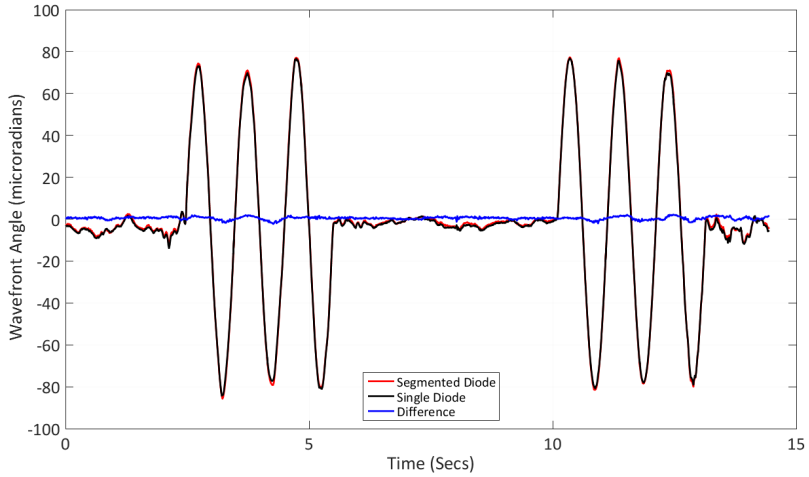

Figure 5.15: The Performance of the Single and Segmented Diodes Compared. The figure illustrates the recovery of a sinusoidal tilt of the wavefront introduced by a fast steering mirror at 1Hz. The red trace shows the measurement made by the segmented diode (the conventional approach) and compares it to the black trace, that measured by the single diode (the digital interferometric approach). The FSM is turned on and off showing that both systems can recognise and track the wavefront tilt it introduces.

Figure 5.16 shows the noise contributions and relative performance of each method. The blue and black curves show the frequency spectra of the phase difference between the two halves of the beam, as measured using the individual segments of the quadrant diodes (top) or by summing the quadrants together to mimic a single diode (bottom). The dominant source of wavefront fluctuations is environmental noise in the laboratory. The red curves show the difference between the measurements from the two diodes and represent environmental noise that is not common-mode. The effect of diode dark noise (green curve) was determined by blocking the light to the diodes and adding the signals from the ADCs to a simulated signal generated inside the FPGA. The ADC noise (cyan curve) was measured by adding the signal from terminated ADCs to the internally generated signal. When using the diodes as quadrant diodes, the segmented mirror was flat and switched off to ensure that mechanical noise from the segmented mirror does not contaminate the measurement.

Figure 5.17 shows a direct comparison of the difference between two measurements of relative phase using quadrant diodes (black) and the single diode (blue) technique. The single diode technique was less susceptible to non-common mode environmental noise at low frequencies (less than 5Hz). This was to be expected, since when all quadrants of the diodes were summed the position of the beam on the active area was unimportant as long as all the light was collected. However, the single diode technique reaches a white noise floor at around 2 Hz, and above 5 Hz the quadrant diode exhibits better common mode noise rejection. This was a result of disturbance from the segmented mirror. In fact, the quadrant diode technique showed the same high-frequency noise performance if the mirror was turned on during the measurement (the red curve) and correlation analysis showed it not simply the result of a gain mismatch between the two diodes. Minimising this noise

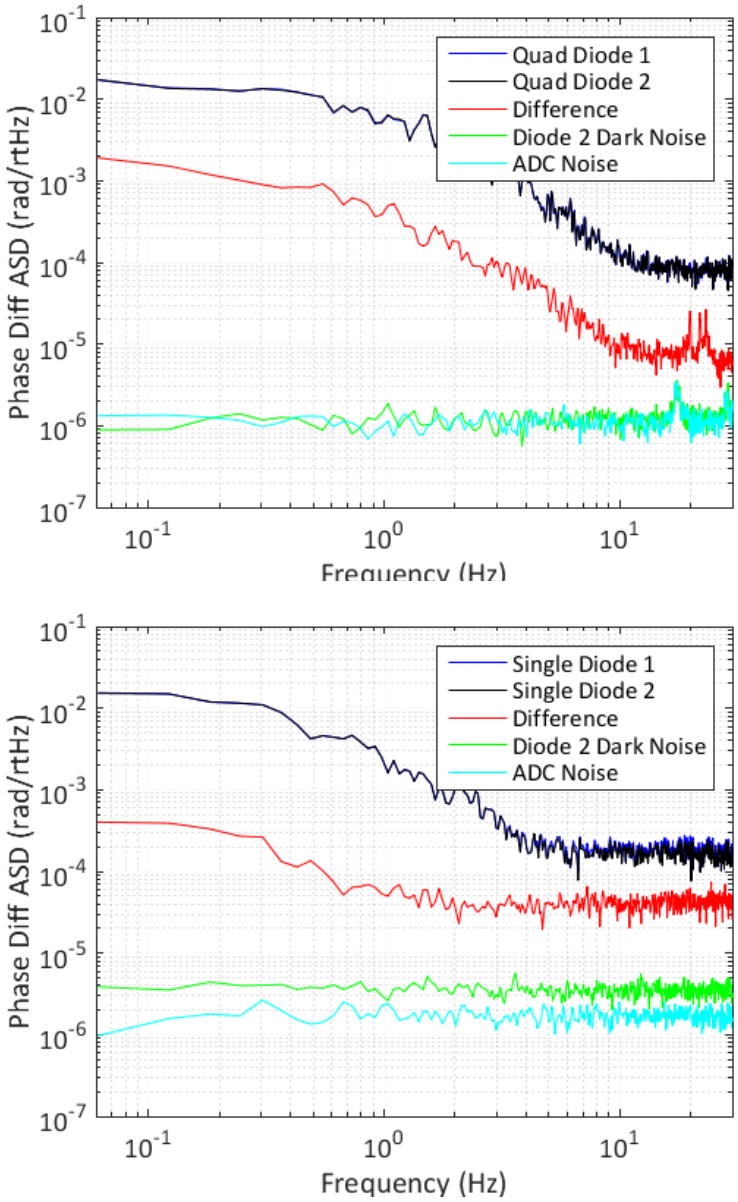

Figure 5.16: Noise contributions for quadrant (top) and single (bottom) diodes. The blue and black curves are the noise performance of live system for two perpendicular diodes, the red curve is the difference between them, the green curve is generated by blocking light to the diodes, and the cyan curve removes ADC noise by terminating the ADCs.

will require some focus in any future set-up, and could be accomplished by better isolation of the mirror or use of a liquid-crystal SLM. (It is not clear why the noise on the segmented diode was lower with the SLM on below 1Hz.)

While sufficient for the proof-of-principle demonstration presented here, a mechanical SLM would not be the preferred choice for a real-world application of this technique for several reasons, including mechanical noise, cross coupling between pixels (allegedly zero for the segmented SLM) and its limited response time. A far superior alternative would be to use an optical phased array (OPA), which uses electro-optic phase modulators to achieve modulation rates of many GHz with no mechanical noise. A recently demonstrated OPA

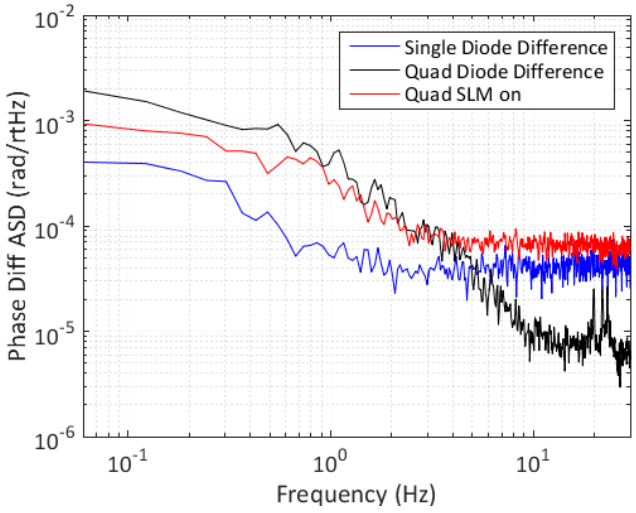

Figure 5.17: A comparison of the performance of the segmented (black) and single (blue) diodes for the live system. The red curve shows the performance of the segmented diode with a minimum modulation on the SLM, consistent with the SLM adding mechanical noise to the set up.

has 61 separate emitters, each of which can be modulated with a pseudo-random noise code at up to 125 MHz14. With a suitable FPGA, such a system would be capable of measuring wavefront phase at 100 fps with a resolution comparable to that achieve in section 5.7.

5.7 Higher Order Wavefront Mapping

The technique presented above is easily extended to develop a more detailed spatial map of the wave front. By modulating each segment of the SLM with a different orthogonal code, the phase and amplitude of many points in the beam can be recovered rather than just the four available from a quadrant photodiode. Figure 5.18 shows a demonstration of this for a 3 x 5 grid of 15 orthogonal codes. A 1 Hz oscillation applied by the fast steering mirror is clearly visible.

5.8 Correcting a Wavefront

One of the promises of the new technique is to not only detect the wavefront but correct it. That was it should be possible to use the SLM to first determine the shape a wavefront and then make small corrections in near real-time. This capability might be used to correct for flaws it has found or to reinforce a particular cavity mode. This is explored in detail in the second experiment, see chapter 6, which looks at the transmission through a cavity and then makes minor adjustments to the wavefront to optimise this. An indicative demonstration was developed during this experiment however.

The SLM was used to apply DC phase offsets to different parts of the beam along with the orthogonal codes, which can then be recovered in the single diode measurement. Figure 5.19 shows how a DC offset and Hadamard codes could be implanted onto the the surface of the SLM so that both act together to modify the shape of the wavefront. The Hadamard

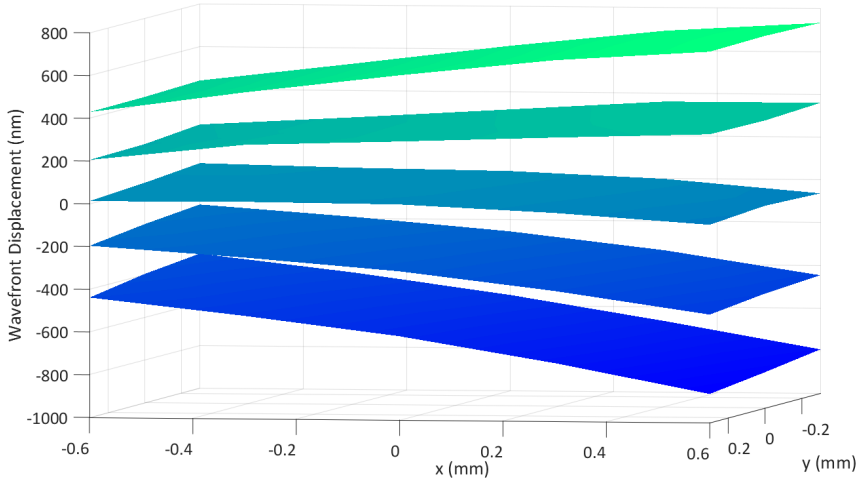

Figure 5.18: A sequence of measurements of wavefront phase over 15 elements of the SLM arranged in a 5 by 3 matrix. The fast steering mirror was used to apply a 1 Hz tilt oscillation to the beam, and the five planes show the wavefront phase as it changes over about half an oscillation. The planes have been separated vertically for clarity.

codes changed much more rapidly than the offset and disappeared during the demodulation process allowing resolution of the shape defined by the DC offset, in this case the numeral '2'. If these offsets were also cycled, a countdown became visible in the wavefront figure 5.20. The entire image was captured by the single photodiode.

Figure 5.20 shows sequential examples of this, with patterns imprinted on a 9 by 7 grid using 63 orthogonal codes (the maximum or small FPGA could process). Within this a central area of 15 pixel (5 by 3) was used to create a countdown by cycling through the numbers from 5 to 1. The orthogonal codes were applied across the 63 pixel grid regardless of the pattern currently resident in the middle 15 pixel grid. Demodulation removed the orthogonal codes but left the numbers intact. These created phase variations across the wavefront and were recovered at near real time in post processing. This illustrates the potential of this technique for the detection and correction of higher-order spatial modes in optical interferometers.

5.9 Conclusion

Digitally enhanced spatial heterodyne interferometry is a new technique which transfers the identification of wavefront information from analogue analysis tools to digital processing. It senses the wavefront phase using a single diode, a spatial light modulator and digital interferometry. It provides an alternative way to measure wavefront tip and tilt with comparable accuracy to conventional approaches based on a quadrant photodiode. The approach is immune to various sources of low-frequency noise that affect quadrant diode techniques because all measurements share a common analogue signal chain. In the implementation presented here, the noise performance above 5 Hz was limited by noise from the segmented mirror, but this could be avoided using a non-mechanical SLM or an optical phased array. This technique is likely to find application in alignment systems for satellite interferometers.

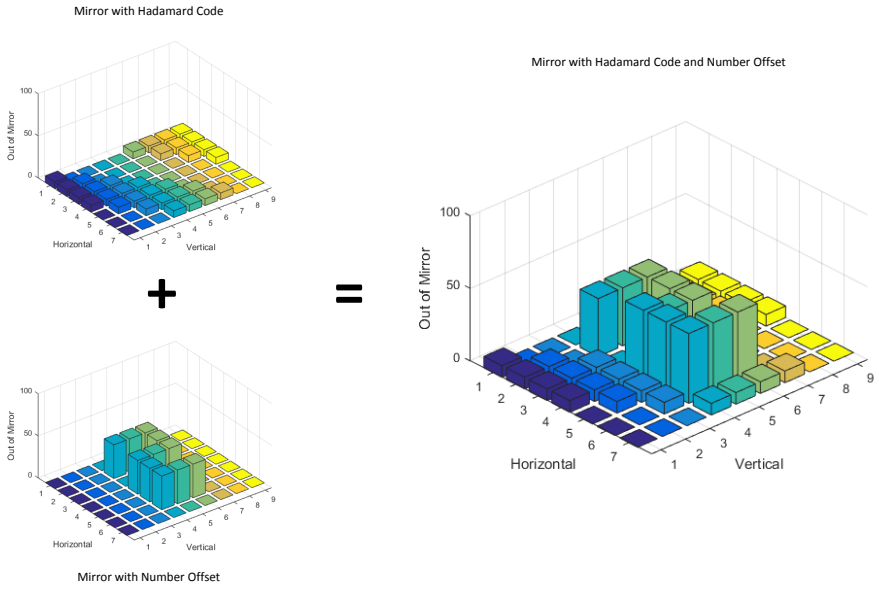

Figure 5.19: An illustration of how a DC offset and a Hadamard code could be added on the surface of the SLM so that both act together to change the shape of the wavefront. The Hadamard code cycles many times for each DC offset and disappears during the demodulation process resolving the shape defined by the DC offset, in this case a 5.

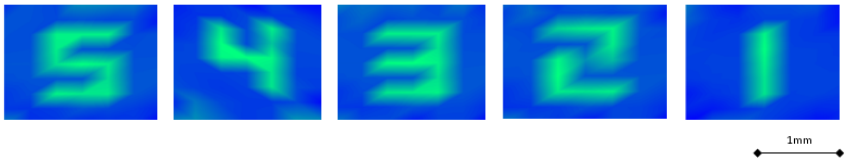

Figure 5.20: Demonstration of higher-order spatial phase manipulation. The blue rectangle shows the individual phases across a 9 by 7 matrix of elements each of which has been spatially tagged with an orthogonal code. The numbers are created by adding a fixed offset to the code modulation on some of the segments of the SLM.

Furthermore, we have shown that the technique can measure phase at many points and so sense and manipulate higher-order spatial modes, limited only by the resolution of the SLM and the capabilities of the digital signal processing hardware. Potential applications include isolating different spatial modes in an interferometer or continuously maximizing input coupling to an optical cavity.

Controlling Wavefronts in Real Time

6.1 An Overview of the Experiment

The goal of the experiment reported in this chapter was to maximise the amount of light transmitted through a cavity by reshaping the laser wavefront so that its mode matches the mode of the cavity as closely as possible. The goal of the second experiment, reported in section 6.5, was to harness these wavefront shaping techniques to select a specific, higher order mode (the TEM_{10} mode) and have this preferentially transmitted through the cavity.

The work has been reported in a draft paper 'Automatic alignment of a Fabry-Pérot cavity using two spatial light modulators, a single photodiode and digitally enhanced heterodyne interferometry', submitted for publication and is currently undergoing peer review. I designed, set-up and ran the experiment on my own. This included building the software and resolving the issues associated with running two spatial light modulators from the same PC with Boston Micromachines. Dr Paul Altin, helped me align the cavity for the first time and provided advice throughout as requested. Professor Daniel Shaddock reviewed my progress and suggested that I run the sensing routine in LabVIEW asynchronously, which saved a lot of coordination work.

There have been previous attempts to optimise transmission through a cavity automatically but these have been focused on physically moving optical components. To date approaches have relied on the differential wavefront sensing outlined in subsection 2.2.1 which uses one or more multi-element photodiodes to measure the varying phase difference between two beams and derive an error signal to improve alignment, see Ref [26]. In 2005 a group at ANU developed a system which used galvanometer actuated mirrors driven by an error signal to control the tilt and offset of the signal beam and optimised them to maximise cavity transmission. The optical arrangement used a Gouy phase invariant telescope to decouple these actuators as far as practicable so that tilt and offset could be controlled separately Ref [101]. In 2010 another group, based in Colorado, developed a simplified actuator with a flat amplitude and phase response to frequencies above 200 kHz, which stabilised the lock of a Fabry-Pérot to a continuous wave laser, see Ref [102]. The cavity was locked using PDH locking.

This experiment includes not only tilt and offset but also waist size and position and, to a degree, beam shape. At first sight it is quite straightforward. The essential optical arrangement is shown in figure 6.1. A laser was directed onto two spatial light modulators and thence along the optical axis of a cavity from which the transmitted power was measured. The surfaces of the SLMs were dithered using predetermined shapes, variations in

transmitted power measured and an algorithm applied to determine the direction in which each movement increased transmission. This information was used to develop an error signal which was multiplied by a gain and applied to make small, automated adjustments to the surfaces of the SLMs and so reshape the wavefront accordingly. The transmitted power was incrementally increased until the maximum value was reached. The individual shapes were dithered with orthogonal codes using the techniques demonstrated in chapter 5 so that their individual contributions could be isolated and evaluated simultaneously.

The first experiment focusses on adjusting beam tilt and offset and waist size and position. This is accomplished by dithering tilts and curves in both the horizontal and vertical planes (see figure 6.8), a total of four degrees of freedom for each mirror. Tilting the surfaces of the SLMs as planes affects the tilt and offset of the beam while curving them, either positively or negatively, changes the incident beam's waist size and position. More complex patterns are introduced later in this chapter in section 6.5.

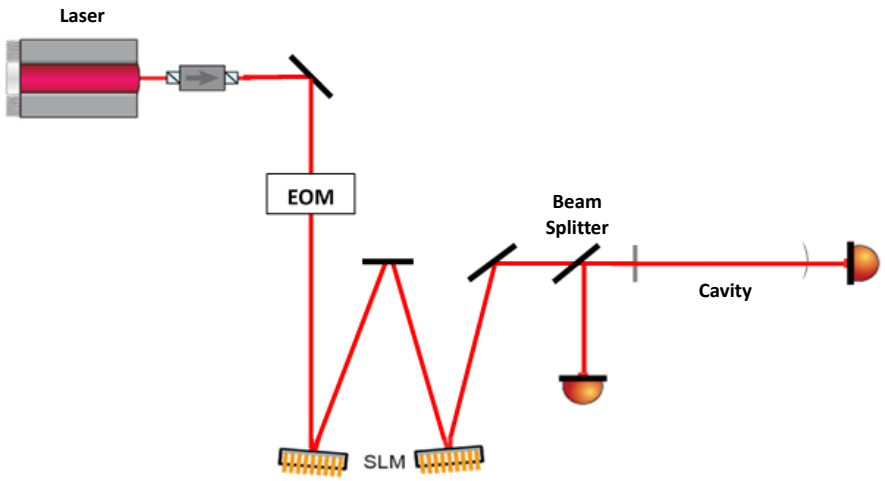

Figure 6.1: A simplified schematic of the essential experimental optical arrangement.

Section4.3 identified some of the factors that had to be optimised to maximise cavity transmission. The first two were physical attributes of the cavity - the quality of its construction and the degree to which it was impedance matched - and these were accepted as given for the cavity used in this experiment. The cavity was of high quality and nearly impedance matched (see its detailed specifications in figure 6.4).

Factor three - the continual match of laser frequency to cavity resonance - was accomplished using PDH locking although future work might assign this work to the SLMs as well. As the cavity was fixed, the PDH error signal was used to feed back to the laser piezo.

The last five - tilt, offset, waist size and waist position and beam quality (the degree to which the beam has a Gaussian profile) - are all qualities of the laser wavefront and were controlled using the two SLMs.

6.2 The Experiment in Detail

6.2.1 Optical Set Up

The detailed optical arrangement is shown in figure 6.2. Figure 6.3 shows a couple of views of the optical bench.

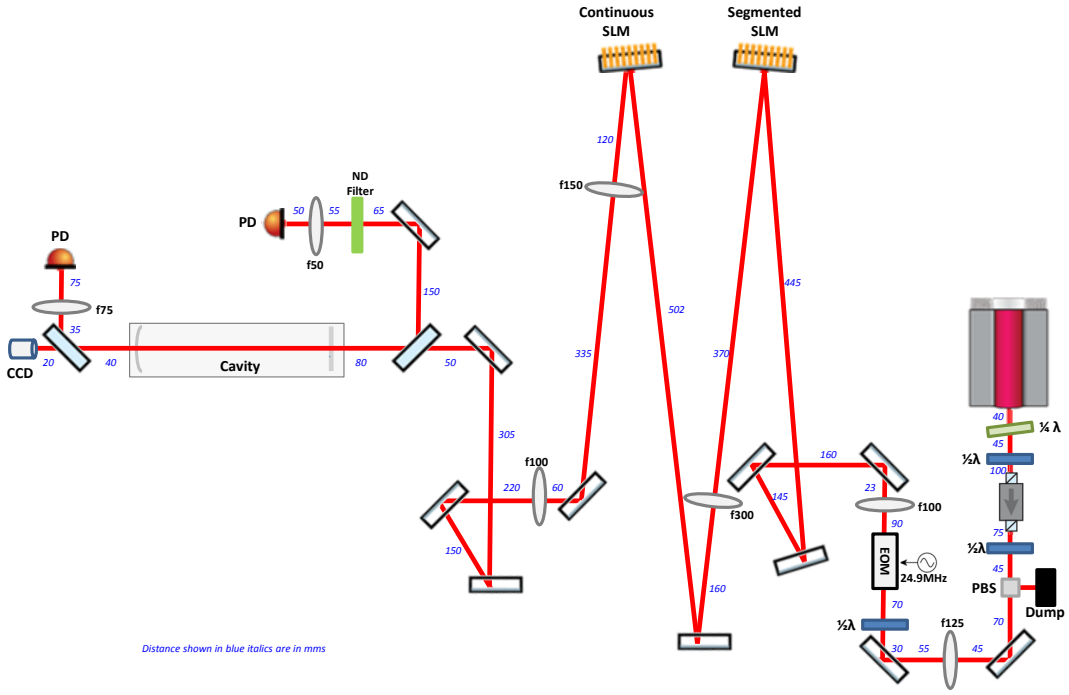

Figure 6.2: A detailed schematic of the optical layout. The distances between components were selected to maximise the distance between the SLMs to maximise the decoupling of tilt and offset while ensuring the beam had the right dimensions to cover the active areas of both SLMs and arrive at the cavity with a waist of 0.55 cm. Lenses with low f numbers were avoided to minimise distortions. The distances between optical components are shown by the blue number in millimetres.

A 1064 nm laser from an Innolight Prometheus laser was directed through a quarter wave plate and a half wave plate to ensure that the emerging beam was plane polarised. From here, it was passed through an isolator, which prevented light from the experiment reflecting back into the laser, and a second half wave plate which acted as a power control by allowing more or less of the laser to be dumped into an adjacent beam dump. The power recorded after the isolator was 670 mW, far to high for this experiment so the majority was dumped and about 82 mW preserved to make the optical measurements. Thereafter it passed through a series of lenses into an electro-optic modulator (EOM) which provided the phase sidebands necessary to create the PDH locking signal. The EOM was modulated at 24.9 MHz.

After that a telescope enlarged the beam so that it covered the active area of the first SLM, the segmented SLM, a square of 3.6 mm, and thence through another lens and off a mirror onto the deformable SLM which had an active area of 4.4 mm^2. The large separation between the SLMs and use of an intervening lens made it possible to decouple changes to beam tilt and offset sufficiently to allow management of both attributes while maintaining

control over waist size and position, see 4.3.3. (The large beam size made achieving complete decoupling impractical within the confines of the bench space available.)

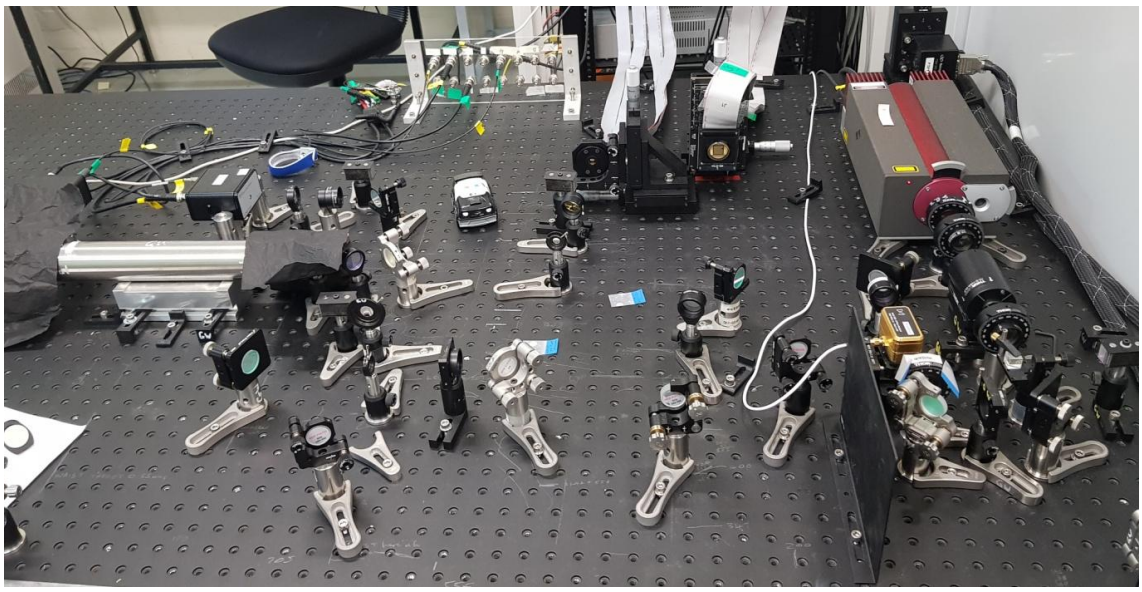

Figure 6.3: A photograph of the optical bench used to conduct the second and third experiments. The entrance and exit from the cavity are covered with flexible, black sheeting to reduce background light and enhance the images on the CCD camera.

From the second SLM the beam was passed through another telescope to create a waist of 0.55 mm on the plane mirror of the cavity. Just in front of the cavity a beam splitter allowed some of the light reflected from the cavity to be directed onto a high speed photodiode (a Newport 1811) as an input to the PDH locking process. Beyond the cavity the beam was split again with 80 percent passing into another high speed photodiode, also an 1811, and the remaining 20 percent being directed to a CCD camera so that the transmitted beam could be examined visually.

The cavity chosen was a plano-concave, fixed spacer cavity with characteristics as outlined in figure 6.4. The reflectivities of the two mirrors were both close to 99% so the cavity was low loss and comprised mirrors of close to equal reflectivity. This meant that it was nearly impedance matched and that it should be possible for nearly all the power entering the cavity to emerge as transmitted light when appropriately mode matched.

Attribute	Value
Length	30 cm
Cavity Type	Plano-concave
Cavity Structure	Aluminium, rigid spacer (No Piezo mirror control)
Waist	0.55 mm diameter
Reflectivity of plano mirror (r_1)	~99%
Reflectivity of curved mirror (r_2)	~99%
Radius of curvature of curved mirror	~3 m
Finesse	~314

Figure 6.4: Table 1: The specifications of the cavity used in the experiment.

Spatial Light Modulators

The spatial light modulators used were from Boston Micromachines. The first was the same used in the experiment described in Chapter 4 and had a segmented surface. Briefly the segmented SLM comprised a 12 by 12 array of gold coated mirrors each of which could be advanced by up to 1.5 microns. (The mirrors in the corners were fixed making the actual number of active elements 140.) Elements were controlled by sending an array of 140, 14-bit numbers to the SLM controller which, once received initiated a change of state. A 14-bit number allows 16384 values suggesting minimum actuator steps for the segmented and continuous SLMs of 9.16 nm and 21.4 nm respectively.

This change took 31 micro-seconds although tests showed that there was also a settling time up to 40 micro-seconds which depended on the magnitude of the movement required.

The second SLM was similar but its actuators were covered with a continuous gold surface which removed phase discontinuities between neighbouring pixels at different offsets. It had a full stroke length of 3.5 microns and had an active area of 4.4 mm. The respective attributes of the two SLMs are shown in figure 6.5.

Attribute	Segmented SLM	Continuous SLM
Boston Micromachines Designation	Multi-DM 1.5 µm SLM	Multi-DM 3.5 µm CDM
Maximum Stroke	1.5 µm	3.5 µm
Surface	Segmented	Continuous
Actuator Pitch	300 µm	400 µm
Active Aperture	3.6 mm x 3.6 mm	4.4 mm x 4.4 mm
Step Size	<1nm	< 1nm
Surface Covering	Gold	Gold
Surface Quality	< 30 nm RMS	< 30 nm RMS
Inter-actuator coupling	0%	20-40%
Minimum step size (14-bit granularity)	9.16nm	21.4nm
Power limits (Unpublished)	100W/cm^2	20W/cm^2
Frame Rate	8kHz with 34kHz burst	8kHz with 34kHz burst
Tested Continuous Max Frame Rate	10.374Khz	10.396kHz

Figure 6.5: Table 2: A comparison of the performance specifications of the SLMs

The SLMs controllers were matched to each SLM individually. This ensured that the responses to arrays sent to the controllers were within the tolerances promised by the manufacturer. They could be driven from a PC using either Matlab or LabVIEW. Two types of controllers were available - Multi Drivers and X-Drivers. The Multi Driver Controller linked to the server PC using a USB connection. It was considerably cheaper as a control option but introduced a degree of uncertainty about when new instructions for the SLM to change shape would actually be sent. This was because the USB connection continually competed with other resources being controlled by the operating system (Windows 10). The X-Driver, in comparison, connected to the PC via a PCIe card meaning its response was far more deterministic and dependent only upon the accuracy of the PC's clock (2.7 GHz).

In the end these choices were academic. According to Boston Micromachines, we were the first to use two SLMs in the same experiment and, in the process, unearthed significant problems with the vendor's driver software. As initially delivered the drivers were unable to distinguish between the two SLMs in the system. A command directed to one SLM affected both simultaneously. This took nearly 6 months to resolve, working closely with

Boston Micromachines. After considerable work we managed to get Matlab drivers which allowed independent operation for the Multi Driver controllers only.

This created a new set of challenges unique to our implementation. Our operational software of choice was LabVIEW as we needed to measure the power of the beam transmitted from the cavity in real time. This was used to create the error signal that would reconfigure the SLMs and maximise transmission. As both SLMs had to be controlled together and this was only possible with Matlab, our only option was to create Matlab instances within our LabVIEW code. The final arrangement, and the only arrangement that will currently allow the independent operation of two Boston Micromachines SLMs from the same program, was to run Matlab control scripts within LabVIEW and use these to manage to Multi Diver Controllers through USB connections. Having solved the problem of when a pulse had actually arrived and when the mirror had settled in the first experiment, see 5.4.1, we were able to use this logic to ensure that we made measurements only during periods when the SLMs had settled into stable states. With a margin for error this allowed us to use about 72% of operational time for measurement, like that achieved in the first experiment.

The Management of Beam Size and Quality

In this proof of concept experiment the SLMs were asked only to make relatively minor adjustments to beam shapes and quality, although these had large effects. Consequently the optical set up was aligned as carefully as possible within the limitations of space and the requirement to decouple actuation of beam tilt and offset and waist size and position.

The management of beam size was important. First it had to be expanded to cover the active areas of both SLMs and then reduced it to a waist of 0.55 mm at the cavity. Ideally this had to be accomplished without using extremely short focal length lenses because while these shorten the optical path they can cause discernible spherical aberration in beam quality, Ref [103].

First, great care was taken to create a polarised beam with a circular cross section from the laser and maintain beam quality by careful alignment of optical components.

The Segmented SLM was the first non-standard component in the optical chain. It had a major impact on wavefront quality. This can be seen in Figure (6.6) which shows the shape of the beam emerging from each SLM as measured by the Thorlabs Scanning-Slit Optical beam scanner. The segmented SLM was placed farthest from the cavity to try to minimise these effects as shown in figure 6.7. The segmented SLM acts as a diffraction grating and the high spatial frequencies corresponding to the grating structure travel at large angles to the optical axis and so are not collected by the optical system as the laser beam progresses.

The Need for Two SLMs

Subsection 4.3.3 explained that full control over tilt and offset, and waist size and position was only possible using two SLMs. Tilt and offset are orthogonal. The limitations of space available and the need to manage beam size meant that, while these were not decoupled completely, they were decoupled sufficiently to allow preferential adjustment of each independently using different mirror configurations. In summary:

- The segmented SLM was farthest from the cavity and affected beam offset more than beam tilt;

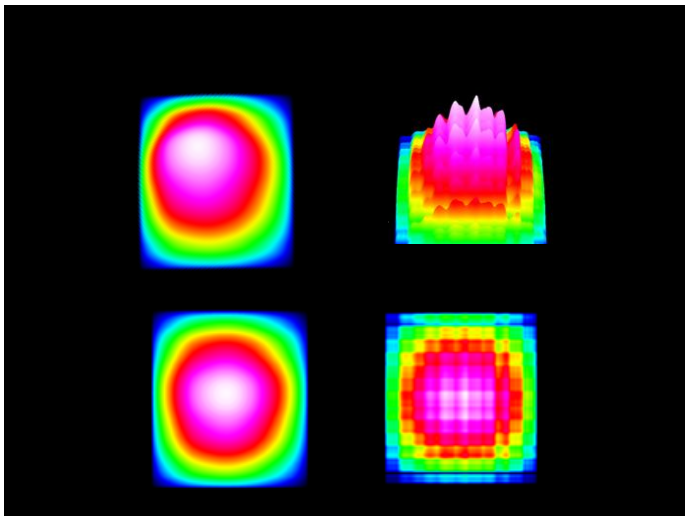

Figure 6.6: The effect of the segmented SLM on beam quality. The images above were taken using a ThorsLabs Scanning-Slit Optical Beam Profiler and accompanying software. The images on the left show the beam profile before it strikes the segmented SLM while those on the right show the beam profile after reflection from the SLM. The effect of the segments is clearly visible as a disruption to the beam's Gaussian shape. The top and bottom rows show the same scans. The top row shows the scans viewed from the side and above while the bottom row shows scans viewed from directly above.

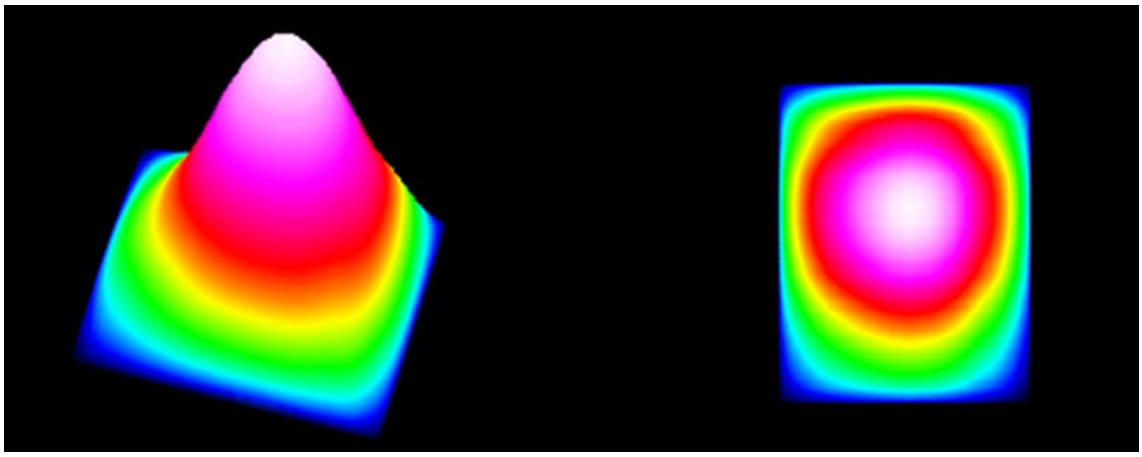

Figure 6.7: The quality of the beam after the second SLM. By ensuring the segmented SLM as far from the cavity most of the aberrations shown in Figure 6.6 have been minimised.

- The continuous SLM was closer to the cavity and affected beam tilt more than beam offset

For tilt and offset, the SLMs need only act as flat, adjustable mirrors which were tilted in the horizontal and vertical planes to achieve the small adjustments required.

As subsection 4.3.3 shows, the relationship between waist size and waist position is more complex. Here the SLMs were required to act as mirrors with variable curvature, refocusing the beam in the near and far fields. The algorithm was required to move the position and size of the waist simultaneously until it was optimised at a diameter of 0.55 mm on the

plane mirror that formed the front of the cavity.

Beam Shaping and Modulation Depth

During the experiment the power transmitted through the cavity was measured as each SLM was dithered to make small adjustments to the beam's tilt and offset, and waist size and position. From this an error signal was produced and used to adjust each SLM. The process was repeated until the transmitted power reached a maximum value and additional changes had no further effect.

Modulating Tilt and Offset

To achieve the changes to tilt and offset the mirrors were adjusted as plane mirrors in the horizontal and vertical directions, see the left hand column in figure 6.8.

So if, for example, the first SLM was oscillated in the horizontal plane by a small amount (determined by the modulation depth) and the power increased every time the left hand side of the SLM advanced and decreased every time the right hand side advanced, the algorithm would conclude that the mirror should be angled slightly more with the left hand side forward. It would multiply the increase in power by a gain (adjustable in the program) and reset the mirror accordingly. Each time an attribute was changed, the transmitted power was changed. The algorithm used the change in transmitted power, to determine whether to add more or less of this change (multiplied by a gain) to the shape of the mirror, always moving the mirror in a direction that increased the transmitted power.

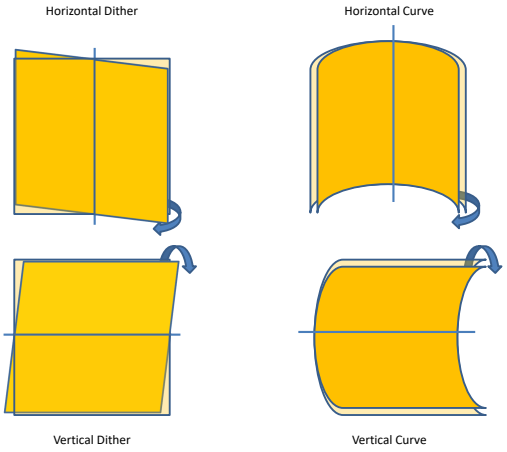

Figure 6.8: The independent dithers applied to each SLM allowing the four degrees of freedom on each mirror to modulate beam tilt, beam offset and waist size and position.

The actual size of the dither was taken by applying the modulation depth, set in the LabVIEW system, across the range of movement of the mirror. Figure (6.9) shows the maximum tilts achievable for an SLM assuming one edge is left unmoved while the opposing edge is extended to the maximum extreme. As figure 6.9 shows, the segmented mirror can be tilted by up to $4.5 \cdot 10^{-4}$ radians but in practice a much smaller modulation depth was used. The modulation depth used was 0.2% which moved the extreme edges of the mirrors forward and back by 0.03 microns and 0.07 microns for the segmented and continuous SLMs respectively. This provided a tilt modulations of $1.6 \cdot 10^{-6}$ radians and $3.04 \cdot 10^{-6}$ radians, again, respectively.

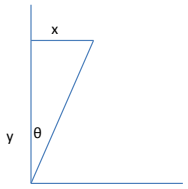

$\Theta = \tan(x/y)$

Calculation of Maximum Tilts		
Variable	Segmented Mirror (m)	Continuous Mirror (m)
x	1.50E-06	3.50E-06
y	3.33E-03	4.03E-03
Θ rad	4.50E-04	8.68E-04
Θ deg	2.58E-02	4.97E-02

Figure 6.9: The maximum tilts available across the two SLMs.

The active area of the segmented SLM is 3.6 mm^2. The tilt occurs between the middle of the first row of micro-mirrors and the middle of the last row of micro-mirrors making the effective length of the mechanical lever 3.3 mm not 3.6 mm. Similarly the effective length for the continuous SLM was 4.033 mm not 4.4 mm. (Note larger tilts are possible by employing saw tooth SLM surfaces but these introduce additional aberrations in the beam and were unnecessary.)

To create changes in waist size and position, the curvatures of the mirrors were increased or decreased in the horizontal and vertical directions, the right hand images in figure 6.8. Figure (6.10) shows the geometry for curving the surface of the SLMs, while figure (6.11) shows the focal lengths that can be achieved for each SLM at different actuation percentages.

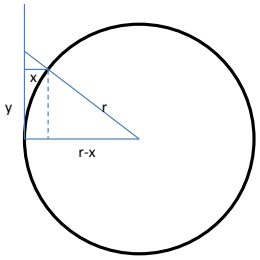

$r^2=y^2+(r-x)^2 = x^2+y^2+r^2-2rx$

$r = (x^2+y^2)/2x$

Calculation of Maximum Focal Lengths		
Variable	Segmented Mirror (m)	Continuous Mirror (m)
x	1.50E-06	3.50E-06
y	1.65E-03	2.02E-03
r	9.08E-01	5.81E-01
Max Focal Length	4.54E-01	2.90E-1

Figure 6.10: The maximum curvatures available on the SLMs turn them into concave mirrors with focal lengths of 454 mm for the segmented SLM and 290 mm for the continuous SLM.

Initially the attributes were adjusted sequentially. Once this was working the attributes were dithered simultaneously (using orthogonal codes), a measurement made and adjustments applied to all 8 degrees of freedom before the next measurement was taken. In this way the transmitted power quickly rose to a maximum where it remained, adjusting to compensate for any small variations (expansions, vibrations etc.) in the optical circumstances as necessary.

6.2.2 Electronic Set Up

The electronic support for the experiment comprised three subsystems, namely:

- the PDH locking and laser beam scanning system;
- the sensing and LabVIEW delivery system.

Segmented SLM					Continuous SLM				
% Curve	x	y	r (m)	f(mm)	% Curve	x	y	r (m)	f(mm)
100	1.50E-06	1.65E-03	0.91	454	100	3.50E-06	2.02E-03	0.58	290
80	1.20E-06	1.65E-03	1.13	567	80	2.80E-06	2.02E-03	0.73	363
60	9.00E-07	1.65E-03	1.51	756	60	2.10E-06	2.02E-03	0.97	484
40	6.00E-07	1.65E-03	2.27	1134	40	1.40E-06	2.02E-03	1.45	726
20	3.00E-07	1.65E-03	4.54	2269	20	7.00E-07	2.02E-03	2.90	1452
15	2.25E-07	1.65E-03	6.05	3025	15	5.25E-07	2.02E-03	3.87	1936
10	1.50E-07	1.65E-03	9.08	4538	10	3.50E-07	2.02E-03	5.81	2904
9	1.35E-07	1.65E-03	10.08	5042	9	3.15E-07	2.02E-03	6.45	3227
8	1.20E-07	1.65E-03	11.34	5672	8	2.80E-07	2.02E-03	7.26	3631
7	1.05E-07	1.65E-03	12.96	6482	7	2.45E-07	2.02E-03	8.30	4149
6	9.00E-08	1.65E-03	15.13	7563	6	2.10E-07	2.02E-03	9.68	4841
5	7.50E-08	1.65E-03	18.15	9075	5	1.75E-07	2.02E-03	11.62	5809
4	6.00E-08	1.65E-03	22.69	11344	4	1.40E-07	2.02E-03	14.52	7261
3	4.50E-08	1.65E-03	30.25	15125	3	1.05E-07	2.02E-03	19.36	9682
2	3.00E-08	1.65E-03	45.38	22688	2	7.00E-08	2.02E-03	29.04	14522
1	1.50E-08	1.65E-03	90.75	45375	1	3.50E-08	2.02E-03	58.09	29045

Figure 6.11: The focal lengths of the SLMS for different percentage curvatures.

The PDH Locking and Laser Beam Scanning System

The PDH locking system was the conventional arrangement and is summarised in figure 6.12. A signal of 24.9M Hz was passed from a signal generator to the EOM to put the required phase sidebands onto the laser beam striking the SLMs. This was mixed with the power reflected from the cavity and detected by a single diode, low pass filtered (DC to 1.9 MHz) and passed to a PID controller before being amplified and low pass filtered again to create the PDH error signal. This was fed to the piezo input on the laser and adjusted the frequency of the laser to match minor changes in cavity length as described in chapter 4.

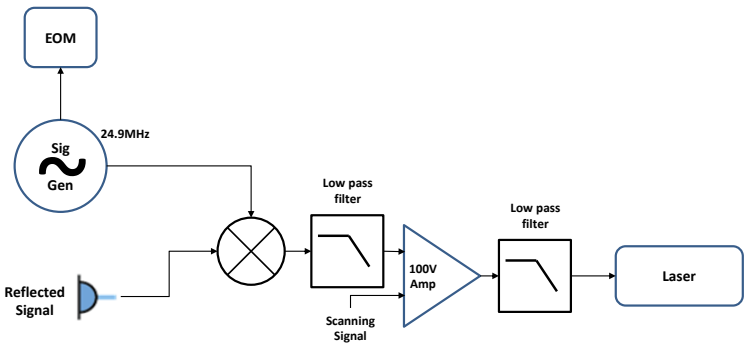

Figure 6.12: The individual electronic systems used to support the experiment.

Locking the laser to the cavity length involved:

- Scanning across an FSR to find the TEM_{00} mode. The FSR for the cavity at 1064 nm was around 500 MHz while the laser imposed a maximum scanning width of

about 200 MHz. Thus only 40% of an FSR was visible at any one time. This window was displayed as the power output by the photodiode receiving the reflected power and moved along the FSR by slowly changing the laser's temperature. The TEM_{00} was identified by examining peaks in power output and using the image collected on the CCD camera to confirm to which mode the peak referred.

- Once the TEM_{00} mode was identified the scan was narrowed and the peak centred in the middle of the scanning pulse.

- Here the scanning signal generator was turned off and the 26.4MHz generator turned on to create a PDH locking signal. Further small adjustments were made to make sure this signal was at it's maximum.

- The PID controller was engaged to lock the laser to the error signal.

The laser scanning system is shown is figure 6.13. It allowed the frequency of the laser to be scanned during PDH lock acquisition and then for maintenance of the lock. The laser controller facilitated direct changes to laser temperature and injection current and was used during the scanning process at the start of the PDH lock. The PDH signal fed directly into the piezo in the laser and allowed the laser frequency to be adjusted to follow changes in cavity length.

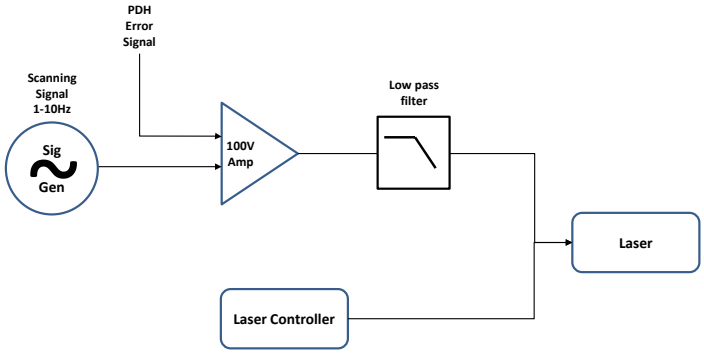

Figure 6.13: A schematic of the laser scanning system.

The Sensing and LabVIEW Delivery System

The sensing and LabVIEW delivery system was quite straightforward and is shown in 6.14. It connected the DC output from the single photodiode which sensed the transmitting power to both the oscilloscope and the LabVIEW chassis. It also connected the AC and DC outputs from the single photodiode capturing the reflected signal and passed these to the oscilloscope as well as passing the AC signal to the PDH locking system. The system also sensed the sync pulse delivered by the continuous SLM. As in the previous experiment, the SLM controllers were driven by Matlab and linked by USB to the PC.

A third system, which used a servo to directly vary the operating temperature of the laser was trialled and proved too noisy to be useful.

Figure 6.15 shows the entire system used to support the experiment.

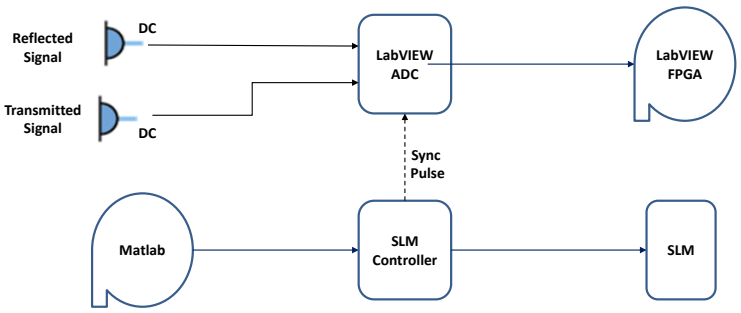

Figure 6.14: A schematic of the sensing and LabVIEW delivery system.

6.2.3 Data Processing

This experiment used digitally enhanced heterodyne interferometry to interrogate spatially distinct regions of the wavefront so, once the optical arrangements are in place, the focus switched to data processing. Once again, both Matlab and LabVIEW were employed to manage the experiment - Matlab to drive the SLM and LabVIEW to collect and analyse data. However the difficulties with the SLM drivers and the requirement to operate both mirrors simultaneously from the same set of error signals meant that the system had to be run from one PC and that, as a consequence, the Matlab code had to be run inside the LabVIEW code. This turns out to be relatively straightforward although debugging can be problematic and it is worth running the code externally first to check for errors.

The architecture of the software environment is shown in figure 6.16. The system sensing and time sensitive signal separation and demodulation was completed on the FPGA while creating the error signal and developing the SLM control signals were managed by the host PC.

Data Flow

The software architecture and basic data flow is shown in figure 6.16. The ADC continually read the output of the photodiode which monitored the power transmitted through the cavity. This monitoring loop was run asynchronously from the rest of the system which simply took the current value of the transmitted power whenever it was ready to create a new error signal. The measurement cycle was triggered by a sync pulse from the continuous SLM (see below) which was detected by the Sync Pulse Detection module. All the same controls that were employed in the initial experiment were used to ensure that it detected a real sync pulse and not a false positive.

Once alerted by the sync pulse that both mirrors had changed state and settled, the Signal Separation and Integration module to read the current value of the transmitted power from the ADC and separated it into its different spatial components by spreading it across up to 32 pairs of channels, demodulating each pair of channels with sine and cosine and the appropriate Hadamard codes and then integrating the resulting signals to obtain the current power for a given spatial element. This data was assembled into an array and passed to the FIFO for transfer to the host.

Once in the host, the Error Signal Creation Module multiplied these power data by nominated gains and used these to create feedback signals for different spatial components. These were transferred to the Matlab Mirror Control Module which translated these error

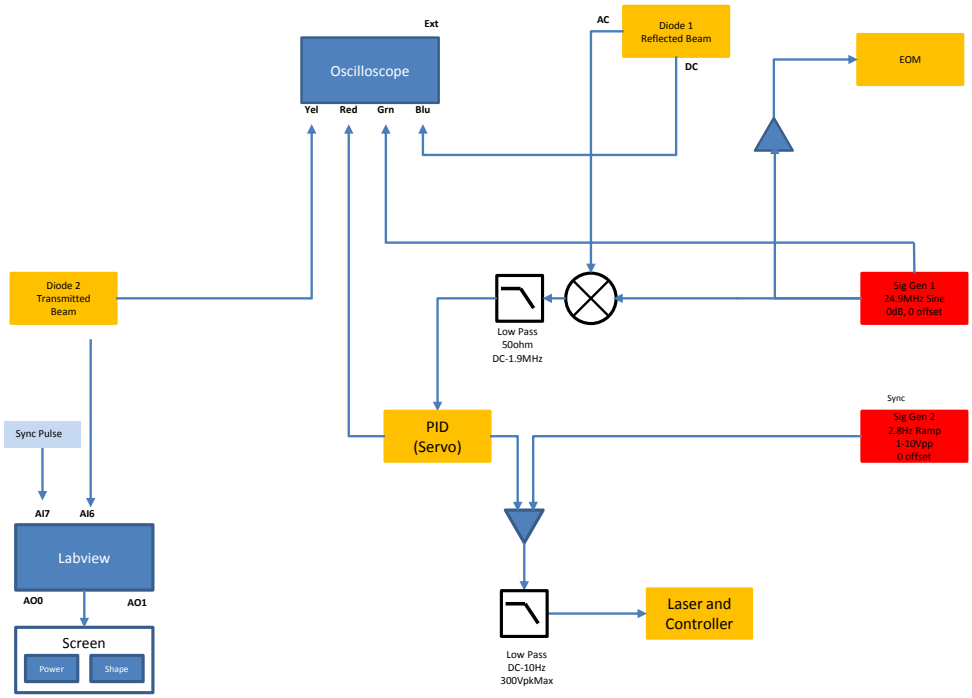

Figure 6.15: The Electronic Set Up for the Second Experiment.

signals into new mirror configurations and passed them to the USB bus for transfer to the mirror controllers.

Synchronising the SLMs

Measurement events were triggered by synchronisation pulses from the two SLMs and both had to have settled into their new configurations before a reliable measurement could be taken. Section 5.4.1 explains the timing difficulties associated with the USB interface, the only option available for managing two SLMs simultaneously, and how safety margins had to be set which reduced reliable sampling time to about 75% of run time. The initial plan had been to wait for sync pulses from both SLMs but this proved unnecessary. The relative speeds of the two mirrors were tested using a program that continually moved their surfaces from flat at the 0% offset to flat at the 100% offset and after one million runs it was clear that the mirrors operate at substantially the same rate and that any timing errors were well within the tolerances that had to be allowed for safety margins. (Specifically one million cycles took just under 191 seconds, an effective operating speed of 10.5kHz, with the continuous mirror being marginally slower and running at 97.7% the speed of the segmented mirror.) Consequently measurements were triggered by identifying the sync pulse from the continuous mirror and applying the safety margins identified in Chapter 5.

The ADC continually read the input from the photodiode which monitored the power transmitted through the cavity. This sensing loop was run asynchronously from the rest of the system which simply took the current value of the transmitted power whenever it was ready to create a new error signal.

LabVIEW The experiment was managed using LabVIEW which used sync pulses to initiate measurements of the power transmitted through the cavity and develop an error

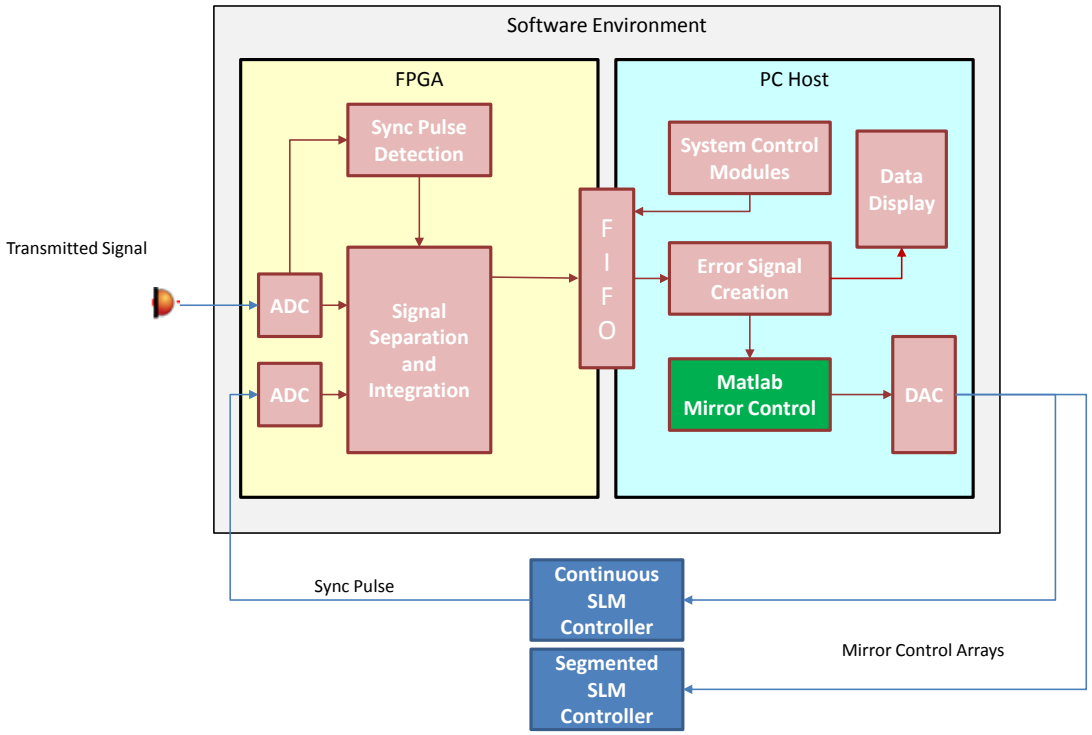

Figure 6.16: The data processing architecture used to support the cavity optimisation experiment.

signal which it then processed and passed back to the SLM controllers. The details of the code are included in Appendix 8.4 but, once again, it is useful to consider the inputs provided to the LabVIEW system. Figure 6.17 shows the front panel of the instance of LabVIEW running on the PC host. Variables were passed to the FPGA as necessary.

The inputs can be grouped as follows:

1. Stop and Reset - The *Stop* button ceased operation of the both the host and FPGA instances of LabVIEW. The *Reset* button allowed the mirror to be reset to flat at 50% displacement during runs for times when the system ran away from resonance (if, for example, the gains were set too high.)

2. Run Algorithm - The *Run* button turned on the algorithm and began the optimisation process.

3. Turn on individual dithers and display values - These controls allowed different dithers to be turned on and off so that the effect of, for example, just a horizontal dither on the continuous SLM could be explored. They were particularly useful in comparing the performance of individual SLMs to their combined effect.

4. Integration parameters - This group of controls managed the Integration (Sum and Dump) algorithm, in particular, how quickly it started after a sync pulse, how long it ran and how much it was scaled down to avoid data loss during the transfer to the host through the FIFO.

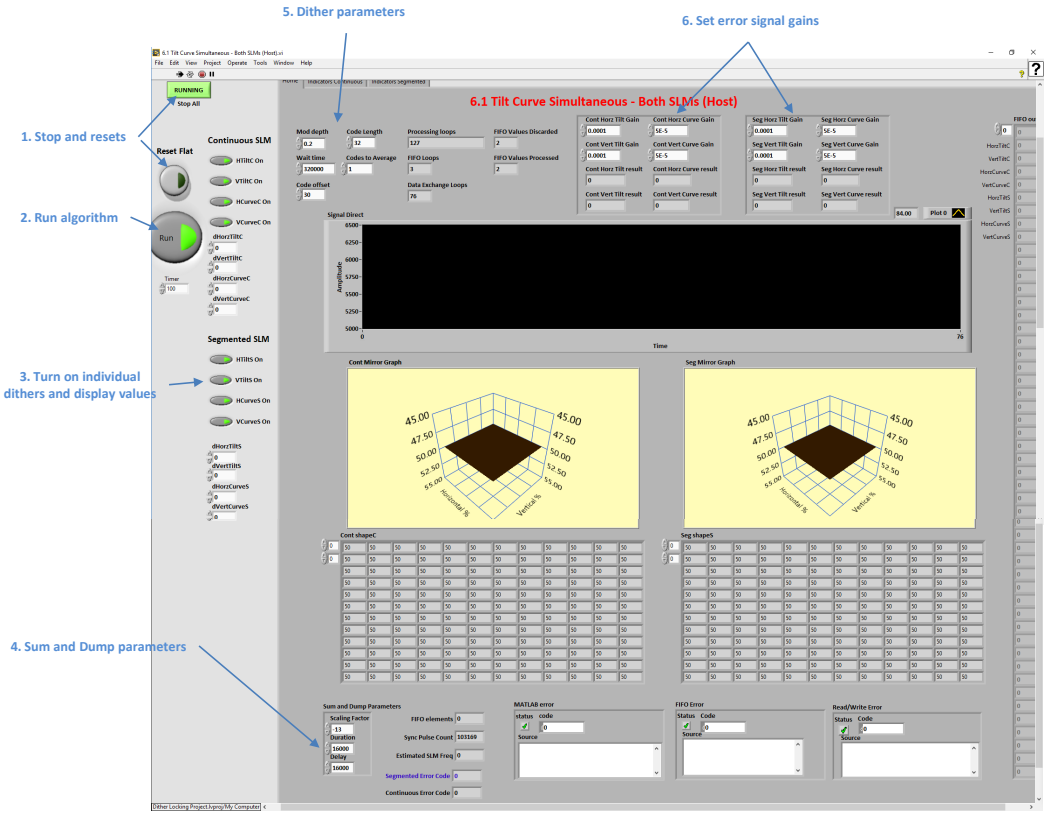

Figure 6.17: The LabVIEW Host Control Panel for the Second Experiment. Data was passed to the FPGA as necessary.

5. Dither parameters - These controls managed the dither setting variables like the length of Hadamard code to be used, the number of code iterations in the cycle and the modulation depth. The system was remarkably responsive to small modulation depth which had the advantage of minimising instantaneous noise on the output.

6. Set feedback signal gains - These controls allowed individual gains to be set for each of the dithers applied to the SLMs.

The front panel also displayed the power transmitted through the cavity in real time as a time series plot and provided information about the current shape of each mirror as both a graphical surface and a matrix of percentages showing the displacement of each SLM actuator. (It is worth noting that the actual movements of the SLMs' surfaces were very small and the scale of the graphical surfaces were selected to highlight changes to these surfaces and made them appear much more dramatic (larger) than they actually were.) It also provided information about the raw channel data received from the FIFO and error messages as necessary. Additional tabs provided feedback about the ongoing status of the individual SLMs.

Matlab

Because of the constraints imposed by the SLMs' drivers, control of the SLMs was managed using Matlab running within LabVIEW. The Matlab code appeared in three modules:

- the *Open Module* opened the USB connection the the SLM controllers and set up

the parameters that would be used to convert the error signals into dither as outline in section 6.2.1.

- the *Operating Module* received the error signals for the eight degrees of freedom, converted these into tilts or curves by multiplying them by specified gains and adding the results to any existing adjustment, before sending them to the SLM controllers. They also refreshed the status information on the front panel of the LabVIEW host.
- the *Close Module* made sure that the SLM controller were reset to flat and closed down properly when the system were terminated.

Details of the Matlab code are included in Appendix 8.4.

6.3 System Operation

Once set up, the operation of the system was straightforward. It involved the following steps.

1. Roughly align the optical components so that the TEM_{00} mode is being transmitted through the cavity.
2. Lock onto the TEM_{00} mode using the PDH locking.
3. Turn on the LabVIEW system and set the initial operational parameters.
4. Push the run button and to start the algorithm and allow it to reshape the SLMs so that they optimised the power transmitted through the cavity for the selected mode (initially the fundamental).

6.4 Results

Figure 6.18 shows a typical result from running the system. The curve shows the power transmitted through the cavity before and during optimisation. The x-axis plots individual samples collected during the run and is a proxy for time (although the uncertainty associated with the USB means that there is not a precise correlation.) The noise on the trace is the effect of the modulation depth of the dither signal which disappears during processing (because of the integration) but remains an instantaneous feature of the process. The first and third peak show the effect of the continuous SLM acting alone while the second shows the effect of the segmented SLM acting alone and the fourth both acting together. The segmented SLM has a smaller effect. Together the mirrors create the greatest effect increasing the power by approximately 39% (from 2800 to 3900) although the starting level is entirely dependent upon the degree to which the optical alignment has been optimised before the system is engaged. The long tail from 8000 sample onwards shows the stability of the system and other executions confirmed that it was able to sustain the optimisation for a couple of hours.

Figure 6.19 shows the CCD image of the transmitted beam before before and after optimisation.

Figure 6.20 shows the shape of the mirrors after the optimisation routine noting, again, that the z-axis is heavily exaggerated to make changes in the surfaces of the SLMs clearly

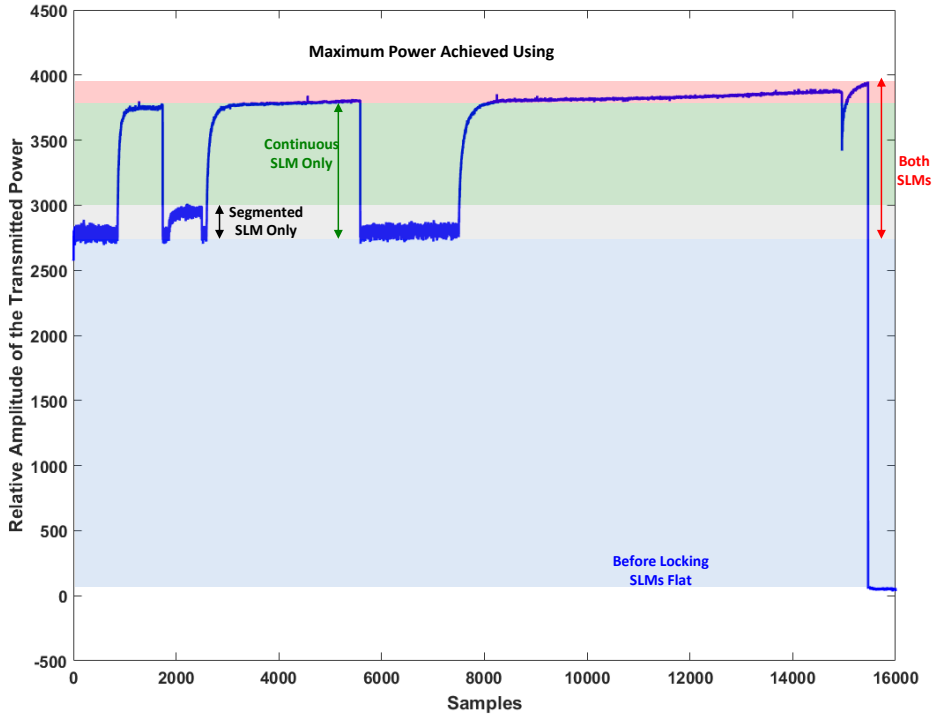

Figure 6.18: The curve shows the power transmitted through the cavity before and during optimisation. The first and third peak show the effect of the continuous SLM acting alone, the second the effect of the segmented SLM acting alone and the fourth both acting together. Samples is a proxy for time although not precise because of the variability of the USB interface. In this case the system ran at approximately 30 samples a second.

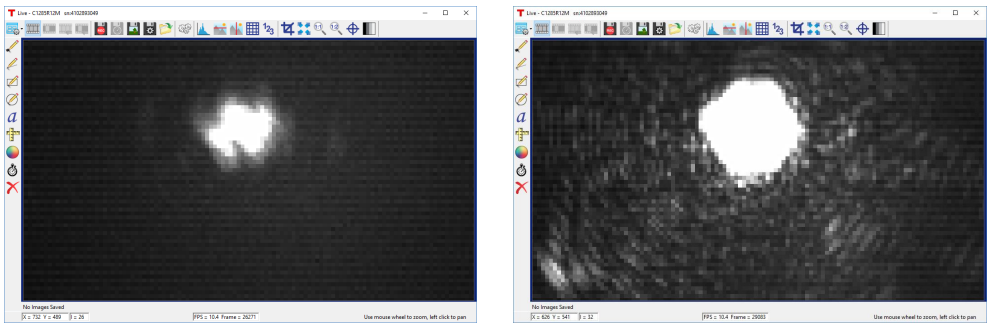

Figure 6.19: The image of the transmitted beam before (left) and after (right) optimisation. While the CCD is saturated, hence the odd shape of the fundamental mode, the plot does illustrate the noticeable increase in power caused by the optimisation process.

visible. (100% is 1.5microns for the segmented mirror and 3.5 microns for the continuous mirror, while the active areas are 3.6 mm square and 4.4 mm square.)

The system has also demonstrated it capacity to recognise and respond to changes in the optical layout. Figure 6.21 shows the system starting in a low, unoptimised state, being optimised and then responding in real time to manual adjustments made to the mount of

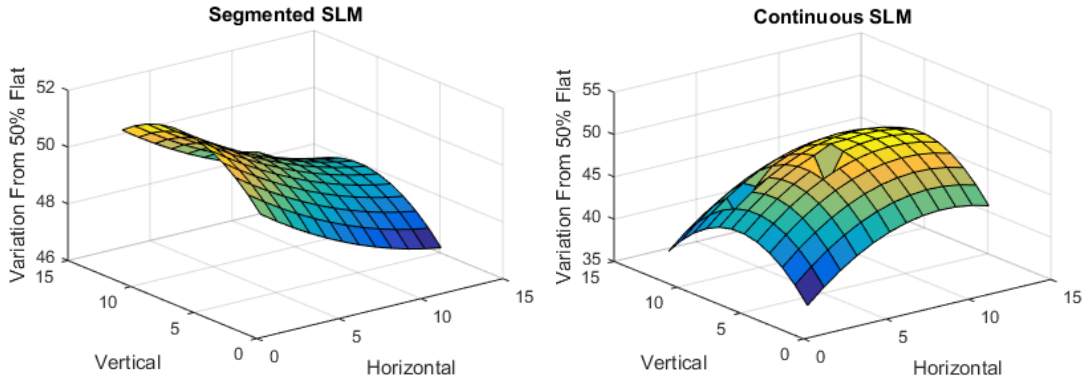

Figure 6.20: The shape of the segmented (left) and continuous (right) SLMs after optimisation. The initial state of the SLMs before optimisation was flat at the 50% position. The plots show how much they have moved from this position as a percentage of the total distance they can move (50% in either direction). For the segmented mirror 50% = 0.75 microns while for the continuous SLM, 50% = 1.75 microns.

the continuous SLM. As the a mount is moved, the beam is no longer optimised along the optical axis and the algorithm responds by adjusting the SLMs to re-optimise it. It also shows the limitations of the system while it relies on an external PHD locking to maintain the match between cavity length and laser frequency. When the manual adjustment is too large, the system loses lock and is unable to recover.

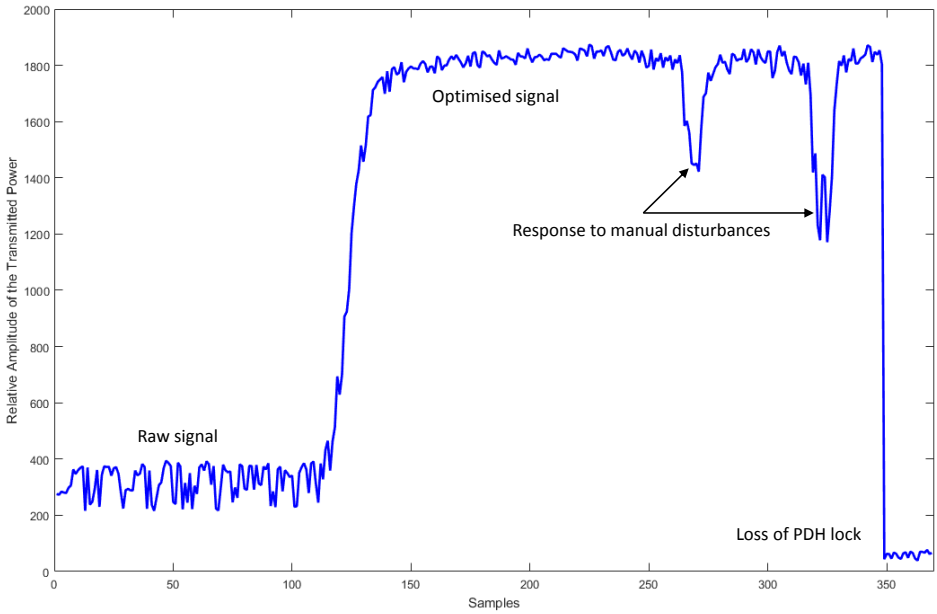

Figure 6.21: This plot show how the system is able to compensate for real time adjustments to the optical path.

Figure 6.22 shows the power in the transmitted and reflected signals before and after an optimisation run. In each case the signal is being scanned across a 200MHZ band around

the TEM_{00} mode (remembering that the cavity has an FSR of 500 MHz but that the the complete range cannot be scanned without the risk of damaging the laser's piezo). The scale of all four signals is the same but their plots moved laterally to align the peaks of the the TEM_{00} modes. The figure shows how optimisation redirects a considerable amount of the transmitted power from the higher order modes into the the TEM_{00} mode making both the transmitted and reflected peaks broader and larger. After optimisation, the power in the TEM_{00} mode increased by a factor of 4.9, consistent with the power increase shown in Figure 6.21 and to which similar initial conditions applied. (The apparent movement of the TEM_{10} peaks is because we were only able to scan 40% of an FSR and the scanning break occurred between the two sets of data making laser frequency drift over the time between measurements look like movement between the modes.)

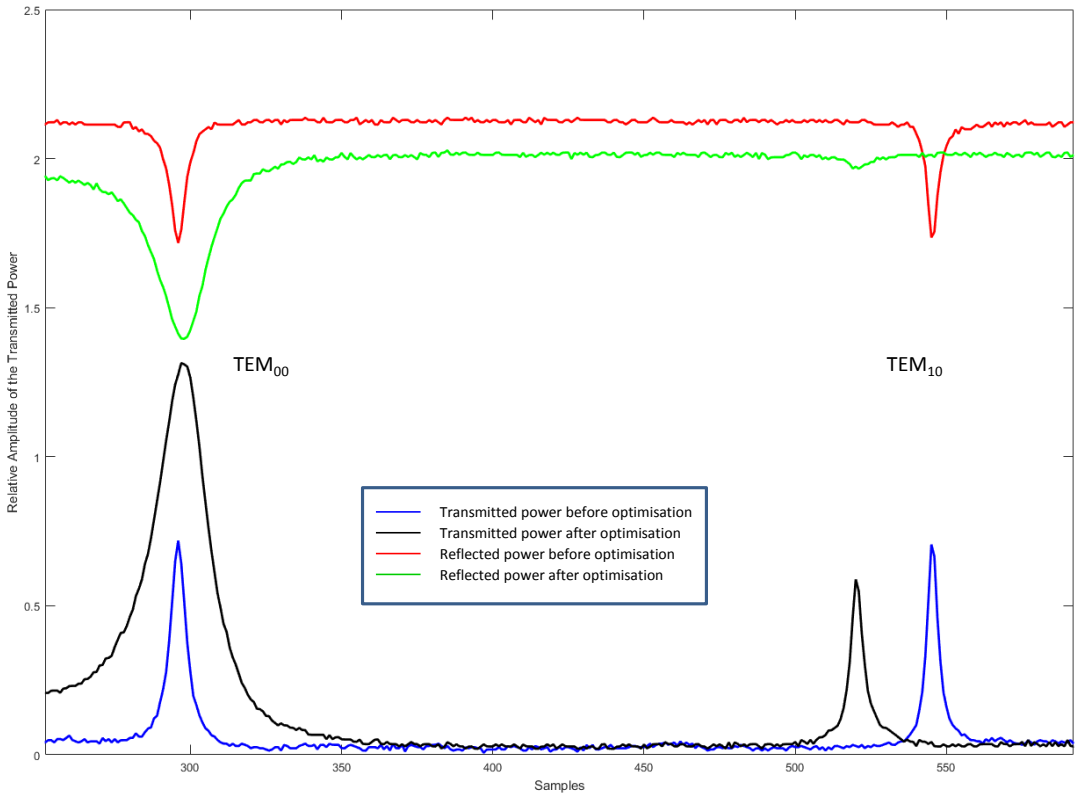

Figure 6.22: Optimisation Transfers Power to the TEM_{00} mode. The amplitude scale is a relative scale and the plots have been realigned to separate them and align the peaks of the TEM_{00} modes. The scales of the two reflected signals and the two transmitted signals are constant however and illustrate how a significant proportion of the power in the higher order mode, in this case the TEM_{10} is transferred to the the TEM_{00} mode by the optimisation process. After optimisation, the power in the TEM_{00} mode increased by a factor of 4.9, consistent with the power increase shown in Figure 6.21 and to which similar initial conditions applied. (The apparent movement of the TEM_{10} peaks is because we were only able to scan 40% of an FSR. As the peaks drift, the apparent distance between them appears to change whereas, in fact, they are only drifting nearer or further from the break in the scanning data.)

6.5 More Complete Wavefront Management

In the experiment reported above the focus was on improving tilt and offset and waist size and waist position, all attributes that can be optimised by careful beam alignment (for tilt and offset) and careful optics placement (waist size and position). However, spatial light modulators are capable of higher order manipulation across the wavefront, something that adjustment of more standard optical components cannot achieve.

The second experiment extended the wavefront sensing and correcting technique work to:

- reshape the beam to sustain transmission of a higher order mode; and
- repair local distortions in the wavefront to optimise power transmission further.

6.5.1 Selecting Higher Order Cavity Modes

Section 4.6 introduced the notion of higher order modes as equally valid solutions to Maxwell's equations. There are occasions when it is desirable to identify and isolate higher order modes, either to employ the mode or remove it from the optical beam. Higher order modes can be employed for mode division multiplexing, where individual modes are used in separate but parallel communication channels, increasing the effective carrying capacity of a fibre, Ref [104]. They can also be engineered to have special dispersion properties, making them ideal for certain jobs including:

- Dispersion compensation in femtosecond fibre lasers, see Ref [105] and Ref [106], and
- Non-linear frequency conversion at wavelengths not feasible with traditional single-mode fibres, see Ref [104].

The TEM_{10} mode is the simplest of these comprising two lobes of opposite phase along the horizontal axis, figure 6.23.

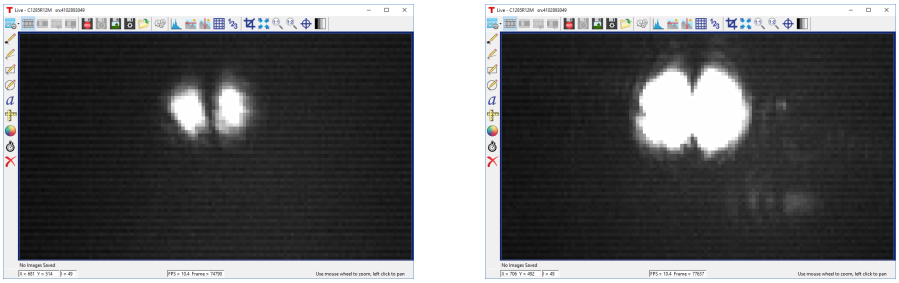

Figure 6.23: The TEM_{10} mode recorded on the CCD camera before (left) and after (right) optimisation.

The system was able to optimise and sustain transmission of the this mode once it had been locked onto using PDH locking. To achieve this an additional two degrees of freedom were added to each mirror, S Curves, figure 6.24.

These S Curves were created by adding a dither in which one half of the SLM's was advanced and the other was retarded by an equal and opposite amount. This advanced the phase of one side of the wave front, while retarding the phase of the other allowing the mode of the beam to be matched to the TEM_{10} mode of the cavity. The geometry

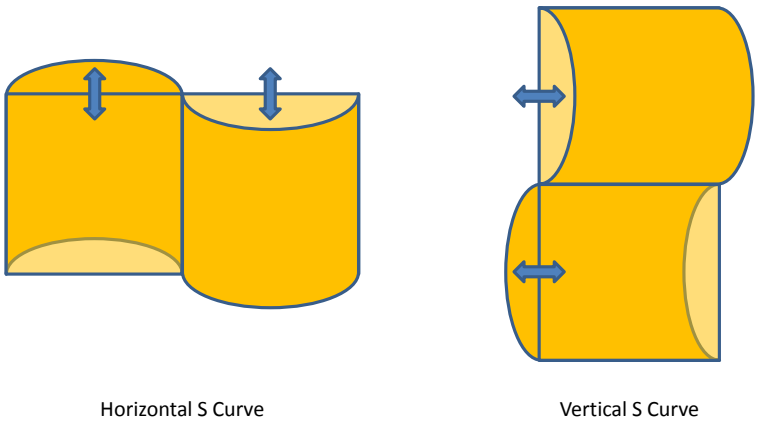

Figure 6.24: The S Curve dither applied in addition to and simultaneously with the tilts and full width curves used in the previous experiment.

of the S Curve dither is shown in figure 6.25. The maximum displacement between the left and right sides of the curve on the segmented SLM is 82% of the full 1.5 micron displacement or 1.16 wavelengths meaning there is more than adequate range to support the π or $1/2$ wavelength shift required to reinforce the TEM_{10}. For the continuous SLM the maximum phase shift achievable is nearly 3.5 wavelengths, nearly seven times that required for sustain the TEM_{10} mode.

Seg	Segmented Full S Curve			Continuous Full S Curve		
	Distance to middle of seg from LHS	Distance from centre of half circle	Offset factor for circle	Distance to middle of seg from LHS	Distance from centre of half circle	Offset factor for circle
	x = r-d	d	h (Seg)	x = r-d	d	h (Cont)
1	0	0.825	0.0000	0	1.0083	0.0000
2	0.3	0.525	0.6364	0.3667	0.6417	0.7778
3	0.6	0.225	0.7937	0.7333	0.2750	0.9701
4	0.9	-0.075	0.8216	1.1000	-0.0917	1.0042
5	1.2	-0.375	0.7348	1.4667	-0.4583	0.8981
6	1.5	-0.675	0.4743	1.8333	-0.8250	0.5798
7	1.8		-0.4743	2.2000		-0.5798
8	2.1		-0.7348	2.5667		-0.8981
9	2.4		-0.8216	2.9333		-1.0042
10	2.7		-0.7937	3.3000		-0.9701
11	3		-0.6364	3.6667		-0.7778
12	3.3		0.0000	4.0333		0.0000

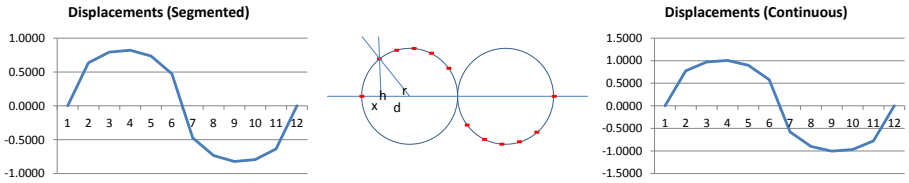

Figure 6.25: The geometry of the S Curve dither.

Figure 6.26 shows the results of adding the S Curve dither to the Tilt and Curve dithers when trying to optimise the TEM_{10} mode.The sequence applied was:

1. No optimisation;

2. Optimisation using just the Tilt dither and Curve dither;

3. Optimisation using just the Tilt dither, Curve dither and S Curve dither;
4. No optimisation once again;
5. Optimisation using just the S Curve dither;
6. Optimisation using just the S Curve dither and Curve dither; and
7. Optimisation using all three one again.

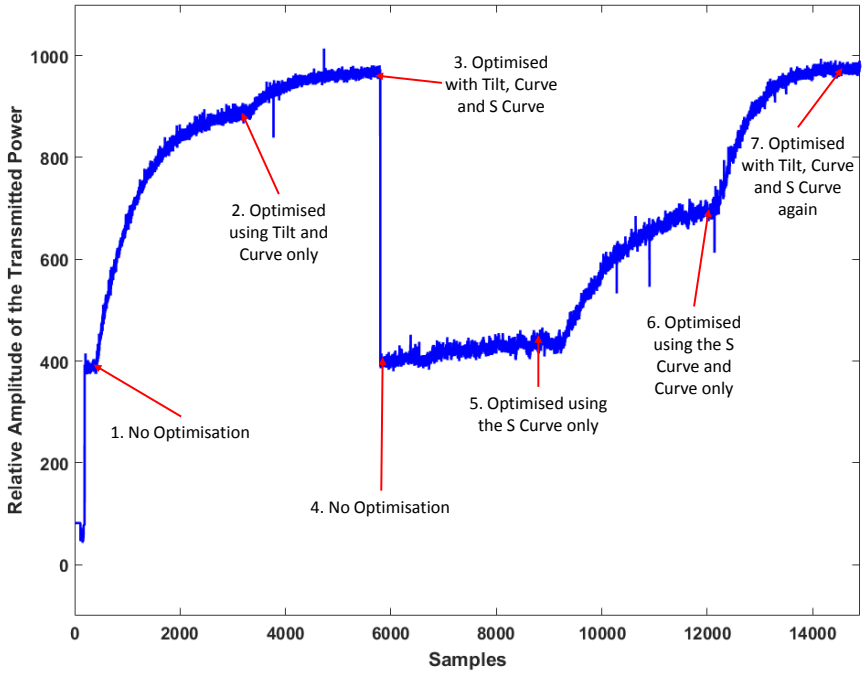

Figure 6.26: The result of applying the S Curve dither when the system is locked on the TEM_{10} mode. The optimisation process first applied just the Tilt and Curve dithers, then added in the S Curve dither before being flattened again and having the Curve dither applied on its own, followed by the addition of the Curve dither follower by all three once again. Each makes a contribution to increasing the power transmitted in the TEM_{10} mode.

From these results it is clear that all three contribute to maximising the amount of power transmitted through the cavity in the TEM_{10} mode. For this particular optical situation, that is the particular tilt, offset, waist size and waist position incident on the cavity at this point in time:

- The Curve dither increased the transmitted power by approximately 50%;
- The Tilt dither increasing the transmitted power by approximately 71%;
- The S Curve dither increasing the transmitted power by approximately 23% ; and
- All three dithers increasing the transmitted power by approximately 145%.

While these percentages depend completely on the initial optical conditions it is clear that all three dithers working together increase the degree of optimisation by more than any individual dither or any pair of dithers is able to achieve. Figure 6.27, taken from the

front panel of the LabVIEW control system, shows the final shape of the mirrors after all three dithers were applied.

The question of how much each mirror and how much each dither type contributes to the final outcome is an interesting if somewhat irrelevant question. Irrelevant because, while two SLMs will always be required for a comprehensive alignment, how much each contributes to the outcome is irrelevant as long as transmission is maximised. As to which contributes what, this depends on the initial condition of the incident beam, the coupling between the SLMs (which is determined by their position in the optical arrangement) and the different gains (which determine how quickly each dithered shape contributes to the total optimisation) and dithers applied.

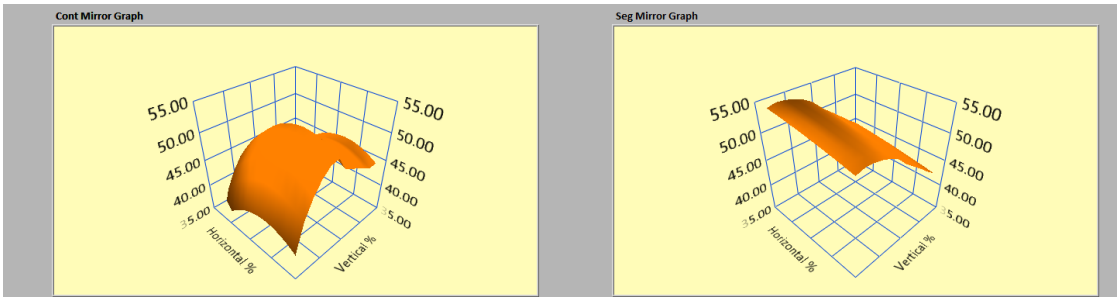

Figure 6.27: The shape of the SLMs after all three dithers had been applied. Once again the vertical scale is exaggerated.

Interestingly the S Curve dither contributes only an 11% increase in optimisation if it is applied on its own. Figure 6.28 shows why. Working alone the S Curve dither has to try to correct tilt, curvature and phase with too few degrees of freedom and some internal conflicts arise. The shape on the continuous mirror, the left hand side, shows how the S Curve has been used to generate a phase shift on the left hand side of the mirror but left nearly flat on the right hand side to create more of a tilt.

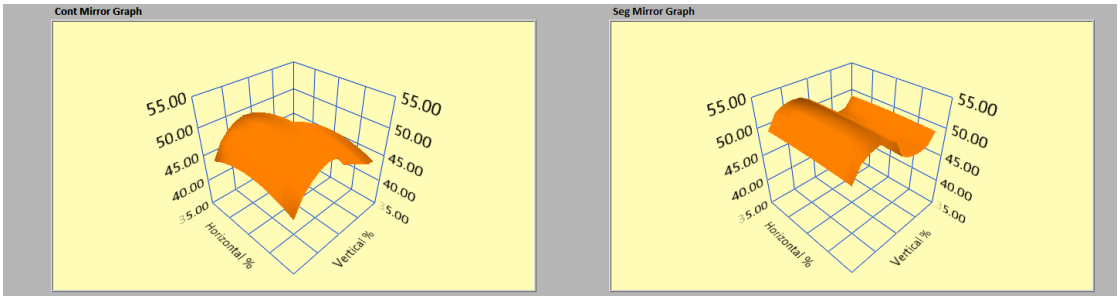

Figure 6.28: The shape of the when optimised using the S Curve dither alone. Once again the vertical scale is exaggerated.

6.5.2 Actuating More Complex Groups of Pixels

Of course, the ideal modulation pattern would be to dither each of the 280 pixels on the two SLMs individually. This would give wavefront control limited only by the number of SLM actuators and any coupling between them. In practice, our experiment was limited

by the size of the FPGA used which was resources limited to implementing 32 spatial regions per mirror.

Figure 6.29 shows the results of a compromise coding pattern designed to explore the greatest definition possible with the FPGA available. Adjacent squares of the same colour were assigned the same Hadamard code so, for example, cells 4, 5, 16 and 17 were each assigned the same code. Cells on the vertical edges were grouped in sixes resulting in the use of a total of 30 codes for each mirror. (The yellow squares in each corner are, of course, fixed flat at 0% offset by the structure of the SLMs.) The Hadamard code set used this final experiment was generated by the Matlab expression $Hadamard(64)$ which generated a matrix of 64 by 64 Hadamard codes. These were applied to each SLM in the pattern shown, the data collected and then integrated over 64 samples to cover a full code set across both SLMs. As the SLMs were run at approximately 2 kHz, each optimisation update took around 32 milliseconds to complete.

	1	2	3	4	5	6	7	8	9	10	11	12
1	1	2	3	4	5	6	7	8	9	10	11	12
2	13	14	15	16	17	18	19	20	21	22	23	24
3	25	26	27	28	29	30	31	32	33	34	35	36
4	37	38	39	40	41	42	43	44	45	46	47	48
5	49	50	51	52	53	54	55	56	57	58	59	60
6	61	62	63	64	65	66	67	68	69	70	71	72
7	73	74	75	76	77	78	79	80	81	82	83	84
8	85	86	87	88	89	90	91	92	93	94	95	96
9	97	98	99	100	101	102	103	104	105	106	107	108
10	109	110	111	112	113	114	115	116	117	118	119	120
11	121	122	123	124	125	126	127	128	129	130	131	132
12	133	134	135	136	137	138	139	140	141	142	143	144

Figure 6.29: The actuator pattern applied alternately to the continuous mirror and the segmented mirror for the more complex dithering. 30 Hadamard codes are used and applied to adjacent actuators in groups of four as indicated by the squares of the same colour. Cells on the vertical edges were grouped in sixes to employ a total of 30 codes for each mirror.

The resulting shape of the two mirrors after optimisation is shown in figure 6.30. The yellow trace shows how the power increase over time as the system is optimised. While the surface looks complex it must be remembered that the z-axis, that representing deviations in the actuator depth, has been exaggerated for clarity.

This compares to a similar but slightly different performance when the Tilt Curve optimisation is used, figure 6.31 on the same set of initial conditions. This is to be expected as the beam has been set up with some care and deviations from the Gaussian form have been

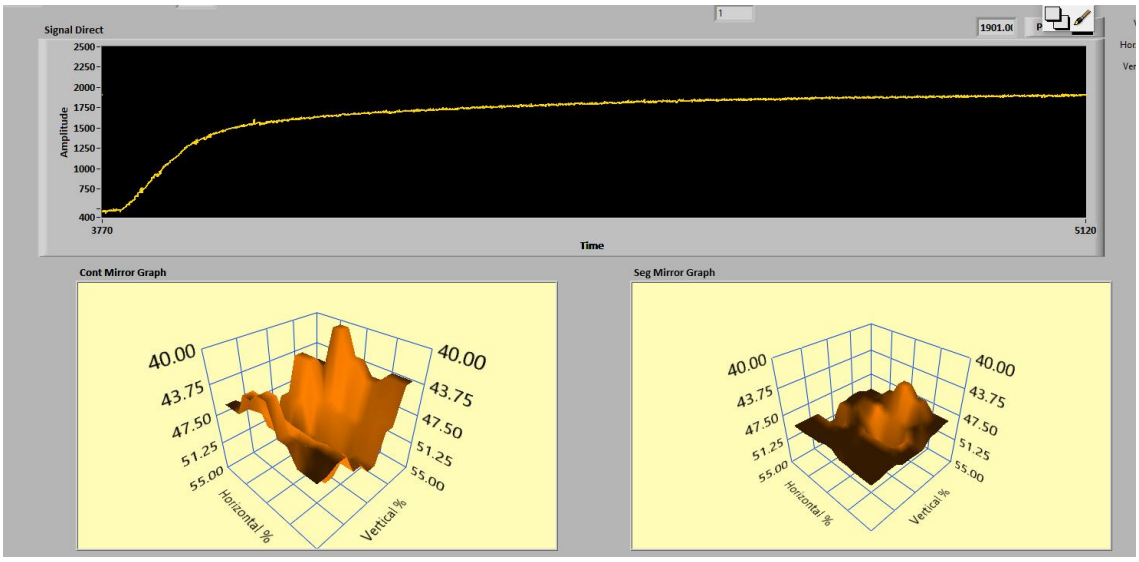

Figure 6.30: The shape of the mirrors when optimised using the more complex dither on both SLMs. The yellow trace shows how the power increase over time as the system is optimised. The vertical scale is exaggerated for clarity with a movement of 66% representing a phase shift of 2π.

minimised to the greatest extent possible (although some are introduced by the surface of the segmented SLM). The basic shape of the two SLMs in each instant is very similar with, in each case, the Continuous SLMs rising to a peak in the furthermost corner of the plot and the Segmented SLM forming a saddle shape.

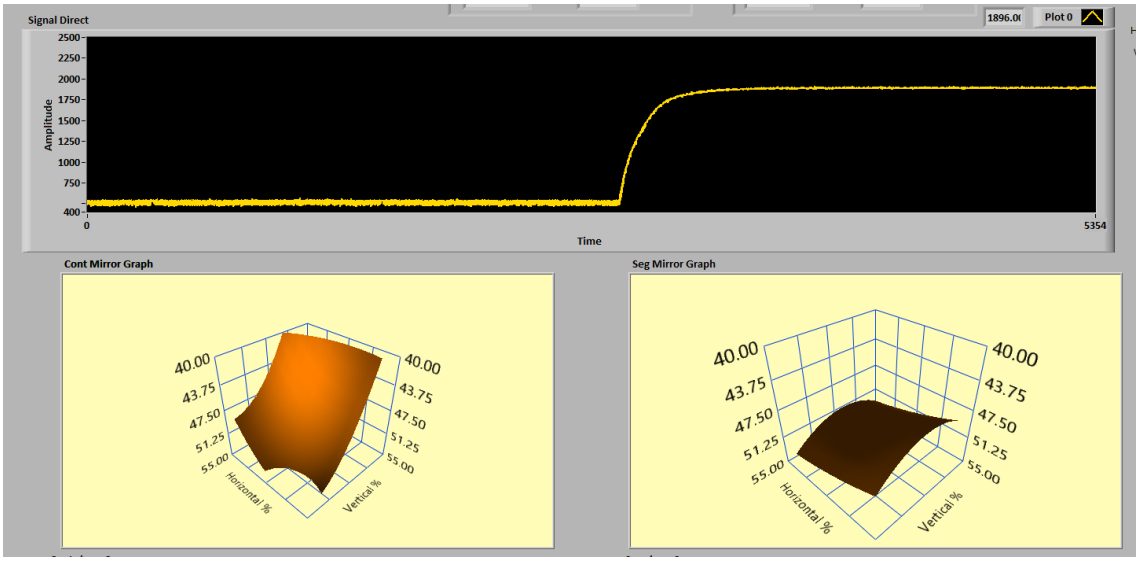

Figure 6.31: The shape of the mirrors when optimised using the standard Tilt Curve optimisation algorithm. The yellow trace shows how the power increase over time as the system is optimised. The vertical scale is exaggerated for clarity with a movement of 66% representing a phase shift of 2π.

Figure 6.32 compares the two optimisation forms. For this set of initial conditions, the Tilt Curve Algorithm, the red curve, starts more slowly but reaches its maximum level faster.

The 30 Code algorithm starts strongly but takes longer to converge on its maximum, perhaps reaching a very slightly higher value (0.65% higher over the last 100 samples), as expected.

Understanding the shape of these curves involves consideration of four factors:

- The SLMs were running at approximately 2 kHz, see 5.4.1. Given samples were the result of integrating across 32 codes, the sampling frequency was approximately 60 Hz, so the 1500 samples along the x-axis in Figure 6.32) represents around 25 seconds. The sampling rate could be increased by pushing the performance of the Multi-Drivers (perhaps to 3 kHz) or switching to the X-Drive if and when the new software drivers become available (the manual claims up to 400k kHz although we have been unable to test this because of the driver issues).

- The time to maximum power was not considered important in this proof of concept and there was no attempt to maximise it. Both reached their maxima after about 20 seconds.

- The gains used were different for each curve. The gain applied to the Tilt dither was $1 \cdot 10^{-4}$ while the gain applied to the Curve dither was $5 \cdot 10^{-5}$ and were selected during early experimental runs as being effective. The gain applied to the 30 codes algorithm was 10^{-4} on each actuator explaining its slightly faster initial rise.

- The noise performance of the system is dominated by two factors - first, the mechanical noise created by the actuation of the SLMs and then their operating frequencies. The system used the same Matlab control tools as the wavefront sensing experiment reported in Chapter 5 and the noise performance was examined in detail in section 5.6. Even if the mechanical noise were removed, laser variations occurring faster than 60 Hz are too fast to be corrected and remain as noise in the incident beam. Both the Tilt Curve algorithm and the 30 codes algorithm shared the same control system so their noise performance is the same.

6.6 Potential Use in Gravitational Wave Detectors

While this proof of concept demonstrates that this new technique works, its value to long baseline gravitational wave detectors, like LIGO, depends on its noise performance. We have not addressed this issue and it would have to be studied in detail before any recommendation could be made. This is suggested as part of the future work program, see 7.7 but a few early comments are made below.

6.6.1 Noise Considerations in LIGO and it Future Variants

LIGO is an extraordinarily sensitive instrument and any additions to the optical layout would have to meet very stringent phase noise requirements. The Instrument Science White Paper of 2017 (see Ref [107]) gives clues about the acceptable level of phase noise for current and future LIGO instruments in its discussion on the introduction of squeezed states. The requirements are:

- For LIGO A+ - 17 mrad;

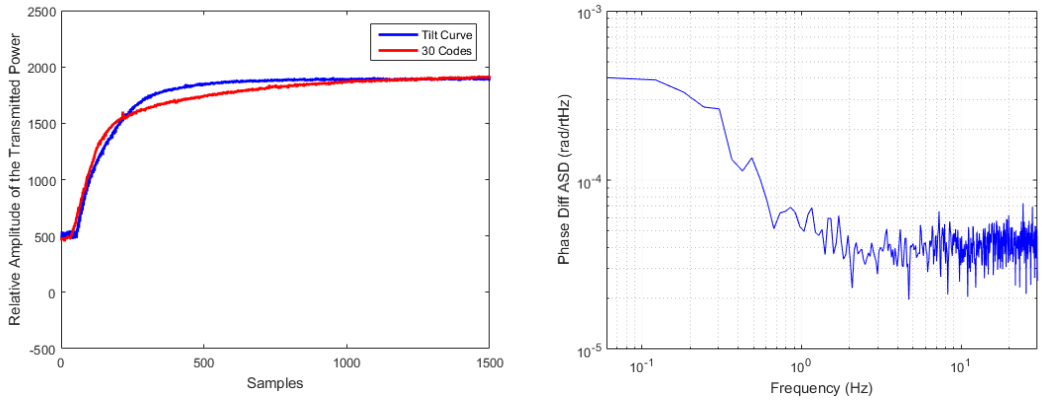

Figure 6.32: The figure of the left is A comparison of the algorithm optimisation rates. The 30 Code algorithm starts more quickly (because of the gains used) but takes longer to converge on the maximum, perhaps reaching a very slightly higher value (0.65% higher over the last 100 samples), as expected. The figure on the right is extracted from Figure 5.17 and shows the operational performance of the single photodiode which is limited by mechanical noise and the 60 Hz operational frequency, see 5.6

.

- For LIGO Voyager - 4.4 mrad; and
- For LIGO Cosmic Explorer - 2.5 mrad.

A first step in forming a view about the applicability of this technique would be to determine the phase noise tolerance it must meet and, if necessary how a total budget would be shared between various innovations.

Two factor mitigate in favour of this technique.

- First, long baseline interferometers, like LIGO, are Michelson Interferometers and if the paths lengths of the arms are the same and the reflectivities of the test masses are equal, most of the noise introduced is common to both arms and is rejected when the beams are recombined. Of course, the arm lengths are never perfectly equal and phase noise is a side band on the signal beam, part of which is redirected to the dark port. The degree to which this happens is governed by $\omega\Delta L/c$ where *omega* is the angular frequency, ΔL is the difference in arm length and c is the speed of light. At low frequencies this effect becomes very small so, for example, for an arm length discrepancy of 1m (huge) and an operating frequency of 3 Hz, it would be 10^{-8}.
- Second, the system would only be required to operate at low frequencies. It's role would be to probe thermal effects which happen over tens of seconds or minutes or look for optical aberrations which are largely static. The system could be operated at very low frequencies, well below the frequencies of interest to LIGO, Ref [107].

6.6.2 Noise Sources on the SLMs

The principle SLM noise sources which would have to be considered include:

1. Mechanical Noise - The SLMs are mechanical devices and introduce a noise floor at around 30 Hz even idling (turned on but not being actuated) and after moving the SLM controllers (which are fan cooled) off the optical bench. This was explored in detail in the first experiment, see Section 5.6 and, in particular, Figures 5.16 and 5.17. Doubtless this could be reduced but clever engineering and we do not know if there are any troubling higher frequency components, but it is hard to imagine an instrument of LIGO's sensitivity tolerating this as it currently stands.

2. Actuator Noise - At first sight the error in the actuators seems less of a concern. The technique removes modulation depth as a factor during demodulation, see equation 3.12, and, being a Michelson Interferometer, this noise is common to both paths and should be rejected. Of course, as pointed out above, 6.6.1, this is only true if both paths length are precisely equal but at low frequencies the effect is likely to be small. We have not tested the performance of the SLMs but Boston Micromachines claims the precision in actuator displacement to be less than 1 nm or $9 \cdot 10^{-4}$ wavelengths for 1064 nm laser light, Ref [97].

3. Sampling Limitations - As currently operating, the sampling frequency (of about 60 Hz) limits the definition with which the wavefront can be examined. This is less of a problem because operating frequency adds no jitter to the beam and the phenomena being examined, such as temperature variations of optical defects typically vary over longer time frames, tens or even hundreds of seconds.

4. Diffraction Effects - SLMs come in two varieties, segmented and continuous. The segmented mirror can have a dramatic effect on the shape of the beam emerging from it (see figure 6.6) and while, once again, this should be common to both arms, the deformable mirror with the continuous surface is likely to be a better option.

5. Power - The continuous SLMs can handle up to 20 W/cm^2, see Ref [108] but we have not explored how powers of these levels may affect its performance. This would have to be investigated before a recommendation could be made.

6. Asynchronous Codes - Unlikely though it sounds, we would need to be sure that applying asynchronous spatial codes (for example pushing the top of the mirror forward while pushing the bottom backwards) does not affect induce any discernible motion in test masses.

In short, therefore, although the technique looks promising, a detailed evaluation of it's noise performance under load is required before it can be considered for inclusion in the LIGO program.

Finally, we do note that the 2017 LIGO Instrument Science White Paper identifies the criticality of mode matching between the interferometer and the Output Mode Cleaner to minimise loses. Because output mode matching varies with the interferometer's thermal state it is a particularly difficult problem to address and the collaboration is investigating the use of deformable mirrors for adaptive mode matching, see Ref [107] and Ref [109].

6.7 Conclusion

This proof of principle experiment clearly demonstrated that it was possible to increase the power transmitted through a cavity by modulating the surfaces of two SLMs in a predefined way and then using an error signal derived from the transmitted power recorded by a single diode to reshape the wavefront. The initial experiment focused on optimising tilt, offset, waist position and waist size and did not attempt to correct higher order wavefront aberrations.

The system was also able to follow and correct for minor disruptions in the optical path caused by manual (or mechanical) interventions.

The second experiment looked at more complex wavefront corrections, initially adding an 'S' shaped curve to support optimising a high order mode, specifically the TEM_{10} mode, and then to adjust the wavefront in greater detail to further improve the TEM_{00} mode. Both proved possible.

The key advantages demonstrated by the system were that:

- The system can detect and correct high-order wavefront distortions with a single element detector. It will function effectively as long as the signal falls somewhere on the detector's surface, making beam alignment less critical;
- These corrections can either be used to adjust the wavefront to accommodate optical distortions in real time or to identify aberrations in optical components so that they can be repaired or replaced;
- The system can manage high powers. While not specified in the documentation, the current SLMs can manage $20W/cm^2$ across all wavelengths. The segmented mirror is much more tolerant because of its thicker mirrored surface and can handle at least $100W/cm^2$ and more at infra red wavelengths Ref [108].
- The system can operate in-line, using the signal beam itself without disturbing the information this beam is carrying. This means that it can follow the precise optical path the signal is travelling and respond to the precise conditions encountered by that beam;
- The system can isolate and support higher order modes for preferential transmission as required;
- The system can respond in real time (or at least within the reaction time of the SLM controllers) and adapt to small changes in the optical path as they arise; and
- The information is returned for processing through a single channel eliminating channel and ADC noise which is now common and disappears in the demodulation process.

The system suffers from two disadvantages, however, which will need to be addressed if it is to become a serious option for ultra sensitive interferometers like LIGO:

- The modulations introduce a degree of noise onto the signal which depends on the modulation depth employed. While it appears to reduce considerably once the system is engaged, this is an effect of the integration process (samples were only reported once an integration was complete) and it remains a source of noise that will be unacceptable in some instruments.

- The system, as currently presented, is dependent upon PDH locking, a well proven but independent process, and can lose lock if a significant mechanical shock occurs.

Further work is required to determine whether enough noise can be removed at the appropriate frequencies to make it useful to long baseline gravitational wave detectors.

Future Work

Section 6.7 outlined the advantages and disadvantages of the new approach and these suggest some problems and improvements for the next stage in this work. These are discussed below. The ultimate goal would be to develop a system which, using the most basic of alignments, can identify and lock onto the TEM_{00} mode (or any higher order modes desired), and maintain and continually adjust this lock across one or multiple cavities with minimal noise. When desirable, it should also be able to identify and correct for aberrations in the components.

7.1 Increasing Speed and Definition

The SLMs used in the experiments reported here were mechanical and the limitations of the current drivers reduced the reliable speed to around 2kHz. If the SLMs were powered by X-Drivers (supported with functioning driver software) Boston Micromachines claims that the operating frequency could be up to 400 kHz, see Ref [99]. If the SLMs were replaced with an Optical Phased Arrays (OPAs) this speed would be limited by the capabilities of Electro-optic Modulators (EOMs) which can operate at frequencies in the tens of megahertz. The EOM used in this experiment, for example, was run at 24.9MHz.

The OPA is a head of tightly packed laser emitters where the phases of the adjacent lasers can be adjusted to steer and focus the beam using interference. While considerably more expensive to operate (each emitter requires it's own EOM at around $6,000 each), the combined, emitted beam behaves very like the reflection from the segmented mirror and changes to the output are fast and predictable Ref [110].

7.2 Incorporating Frequency Locking in the System

Currently the system relies on PDH locking to lock the frequency of the laser and length of the cavity together. This is a well proven and understood process but must be completed before optimisation can begin. If the SLMs were asked to replace PDH locking they would have to oscillate at terahertz frequencies (clearly impossible, even for the OPA) or be integrated into the cavity as either a rear mirror or as the last mirror in a Kahlili cavity, see 7.8. In either event, light cannot be transmitted through the SLM so these arrangements would only work for a limited range of optical arrangements.

7.3 Using Artificial Intelligence to Align Multiple Cavities

The process of optimising transmission through a single cavity is well understood and relatively easy to master either manually or using this system. Once multiple cavities are introduced the problem becomes much harder. Now the frequency of the laser must be matched to the varying lengths of several cavities at once and the optimal solution is not to be found by simply scanning a single FSR. The introduction of artificial intelligence along with the faster locking solution proposed in section 7.2 might make this achievable. The system would be asked to map the complex space created by the power transmitted through the multi cavity system, learn to distinguish the different modes (current achieved manually using a CCD image) and then find a solution that locks all the cavities simultaneously using the PDH signal.

Figure 4.4 provides a clue as to just how difficult this may prove. The reflected troughs (and consequentially the transmitted peaks) produced by the PDH locking process are very sharp and the frequency space between them is largely featureless (although some higher order modes are present). It may be that the gradient descent engines used in most artificial intelligence engines are unable find enough features to build a functional map of the space.

Finding the TEM_{00} mode using a CCD image is quite difficult. Most of the time the screen displays a variety of higher order modes and it can be quite difficult to know which way to make adjustments until the image has been resolved in one plane (say the horizontal plane) and one of the simpler higher order modes has been identified as a starting point for adjustment. Artificial intelligence may be better at identifying and responding to the complex images of higher order modes than humans.

7.4 Supporting Higher Order Modes on Demand

The capacity to identify and preferentially transmit higher order modes is something quite unique about this solution. Currently it relies on identifying these modes first, visually, locking onto them and then using the system to optimise them. Ultimately it would be nice to combine integrated locking, artificial intelligence and cavity optimisation to create a system which can respond to the request for a specific mode, for example, the TEM_{01} mode from a single manual request.

7.5 Spatial Multiplexing

At first sight, it might appear that the spacial separation of channels would allow the increased transmission of data across optic fibre networks by simultaneously transmitting data through parallel spatial channels. Unfortunately the recovery of meaningful data from those channels requires integration by the length of the code employed and this the integration length must always be greater than or equal to the number of spatial channels. There seems to be no capacity advantage over single channel transmission, even if spatial compression algorithms are applied. We do note that higher order modes have been used to multiplex signals, (see Ref [104]), so perhaps further consideration is warranted. It may be that spatial multiplexing offers some advantages in encryption that have yet to be explored.

7.6 Developing New Forms of Phase Based Imaging

The first experiment demonstrated what was effectively a phase camera and allowed the user to explore the microscopic shape of a reflecting surface in some detail. This may provide an interesting, alternative imaging system for examining structures at nanometre resolution.

A pilot flying in thick fog or a scuba divers swimming in murky water, for example, often through scattered light to the object of interest, for example, the end of the runaway. If the target were illuminated by a laser beam one might expect that, in most cases the scattered light would return first and the last reflection would be that from the distant target. If these could be gated and separated using orthogonal codes, the scattered light could be discarded and the target illuminated. This is essentially the technique used to track space debris in Ref [44].

7.7 Assisting the Detection of Gravitational Waves

The detection of gravitational wave requires extremely sensitive interferometry. On land everything is done to remove noise and the operating powers within the cavities are very high. The potential for applying this technique to long baseline interferometers is discussed in detail in section 6.6 and we identify the extensive work that must be done before any recommendation could be made.

This includes understanding the phase noise budget that the interferometers could tolerate and then investigating:

1. The mechanical noise of the SLMs and any possible mitigations;
2. The actuator noise;
3. Any important noise lost in a low sampling frequency;
4. Diffraction effects for the SLM surface;
5. The impact of high power on the SLMs; and
6. The impact of using asynchronous spatial codes.

In space the alignment of different satellites using very faint signals gives this wavefront sensing approach an advantage worth examining further. The technique uses a single diode, a single channel and a single ADC to return the signal to the CPU for processing. The single diode creates a target that is larger and indifferent to where the signal strikes it. The common channel and ADC mean that a source of noise can be removed that has not been mitigated by other means. The SLM means that small, non-linear aberrations in the returning wavefront can be identified more easily.

7.8 Using a Khalili Cavity to Compensate for Mirror Aberrations

A Kahlili cavity is a three mirror cavity in which the transmitted field beyond the second mirror is reduced to zero by adjusting the shape of the third mirror, figure 7.1. While

anecdotally difficult to build, using an SLM as the third mirror in a Kahlili cavity should allow the amplitude of the transmitted field to be reduced to zero and ensure all the light is reflected back into the interferometer even if the quality of the second mirror is suboptimal. This may help reduce the cost of optics in future gravitational wave detectors.

Figure 7.1: A Kahlili cavity. Using an SLM as the third mirror in a Kahlili cavity can reduce the transmitted field to zero and ensure all the light is reflected back into the interferometer even if the quality of the second mirror is poor.

Appendices

The following pages contain samples of the code used to undertake the experiments reported in this thesis, specifically:

- The Matlab codes used to drive the SLM for part of the first experiment (included because it uses the old SLM drivers;
- The LabView Host code used for the 'Tilt Curve dither' part of the second experiment; and
- The LabView FPGA code used in the second experiment; and
- The Open, Main Routine and Close modules embedded in LabView for the 'Tilt Curve and S Curve dither' part of the second experiment.

These summaries should allow a reader to understand and reproduce the software suite used to conduct these experiments, although, frankly, they are more useful as an interpreting key to the entire group of programs. The complete suite is available from tralph@netspeed.com.au.

The system details were as follows:

- Windows PC running Windows 10 Home version 1709, OS Build 16299.371, 64 bit
- Matlab version 2014b
- LabView version 2016

The first experiment was conducted using the mirror Boston Micro Machines drivers from 2016 which used the

- UPDATE_multiDM(driver_info, code(:)); % Send the new matrix to the mirror

groups of commands. These scripts would have to be updated to the new drivers to repeat these experiments. This is a simple substitution process and the new driver format is shown in Appendix8.4.

8.1 Appendix 1 - Matlab Code for Controlling Each Half of the SLM Independently

The Matlab code below was used to control the SLM in the first part of the first experiment. Each half of the mirror was controlled independently using the Hadamard codes [+1 -1 +1 -1] and [+1 +1 -1 -1].

Readers should be aware that this programs uses the controller drivers from 2016. These have now been updated and different driver commands are required. Comments are shown in blue, commands in black and SLM updates in red.

8.1.1 Matlab Code to Manage the Mirror in Halves

```
Author : Tarquin Ralph
Date : 21 November 2014
Purpose : This program cycles the SLM through four states, all flat, LHS advanced, RHS
advanced, all advanced.

function SLMCodedHalves

%% Reinitialise Matlab
close all; % Close all figures
clc; % Clear Command Window:

%% Open and initialise the Multi-DM driver USB connection
mapping_ID=4;

[error_code, driver_info] = OPEN_multiDM(mapping_ID); % From manual
code = 50*ones(12,12); % A matrix to hold the 144 values to be sent to the mirror for a
given set of codes
UPDATE_multiDM(driver_info, code(:)); % Flatten the mirror

% create our clean up object
cleanupObj = onCleanup(@() cleanMeUp(driver_info));

%% Enter diving parameters
moddepth = input ('Code Depth '); % 2.69% is equivalent to pi/4 for a right angle mirror
see Mirror Codes and Depths.xlsx
time = input('Pause length ');
time=time*1000000;

%% Cycle the mirror continuously through the 4 states namely: 1,1 -1,1 1,-1 -1,-1

for j=1:200e6;% Cycles through 10,000 times
j;
% % scan modulation depth
% if (moddepth  = round(j/1000)+1)
% moddepth = round(j/1000)+1;
% disp([int2str(moddepth) ' = ' num2str(moddepth/100*1500*2, '%.0f') ' nm']);
% end

% State 1 = Codes are 1,1 and mirror is flat and advanced a constant 3*moddepth/2
code(:,:)= 50 + moddepth/2;
UPDATE_multiDM(driver_info, code(:)); % Send the new matrix to the mirror
for i = 1:time
```

```
end

% State 2 = Codes are -1,1 and mirror is advanced at an angle from 0 to 2*moddepth
code(1:6,:)= 50 - moddepth/2;
code(7:12,:)= 50 + moddepth/2; % Set bottom to maximum value of offset
UPDATE_multiDM(driver_info, code(:)); % Send the new matrix to the mirror
for i = 1:time
end

% State 3 = Codes are 1,-1 and mirror is advanced at an angle from 0 to 2*moddepth
code(1:6,:)= 50 + moddepth/2;
code(7:12,:)= 50 - moddepth/2; % Set bottom to 0
UPDATE_multiDM(driver_info, code(:)); % Send the new matrix to the mirror
for i = 1:time
end

% State 4 = Code is -1,-1 and mirror is flat and advanced a constant moddepth/2
code(:,:)= 50 - moddepth/2;
UPDATE_multiDM(driver_info, code(:));% Send the new matrix to the mirror
for i = 1:time
end

end; % Repeat the next cycle

finish = input ('Will finish ');

%% Reset mirror to flat and then disable and close Multi-DM driver USB Connection
code(:) = 50;
UPDATE_multiDM(driver_info, code(:));
error_code = CLOSE_multiDM(driver_info);% From manual
end

% Fires when main function terminates
function cleanMeUp(driver_info)
end
```

8.2 Appendix 2 - Details of the LabView Host Code Used in the Second Experiment

8.2.1 Function

The following paragraphs summarise the functioning of the Host program *6.1 Tilt Curve SCurve Simultaneous - Both SLMs(Host)*. Its function is to measure the transmitted power from the cavity across up to 32 spatially distinct areas defined by the Host program by separating, demodulating and integrating the signal collected in the single photo-diode.

An important and simplifying principle of the design is the asynchronous operation of the two main processing loops, Error Correction Creation (paragraph 8.2.3) and Mirror Dithering Routine (paragraph 8.2.3). The system updates the error signal based on current variations in the transmitted power as fast as it can. The current value is sampled by the Mirror Dithering Routine whenever it is ready to process a new value.

The paragraphs below include a picture of the front panel with details the inputs and outputs used by the system, and the FPGA wiring diagram which shows the overall layout and links to the details of the individual modules outlined in the Processes section.

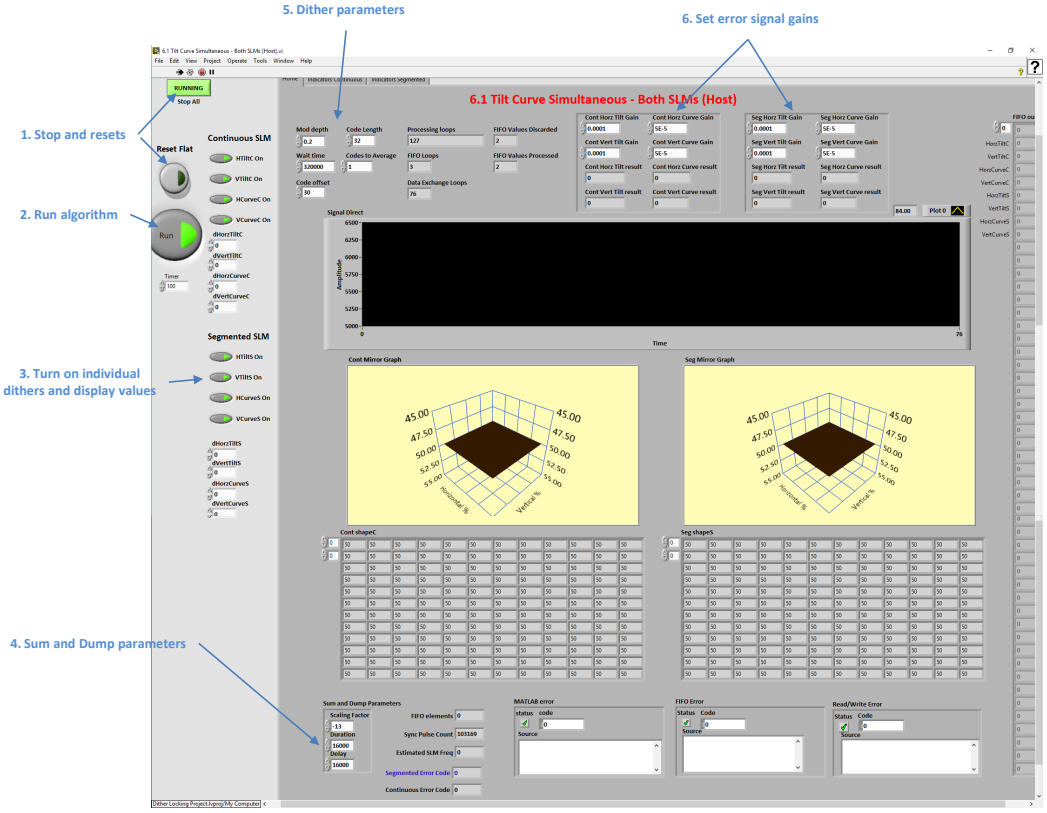

Figure 8.1: The Front Panel of the Host Program used in Experiment 2

8.2.2 Inputs

The program inputs are:

1. Reset Mirror: Sets mirror flat at 50
2. Run: Runs Optimisation or allow mirror to be reset flat at 50
3. HTilt On: Allows optimisation to update the value of dHorzTilt. If off the Horizontal Tilt remains unchanged
4. VTilt On: Allows optimisation to update the value of dVertTilt. If off the Vertical Tilt remains unchanged
5. HCurve On: Allows optimisation to update the value of dHorzCurve. If off the Horizontal Curve remains unchanged
6. VCurve On: Allows optimisation to update the value of dVertCurve. If off the Vertical Curve remains unchanged
7. Mod depth: Changes the depth of dithering for HTilt, VTilt, HCurve and VCurve.
8. Wait Time: A counter used to force Matlab to wait for the mirror to settle after a dither has been sent to the mirror so that this completes before beginning on the next optimisation.
9. Code Offset: Sent to FPGA to determine where to start picking codes from the Hardamard code vector so that they are aligned with those used on the mirror. Recheck 0 is the right default value!
10. dHorzTilt: The output of the optimisation. Multiplied by HTilt to develop the new value for Horizontal Tilt, a component of shape the total arrangement of the mirror.
11. dVertTilt: The output of the optimisation. Multiplied by VTilt to develop the new value for Vertical Tilt, a component of shape the total arrangement of the mirror.
12. dHorzCurve: The output of the optimisation. Multiplied by HCurve to develop the new value for Horizontal Curve, a component of shape the total arrangement of the mirror.
13. dVertCurve: The output of the optimisation. Multiplied by Vcurve to develop the new value for Vertical Curve, a component of shape the total arrangement of the mirror.
14. Stop All: Stops execution of the FPGA Data Exchange and Main Processing loops
15. Code Length: Sets the code length for the FPGA.
16. Codes to Average: Tells the FPGA how many codes to average across.
17. Tilt Gain: Multiplies the new values of dHorzTilt and dVertTilt before they are passed onto the next iteration. (Does this also reduce the 50 flat value, i.e. should it be applied to the gain less the 50 which is then added back.)
18. Curvature Gain: Multiplies the new values of dHorzCurve and dVertCurve before they are passed onto the next iteration. (Similarly, does this also reduce the 50 flat value, i.e. should it be applied to the gain less the 50 which is then added back.)
19. S&D Delay: Defines how many ticks are ignored to let the mirror settle before the Sum and Dump starts - passed to the FPGA program

20. S&D Duration: Defines the length of the Sum and Dump - passed to the FPGA program

21. S&D Scaling Factor: - Defines the power of 2 by which the integrated signal is scaled down to keep it an I16 - passed to the FPGA program

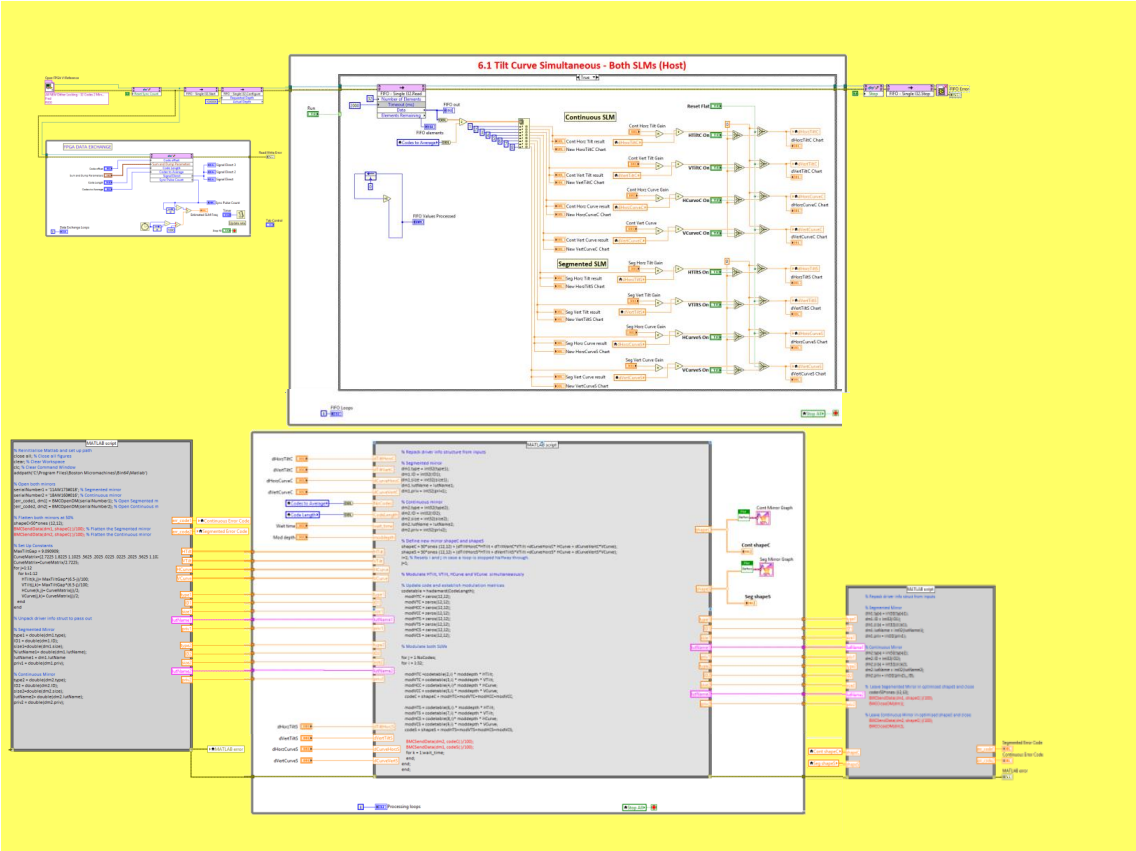

Figure 8.2: The Wiring Panel of the Host Program used in Experiment 2

8.2.3 Processes

The following paragraphs outline the processing modules employed in the program. Each is accompanied by a diagram of the LabView code.

Call FPGA opens the FPGA and calls the FPGA program.

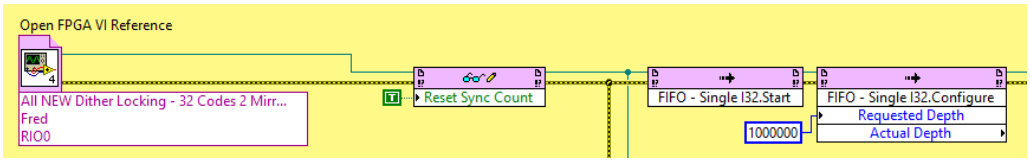

Figure 8.3: Experiment 2 - Call FPGA

Open SLMs opens the USB connection to the SLM controllers and sets up the initial operating state. It:

1. Opens both mirrors;
2. Flattens both at 50%;
3. Decomposes dm1 (Segmented) and dm2 (Continuous) and passes into the main processing loop;
4. Establishes the control matrices for tilting and curving the SLMs; and
5. Reports any Matlab errors.

FPGA Data Exchange :

1. Sends Code offset, Sum and Dump Parameters, Code Length and Codes to Average to the FPGA;
2. Receives Signal Direct and the Sync Pulse Count from the FPGA; and
3. Estimates the SLM frequency by dividing the number of Sync Pulses by duration.

Error Signal Creation :

1. Collects individual error signals from the FPGA;
2. Multiplies them by their respective gains; and
3. Updates the individual parameter for the SLMs.

The module include a Reset Flat option which resets both SLMS to 50% until the latch is released.

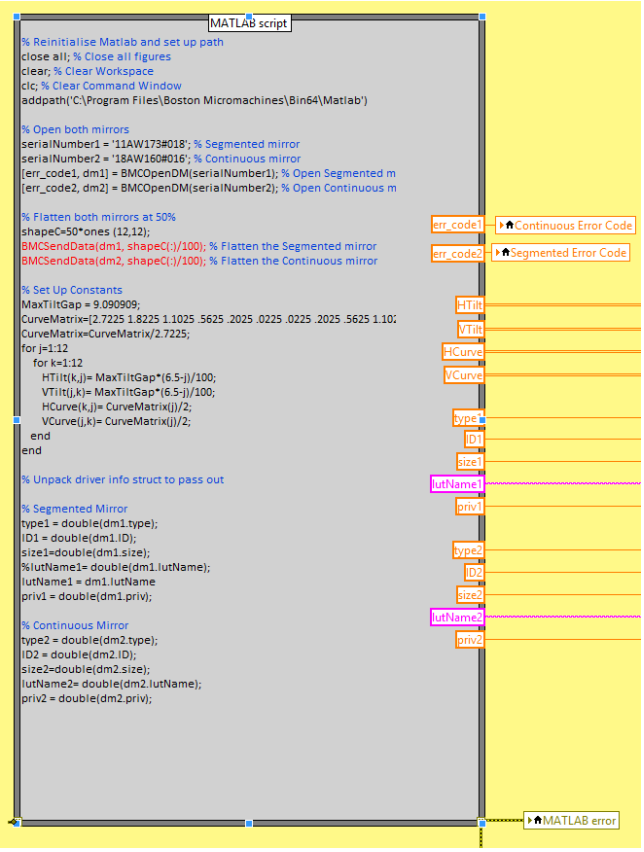

Figure 8.4: Experiment 2 - Open SLMs

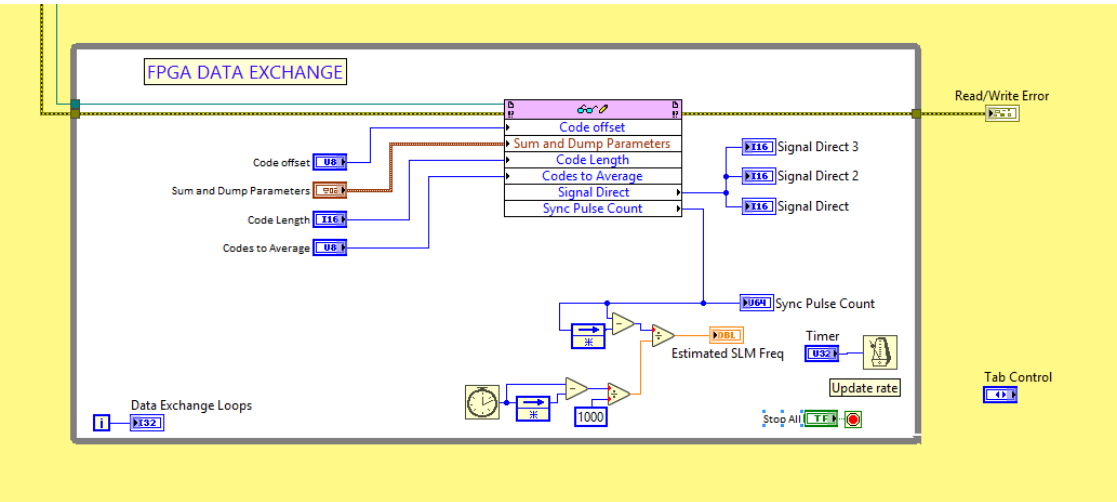

Figure 8.5: Experiment 2 - FPGA Data Exchange

Error Signal Update : passes the error signals from the latest assessment to the SLM Main Routine for the next processing run. Specifically it:

1. Passes the last values for dHorzTilt , dVertTilt, dHorzCurve, dVertCurve, dHorzTiltS, dVertTiltS, dHorzCurveS and dVertCurveS into the processing loop as the new starting point; and

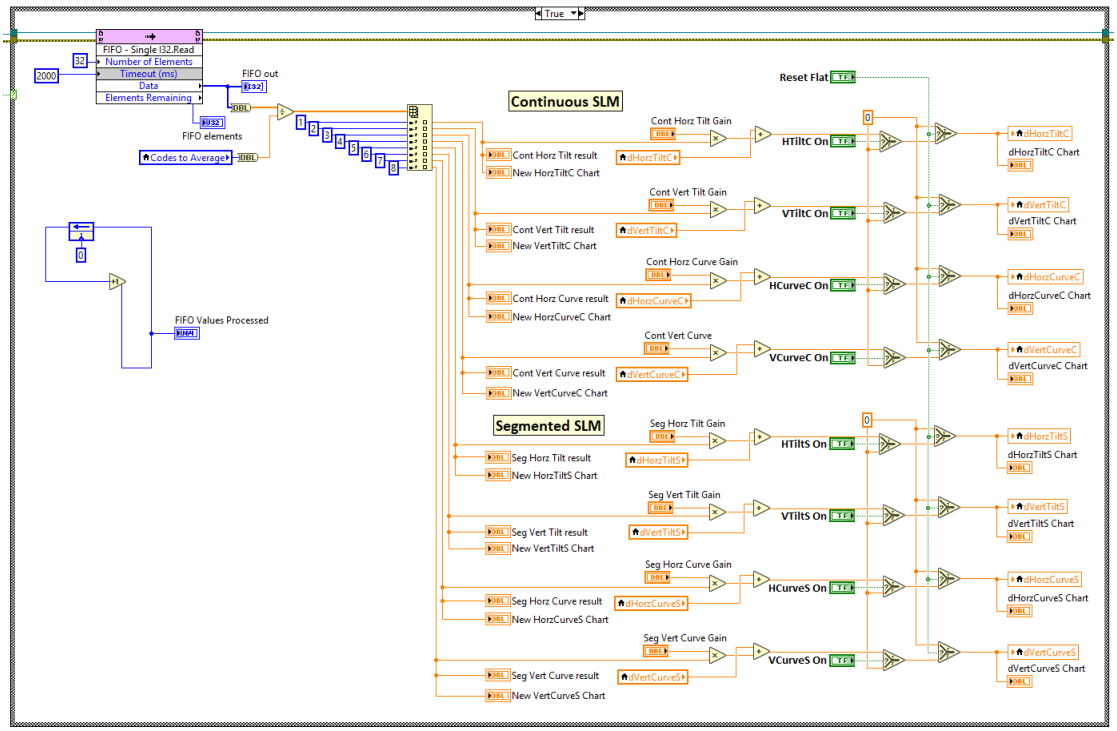

Figure 8.6: Experiment 2 - Error Signal Creation

2. Passes Codes to Average, Code Length, Wait Time, dm1 and dm2 into the main processing loop.

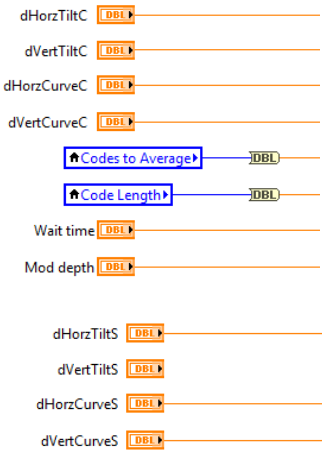

Figure 8.7: Experiment 2 - Error Signal Controls

Dithers the SLMs by:

1. Reconstructing dm1 and dm2;

2. Importing the latest error signal and update matrices to create a new shape for the SLM;
3. Dithers this new shape;
4. Resetting the mirror to the value of shape; and
5. Passes the new data to the graphing and reporting modules.

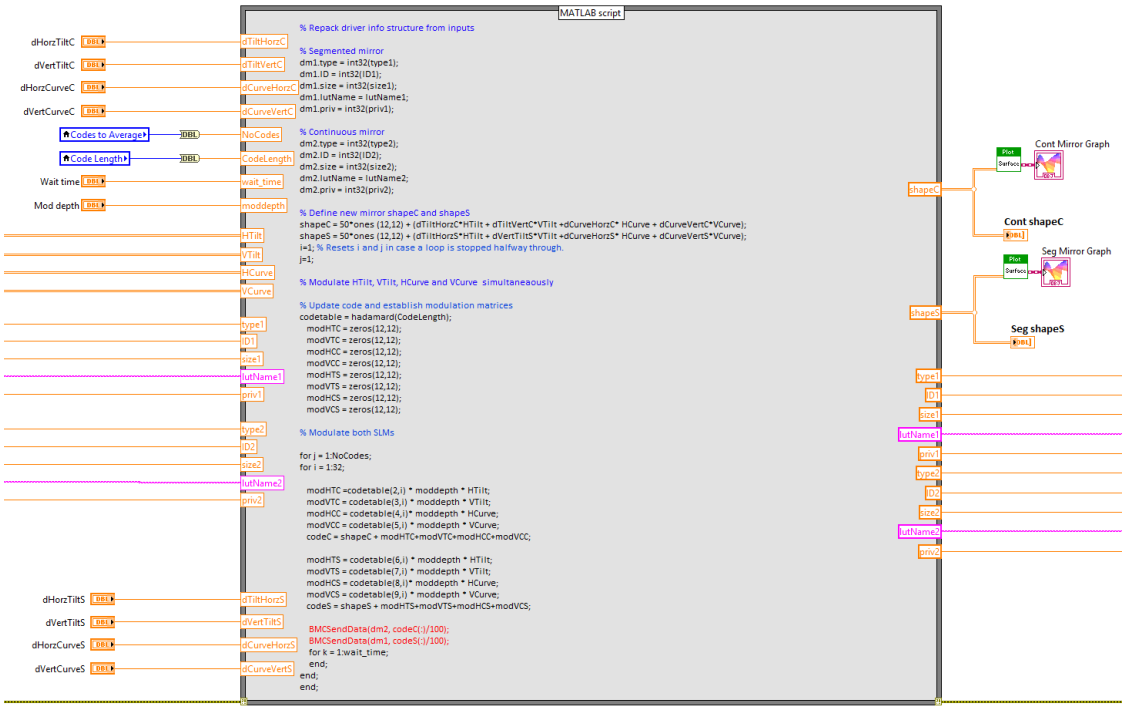

Figure 8.8: Experiment 2 - Main Routine

Close SLMs by:

1. Reassembling dm1 and dm2;
2. Flattening the segmented mirror at 50% and closes it; and
3. Leaveing the continuous mirror at the last value of 'shape' and closes it.

Call FPGA opens the FPGA and calls the FPGA program.

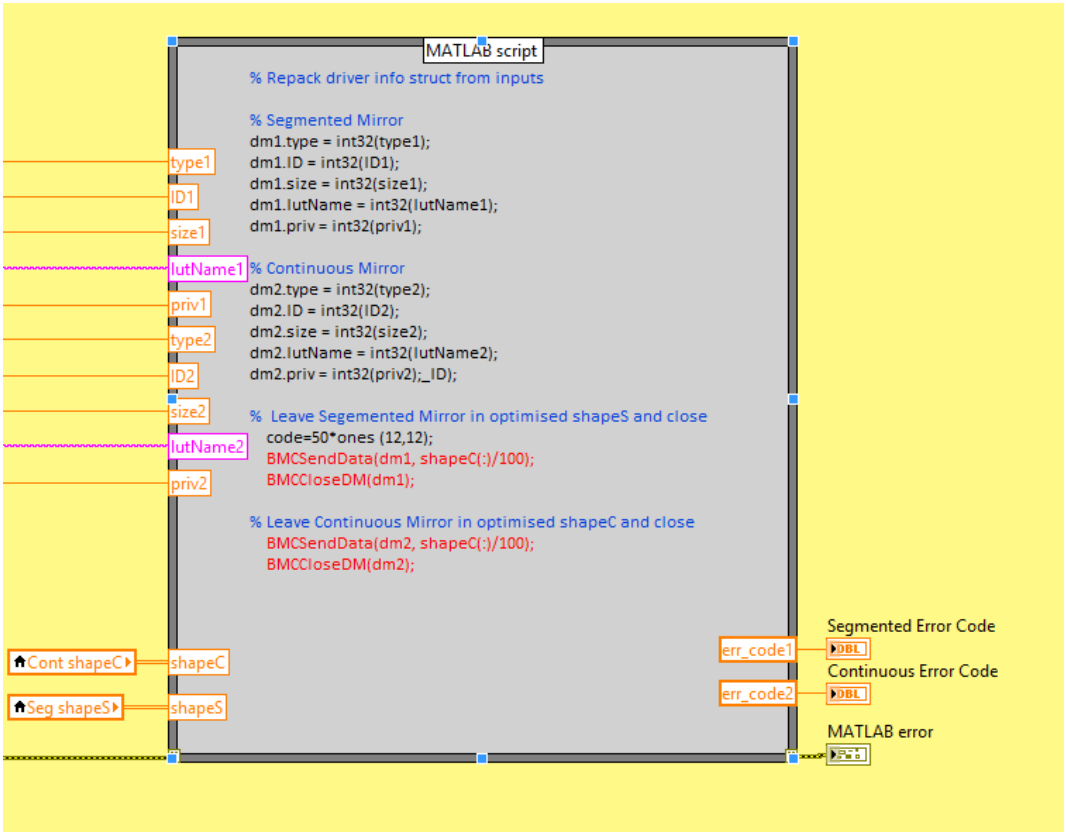

Figure 8.9: Experiment 2 - Close FPGA

8.2.4 Outputs

The program outputs shown on the LabView front panel are:

1. Main Loops: Increments each time the main processing loop runs (even for the False False case.
2. FIFO out: Displays arrays received from FIFO
3. Horizontal Tilt Result: Reports FIFO element 1 (Second element)
4. Horizontal Curve Result: Reports FIFO element 33 (34th element)
5. Vertical Tilt Result: Reports FIFO element 65 (66th element)
6. Vertical Curve Result: Reports on FIFO element 97 (98th element)
7. Signal Direct: Graphs the signal received by FPGA on ADC
8. Current Mirror Map: Plots current shape of the mirror
9. FIFO elements: Shows the number of FIFO elements remaining. (???)
10. Sync Pulse Count: Shows number of Sync Pulses received. From FPGA
11. Estimated SLM Freq: Estimates the SLM frequency by dividing the number of Sync Pulse Count by seconds.

12. Matlab Error: Shows errors in Matlab processing loop
13. FIFO Error: Shows any FIFO errors
14. Read/Write Error: Shows read/write errors
15. Shape: Shows the values of each element of the SLM (as plotted on the Current Mirror Graph).

8.3 Appendix 3 - Details of the LabView FPGA Code Used in the Second Experiment

The following paragraphs summarise the functioning of the FPGA program *32 Codes, 2 Mirrors 1 Sync Pulse (FPGA).vi* which was used to collect and demodulate data from the second experiment. Once again, the paragraphs below include a picture of the front panel with details the inputs and outputs used by the system, and the FPGA wiring diagram which shows the overall layout and links to the details of the individual modules outlined in the Processes section.

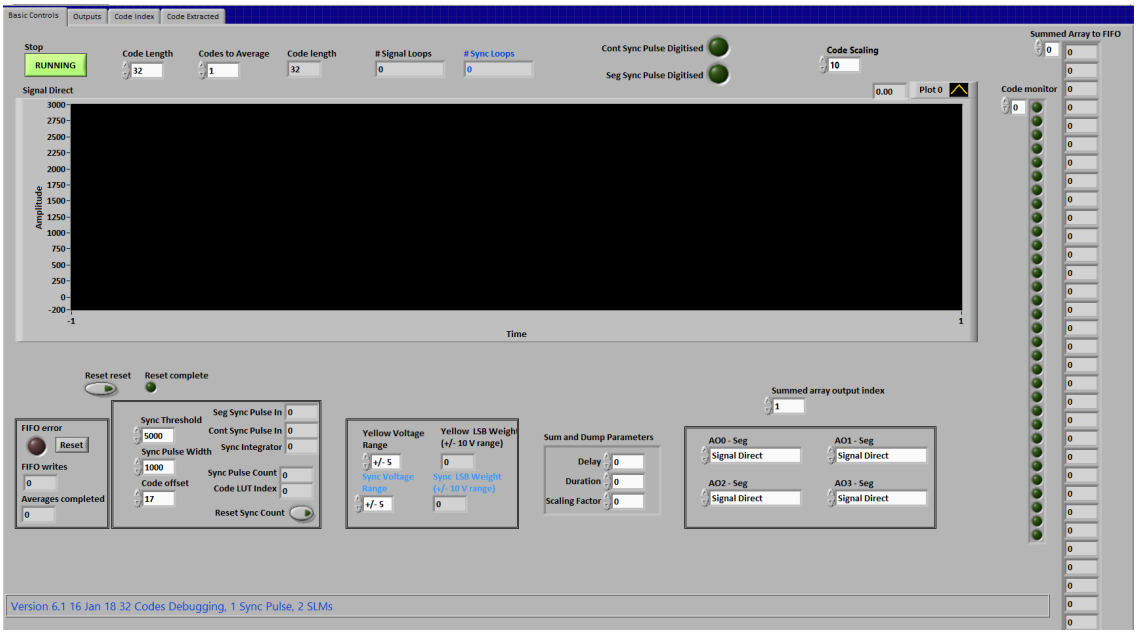

Figure 8.10: The Front Panel of the FPAG used in Experiment 2

8.3.1 Function

The program summarised here is *All NEW Dither Locking - 32 Hadamard Codes Reset Handshaking (FPGA)*. Its function is to optimise the power transmitted a cavity by reshaping the wavefront to match the mode of the cavity more closely.

8.3.2 Inputs

The program inputs are:

1. Stop: Stops execution of all while loops
2. Code Length: Length of Hadamard Code used - should always be 32 - usually controlled by the HOST program
3. Codes to Average: Defines how many codes length the signal is integrated across - usually 1 and controlled by the HOST program

4. Reset reset: Usually controlled by the HOST program and set to TRUE once the reset of the FIFO and Sync Count have been confirmed
5. FIFO error Reset: Resets the FIFO - usually controlled by the HOST program
6. Sync Threshold: Sets the limit the above which the Sync Pulse Digitised Boolean is set to TRUE (i.e determines that a Sync pulse has been received
7. Sync Pulse Width: Sets minimum width of a Sync Pulse to avoid false positives
8. Code Offset: Specifies where in the Hadamard code sequence demodulation should begin - usually controlled by the HOST program
9. Reset Sync Count: Sets the sync pulse count back to 0
10. Yellow Voltage Range: Controls sensitivity of ADC used to collect the transmitted power
11. Sync Voltage Range: Controls sensitivity of ADC used to collect the sync pulse power - set at 5V - Sync pulse is 3.3V
12. S&D Delay: Defines how many ticks are ignored to let the mirror settle before the Sum and Dump starts - usually controlled by the HOST program
13. S&D Duration: Defines the length of the Sum and Dump - usually controlled by the HOST program
14. S&D Scaling Factor: - Defines the power of 2 by which the integrated signal is scaled down to keep it an I16 - usually controlled by the HOST program

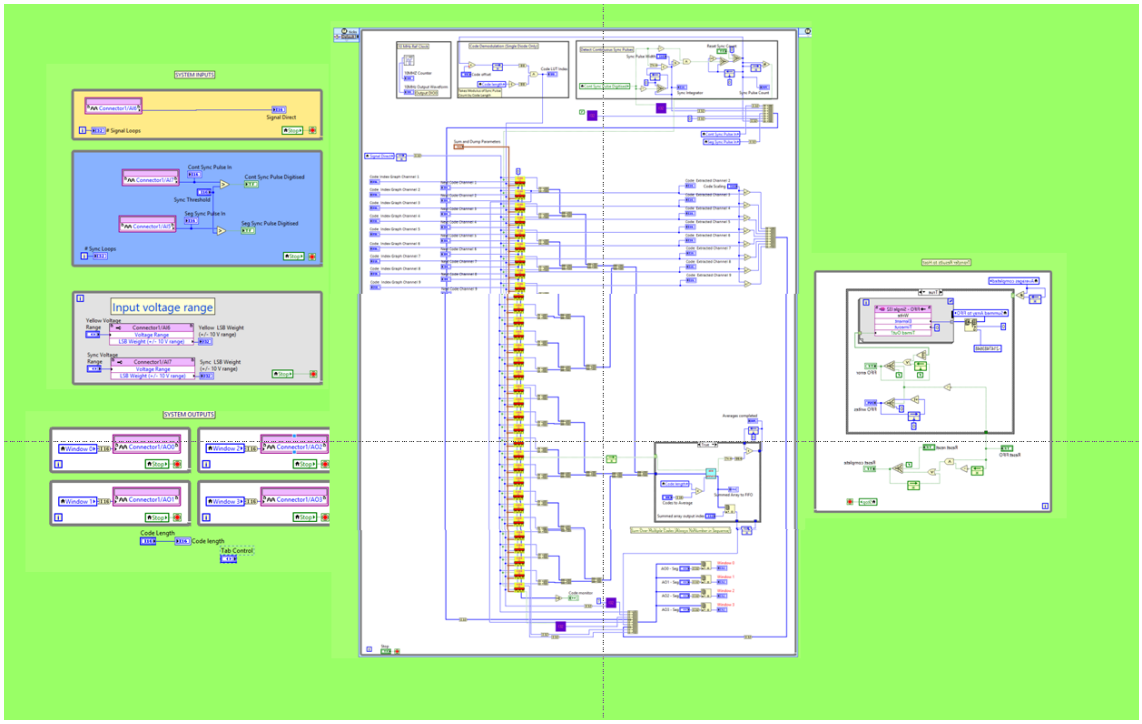

Figure 8.11: The Wiring of the FPAG used in Experiment 2

8.3.3 Processes

The following paragraphs outline the processing modules employed in the program. Each is accompanied by a diagram of the LabView code.

System Inputs allows the system to:

1. Collect transmitted power from diode and makes digitised version available for demodulation as Signal Direct. Also counts the number of samples and report as #Signal Loops;
2. Measure the Sync Pulse from a Mirror driver, reports the raw data as Sync Pulse In, and sets Sync Pulse Digitised to true if the value exceeds the Sync Threshold (typically 5000);
3. Set the inputs sensitivity for the two ADCs; and
4. Receive the nominated Code Length from the Host program and makes it available to the FPGA.

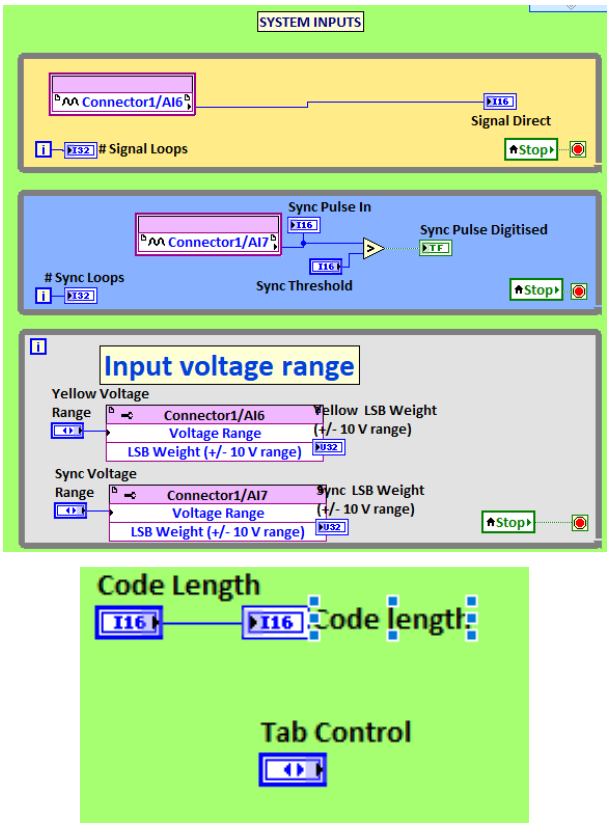

Figure 8.12: Experiment 2 - System Inputs

System Diagnosis allows the system to provide 4 windows on the progressive values of important quantities during the running of the program, specifically (see definitions below):

1. Signal Direct
2. Code LUT Index
3. Summing
4. Sync Pulse In
5. Sync Pulse Digitised
6. Sync Integrator
7. Sync Pulse Connect

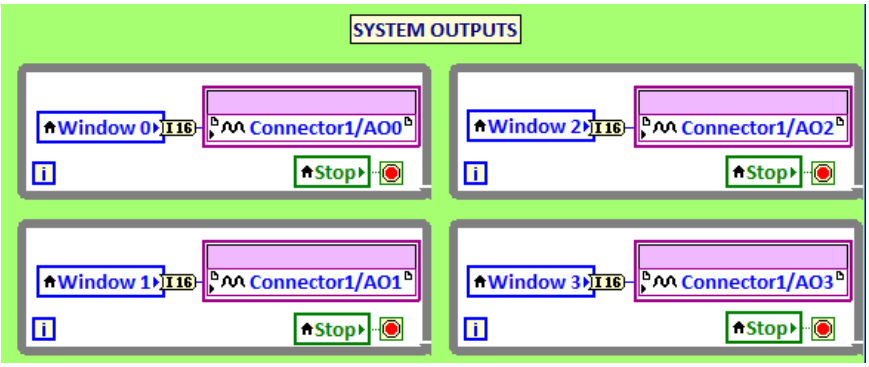

Figure 8.13: Experiment 2 - System Diagnosis

Reference Clock provides a 10MHz reference clock to synchronise various components of the experiment. The *LabView RIO 7856-R* does not have a reference clock input. This routine produces a 10MHz reference clock by flipping a bit between 1 and 0 every 4 ticks (the FPGA runs at 40MHz) that can then be used to link other clock in the experiment. Not used for this experiment. Outputs the 10MHz Output Waveform as a digital signal through DIO0 and displays the 10MHz Counter which runs between 1 and 4

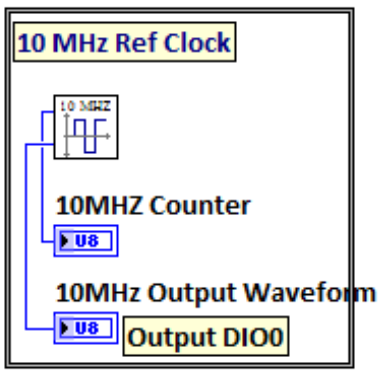

Figure 8.14: Experiment 2 - Reference Clock

Sync Pulse Detection detects each new sync pulse unambiguously. Every time the mirror changes shape (which takes $31\mu s$) there is a drop in the 3.3V Sync Pulse follower by a settling time (measured at around $72\mu s$ depending on the depth of the transition) while it rattles and shudders into position and becomes usable again. Being USB the time between changes is unpredictable which makes the useful sampling time around 75% of the total time of operation (total time - change and settle time - margin for USB variability).

This routine detects a digitised sync pulse above the threshold level and starts a counter, the Sync Integrator, which only integrates when the sync pulse is up. If this exceeds the Sync Pulse Width then we know it's a real Sync Pulse. Only when the pulse has exceeded then Sync Pulse Width and fallen to 0 (i.e. below the Sync Threshold) are both inputs to A TRUE and a real Sync Pulse has been identified and the Sync Pulse Count is increased by 1.

This count can be reset to 0 by pressing the Reset Sync Count. This and Reset FIFO are reset by the Host each time a new configuration cycle is started.

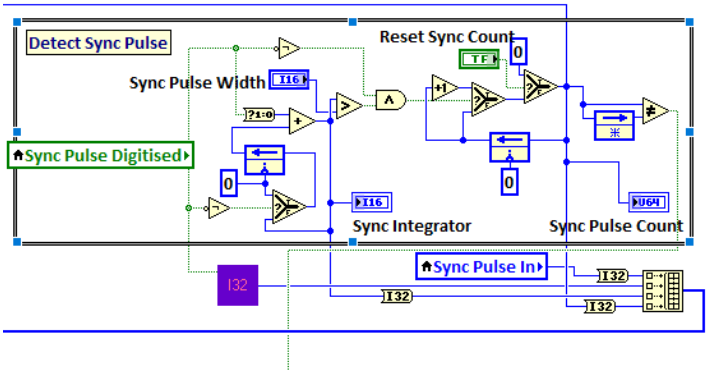

Figure 8.15: Experiment 2 - Sync Pulse Detection

Hadamard Code Management is a pointer that ensures that the right Hadamard code is used in the right channel at the right time. Each time a sync pulse is detected (above) the pointer across the 32 columns of the 32 by 32 Hadamard Code matrix is advanced by one so the next column of codes can be applied to the demodulation.

The position of the pointer is controlled by three things:

1. Sync Pulse Count - each time this is incremented a new mirror configuration has been provided and the demodulation must use the next code in the cycle
2. The Code Offset - an offset which aligns the internally generated code with the codes applied to the SLM and accommodates feedback delays etc.
3. The Code Length, which tells the count when to wrap.

These are combined into the Code Index for use below.

Demodulation and Integration splits the signal into 32 channels, demodulates it and integrates and normalises the result. The signal is now split into 32 channels and put through the Dither Hadamard Code32.vi. The .vi contains the 32 by 32 Hadamard Matrix as a vector 1024 elements long. Inputs are:

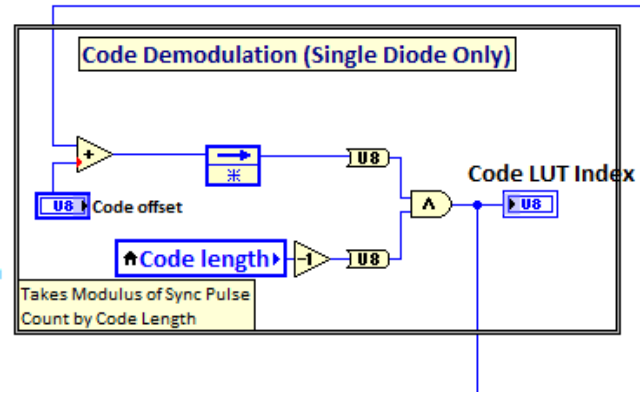

Figure 8.16: Experiment 2 - Hadamard Code Control

1. The digitised transmitted power from the diode;
2. Code Index;
3. Code (default 1);
4. Code monitor in;
5. Sum and Dump Parameters; and
6. Sync Pulse Down.

The outputs are:

1. Next code;
2. Code monitor out;
3. New Value Available;
4. Summing; and
5. Output.

The .vi manages code selection, display the current code as the Code Monitor, and then sums and dumps the signal using a further Sub.vi - Sum Dump and Wait Tester v5.vi which sums and dumps the input signal while a counter is in a specific range and produces the scaled output unique to that channel for that code tick.

- The counter remains on 0 when Sync Pulse is True. (this occurs after the Sync Pulse drops and is usually just for 1 tick.) Once it switches to False the counter starts and keeps incrementing until the Sync Pulse is True again.
- The counter ignores the first 2,000 ticks (set in the Host as Delay)
- It then accumulated the next 10,000 ticks (set in the Host as Duration)
- It divides this sum by $2^{-}13$ (set in the Host as Scaling Factor)) , outputs this number and resets the total to 0
- It then sits and waits until the sync pulse drops again

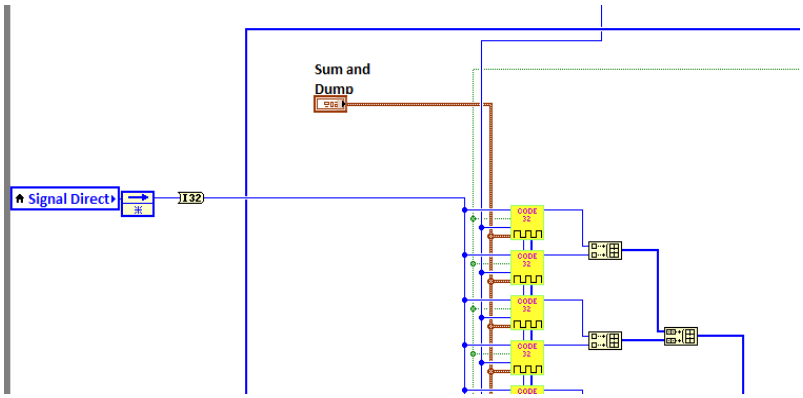

Figure 8.17: Experiment 2 - Demodulation and Integration

Output Consolidation takes the output of each channel and adds it to the sum of the previous values until the total number of sums is the Code length multiplied by Codes to Average. Once this total has been reached, a new value is made available to the FIFO as Summed Array to FIFO. Until this has been reached, the previous version is available to the FIFO.

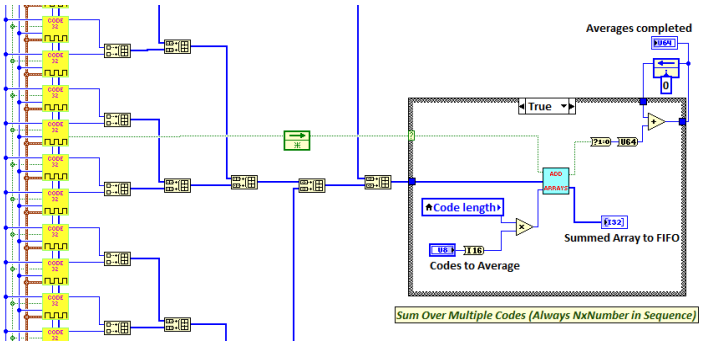

Figure 8.18: Experiment 2 - Output Consolidation

System Interrogation consolidates the 7 working variables that can be assigned to output 0-3 and inspected using the Oscilloscope or Windows 0-4 on the second Tab. Each time a new value is made available Averages Completed increases by one.

FIFO Exchange passes the summed array to FIFO generated in the routine above, writes -2147483648 (-2^31) into the 0^{th} position and passes it to the Host through the FIFO as an I32 number. It is only written when a new value is available as defined by the value of Averages Completed changing from the previous value.

The For Loop has no value of N but takes in summed array to FIFO through an auto indexed tunnel. The default behaviour is that this becomes the value of N so the loop operates N times where N is the size of the array.

Similarly the Timed Out indicator is set to True if the FIFO is not ready for the next set of values. These are gathered into an array and passed to the Or Array Elements. If

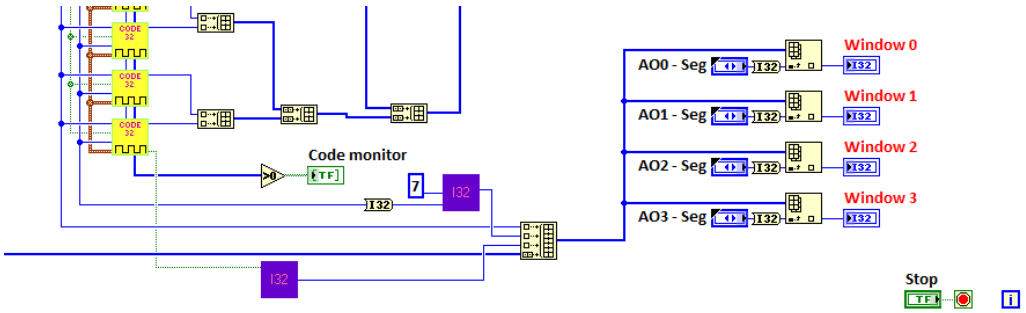

Figure 8.19: Experiment 2 - System Interrogation

any of the writes failed, one element of this array is TRUE which feed through to create a FIFO error. It also includes routines to: a. Count the number of FIFO writes - Blue Box. If this number is different to Averages Completed, writes are being missed. This is reset by the reset routines at which point it will no longer match the Averages Completed. Reset the FIFO and ensure it is ready whenever the Host determines this is necessary, which is after each iteration.

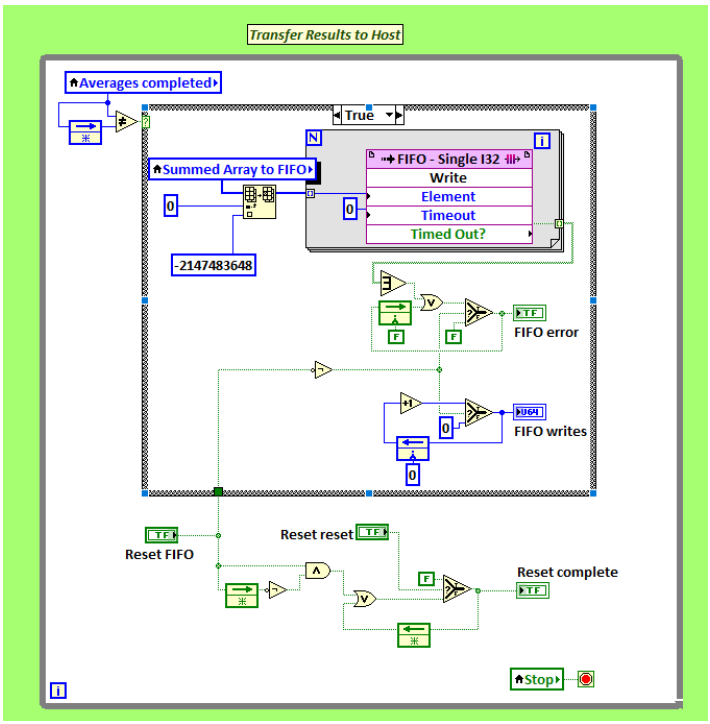

Figure 8.20: Experiment 2 - FIFO Exchange

8.3.4 Outputs

The program outputs shown on the LabView front panel are:

1. Code Length: Confirms length of Hadamard Code used - should always show as 32# Signal Loops: Counts the number of signals the ADC processes - should increase continually as the program runs
2. #Sync Loops: Counts the number of sync loops the ADC processes - should increase continually as the program runs
3. Sync Pulse Digitised: Boolean set to TRUE when the Sync Pulse ADC receives a signal that exceeds the Sync Threshold
4. Signal Direct: The transmitted power through the cavity, used for demodulation
5. Code Monitor: Shows the current values of the 32 Hadamard codes in use
6. Summed Array to FIFO: Shows the 32 values of the demodulated signal passed to the FIFO
7. Reset Complete: Indicates that the FIFO has been reset - passed to the Host program
8. FIFO writes: Counts the number of writes to the FIFO- this should match the number of averages completed
9. Averages Completed: Counts the number of averages completed
10. Sync Pulse In: Shows the raw input of the Sync Pulse ADC
11. Sync Integrator: Increases during a Sync Pulse, must exceed the Sync Pulse Width for a Sync Pulse to be counted
12. Sync Pulse Count: Counts number of Sync Pulses between resets ???
13. Code LUT Index: Combines the code length and code offset to determine which set of values from the 32 Hadamard codes is used for this demodulation
14. Yellow LSB Weight: Shows accuracy of transmitted power main signal ADC in terms of a least significant bit (I think)
15. Sync LSB Weight: Shows accuracy of transmitted power of the Sync Pulse ADCin terms of a least significant bit (I think)
16. AO0: Passes a selected output to AO0 and hence the oscilloscope - options
    - Signal Direct - the transmitted power through the cavity
    - Code LUT Index - current value for each Hadamard code being used for demodulation, cycles through 32 values
    - Summing - Boolean(TRUE) while integration is occurring - Sync Pulses should be adjusted so this comfortably between them
    - Sync Pulse In - the raw sync pulse received
    - Sync Pulse Digitised - Boolean(TRUE) while a sync pulse is present i.e. a noise free version of the raw sync pulse)
    - Sync Integrator - increases as a Sync Pulse is summed, used to check it is bigger than the defined sync pulse width
    - Sync Pulse Connect (Count?) - Counts the number of Sync Pulses between resets???
17. AO1: Passes a selected output to AO1 and hence the oscilloscope

18. AO2: Passes a selected output to AO2 and hence the oscilloscope
19. AO3: Passes a selected output to AO3 and hence the oscilloscope
20. Window 0: Shows AO0 output
21. Window 1: Shows AO1 output
22. Window 2: Shows AO2 output
23. Window 3:Shows AO3 output
24. 10MHz Counter: Counts the number of 10MHz reference pulses produced - not used in this experiment
25. 10MHz Output Waveform: Shows the 10MHz reference waveform - not used in this experiment

8.4 Appendix 4 - Samples of the Matlab Code Used for the Second Experiment

The following pages contain the Matlab modules used within LabView in the second part of the second experiment which applied the Tilt, Curve, S Curve dithers. There are three:

- The opening routine which opens the USB link to the SLMs and establishes the matrices for controlling the dithers;
- The main routine which uses to error signal to dither the SLMs and update their shapes; and
- The close routine which sets the mirrors to their final state and closes the USB link to the SLMs.

Again, comments are shown in blue, commands in black and SLM updates in red.

Tilt, Curve and SCurve Opening Routine

```
% Reinitialise Matlab and set up path
close all; % Close all figures
clear; % Clear Workspace
clc; % Clear Command Window
addpath('C:—Program Files—Boston Micromachines—Bin64—Matlab')

% Open both mirrors
serialNumber1 = '11AW173#018'; % Segmented mirror
serialNumber2 = '18AW160#016'; % Continuous mirror

[err_code1, dm1] = BMCOpenDM(serialNumber1); % Open Segmented mirror

[err_code2, dm2] = BMCOpenDM(serialNumber2); % Open Continuous mirror

% Flatten both mirrors at 50%
shapeC=50*ones (12,12);
BMCSendData(dm1, shapeC(:)/100); % Flatten the Segmented mirror
BMCSendData(dm2, shapeC(:)/100); % Flatten the Continuous mirror

% Set Up Constants
MaxTiltGap = 9.090909;
CurveMatrix=[2.7225 1.8225 1.1025 .5625 .2025 .0225 .0225 .2025 .5625 1.1025 1.8225
2.7225];
CurveMatrix=CurveMatrix/2.7225;
SegSCurveMatrix = [0.0000 0.6364 0.7937 0.8216 0.7348 0.4743 -0.4743 -0.7348 -0.8216
-0.7937 -0.6364 0.0000];
ContSCurveMatrix = [0.0000 0.7778 0.9701 1.0042 0.8981 0.5798 -0.5798 -0.8981 -1.0042
-0.9701 -0.7778 0.0000];
SegSCurveMatrix = SegSCurveMatrix/0.8216;
ContSCurveMatrix = ContSCurveMatrix/1.0042;

for j=1:12
for k=1:12
HTilt(k,j)= MaxTiltGap*(6.5-j)/100;
VTilt(j,k)= MaxTiltGap*(6.5-j)/100;
HCurve(k,j)= CurveMatrix(j)/2;
VCurve(j,k)= CurveMatrix(j)/2;
HContSCurve(j,k)=ContSCurveMatrix(j)/2;
VContSCurve(j,k)=ContSCurveMatrix(j)/2;
HSegSCurve(j,k)=ContSCurveMatrix(j)/2;
VSegSCurve(j,k)=ContSCurveMatrix(j)/2;
end
end

% Unpack driver info structure to pass out

% Segmented Mirror
type1 = double(dm1.type);
ID1 = double(dm1.ID);
size1=double(dm1.size);
%lutName1= double(dm1.lutName);
lutName1 = dm1.lutName;
priv1 = double(dm1.priv);
```

```
% Continuous Mirror
type2 = double(dm2.type);
ID2 = double(dm2.ID);
size2=double(dm2.size);
lutName2= double(dm2.lutName);
priv2 = double(dm2.priv);
```

8.4.1 Tilt, Curve and SCurve Main Routine

```
% Repack driver info structure from inputs

% Segmented mirror
dm1.type = int32(type1);
dm1.ID = int32(ID1);
dm1.size = int32(size1);
dm1.lutName = lutName1;
dm1.priv = int32(priv1);

% Continuous mirror
dm2.type = int32(type2);
dm2.ID = int32(ID2);
dm2.size = int32(size2);
dm2.lutName = lutName2;
dm2.priv = int32(priv2);

% Define new mirror shapeC and shapeS
shapeC = 50*ones (12,12) + (dTiltHorzC*HTilt + dTiltVertC*VTilt +dCurveHorzC*
HCurve + dCurveVertC*VCurve+dHorzContSCurve*HContSCurve+dVertContSCurve*VContSCurve);
shapeS = 50*ones (12,12) + (dTiltHorzS*HTilt + dVertTiltS*VTilt +dCurveHorzS*
HCurve + dCurveVertS*VCurve+dHorzSegSCurve*HSegSCurve+dVertSegSCurve*VSegSCurve);
i=1; % Resets i and j in case a loop is stopped halfway through.
j=1;

% Modulate HTilt, VTilt, HCurve, VCurve, HSCurve and VSCurve simultaneaously

% Update code and establish modulation matrices
codetable = hadamard(CodeLength);
modHTC = zeros(12,12);
modVTC = zeros(12,12);
modHCC = zeros(12,12);
modVCC = zeros(12,12);
modHTS = zeros(12,12);
modVTS = zeros(12,12);
modHCS = zeros(12,12);
modVCS = zeros(12,12);
modHSCurveC = zeros(12,12);
modVSCurveC = zeros(12,12);
modHSCurveS = zeros(12,12);
modVSCurveS = zeros(12,12);


% Modulate both SLMs
for j = 1:NoCodes;
for i = 1:32;
modHTC =codetable(2,i) * moddepth * HTilt;
modVTC = codetable(3,i) * moddepth * VTilt;
modHCC = codetable(4,i)* moddepth * HCurve;
modVCC = codetable(5,i) * moddepth * VCurve;
modHSCurveC = codetable(10,i) * moddepth * HContSCurve;
modVSCurveC = codetable(11,i) * moddepth * VContSCurve;
codeC = shapeC + modHTC+modVTC+modHCC+modVCC+modHSCurveC+modVSCurveC;
modHTS = codetable(6,i) * moddepth * HTilt;
```

```
modVTS = codetable(7,i) * moddepth * VTilt;
modHCS = codetable(8,i)* moddepth * HCurve;
modVCS = codetable(9,i) * moddepth * VCurve;
modHSCurveS = codetable(12,i) * moddepth * HSegSCurve;
modVSCurveS = codetable(13,i) * moddepth * VSegSCurve;
codeS = shapeS + modHTS+modVTS+modHCS+modVCS+modHSCurveS+modVSCurveS;
BMCSendData(dm2, codeC(:)/100);
BMCSendData(dm1, codeS(:)/100);
for k = 1:wait_time;
end;
end;
end;
```

8.4.2 Tilt, Curve and SCurve Close Routine

```
% Repack driver info structure from inputs

% Segmented Mirror
dm1.type = int32(type1);
dm1.ID = int32(ID1);
dm1.size = int32(size1);
dm1.lutName = int32(lutName1);
dm1.priv = int32(priv1);

% Continuous Mirror
dm2.type = int32(type2);
dm2.ID = int32(ID2);
dm2.size = int32(size2);
dm2.lutName = int32(lutName2);
dm2.priv = int32(priv2);_ID);

% Leave Segemented Mirror in optimised shapeS and close
code=50*ones (12,12);
BMCSendData(dm1, shapeC(:)/100);
BMCCloseDM(dm1);

% Leave Continuous Mirror in optimised shapeC and close
BMCSendData(dm2, shapeC(:)/100);
BMCCloseDM(dm2);
```

References

[1] T. H. Maiman, "Stimulated optical radiation in ruby," *Nature*, vol. 187, no. 4736, pp. 493–494, 1960. [Online]. Available: ⟨GotoISI⟩://WOS:A1960ZQ06900019

[2] Empedocles, *The Poem of Empedocles*, vol. CTXT-85.

[3] Euclid, *Optica*, 300BC.

[4] Lucretius, *De Rerum Natura (Nature)*, 55BC.

[5] I. al Haytham, *Book of Optics*, 1020.

[6] 2015. [Online]. Available: http://www.ibnalhaytham.com/discover/who-was-ibn-al-haytham/

[7] 2017. [Online]. Available: https://www.britannica.com/biography/Ibn-al-Haytham

[8] T. S. W. o. P. Gassendi, *Brash, C. B*, 1972.

[9] Newton, "Opticks: or a treatise of the reflections, refractions, inflections and colours of light," p. 68, 1730.

[10] R. Hooke, "Robert hooke's critique of newton's theory of light and colors," *The History of the Royal Society*, vol. 3, pp. 10–15, 1672.

[11] Huygens, "Traité de la lumière (treaty of the light)," *Project Gutenberg Edition*, 1678.

[12] D. Fleisch, *A Student's Guide to Maxwell's Equations*. Wittenberg University: Cambridge University Press, 2008.

[13] L. A. Group, "Advanced ligo," vol. LIGO-P1400177-v5, 2014.

[14] S. Robert, B. Brian, V. Glenn de, D. Jeffrey, K. William, O. Tetsuo, M. Kirk, S. Daniel, R. David, S. Andrew, and W. Brent, "Progress in interferometry for lisa at jpl," *Classical and Quantum Gravity*, vol. 28, no. 9, p. 094007, 2011. [Online]. Available: http://stacks.iop.org/0264-9381/28/i=9/a=094007

[15] A. F. Brooks, T.-L. Kelly, P. J. Veitch, and J. Munch, "Ultra-sensitive wavefront measurement using a hartmann sensor," *Optics Express*, vol. 15, no. 16, pp. 10 370–10 375, 2007. [Online]. Available: ⟨GotoISI⟩://WOS:000248753100050

[16] T. L. Kelly, P. J. Veitch, A. F. Brooks, and J. Munch, "Accurate and precise optical testing with a differential hartmann wavefront sensor," *Appl Opt*, vol. 46, no. 6, pp. 861–6, 2007. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/17279130

[17] D. Tarquin Ralph, P. A. Altin, D. S. Rabeling, D. E. McClelland, and D. A. Shaddock, "Interferometric wavefront sensing with a single diode using spatial light modulation," *Appl Opt*, vol. 56, no. 8, pp. 2353–2358, 2017. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/28375282

[18] S. A. E, *Lasers*. Mill Valley, California: University Science Books, 1986.

[19] G. Goubau and F. Schwering, "On the guided propagation of electromagnetic wave beams," *IRE Transactions on Antennas and Propagation*, vol. 9, no. 3, pp. 248–256, 1961.

[20] L. W. Davis, "Theory of electromagnetic beams," *Physical Review A*, vol. 19, no. 3, pp. 1177–1179, 1979. [Online]. Available: ⟨GotoISI⟩://WOS:A1979GT83200030

[21] M. Lax, W. H. Louisell, and W. B. McKnight, "From maxwell to paraxial wave optics," *Physical Review A*, vol. 11, no. 4, pp. 1365–1370, 1975. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.11.1365

[22] C. G. Chen, P. T. Konkola, J. Ferrera, R. K. Heilmann, and M. L. Schattenburg, "Analyses of vector gaussian beam propagation and the validity of paraxial and spherical approximations," *J Opt Soc Am A Opt Image Sci Vis*, vol. 19, no. 2, pp. 404–12, 2002. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/11822605

[23] O. Svelto, *Principles of Lasers*, 5th ed., 2010.

[24] E. Toombs, "Intensity of a gaussian beam," 2011. [Online]. Available: https://creativecommons.org/licenses/by/3.0/

[25] A. Ourjoumtsev, "Beam tools," 2003.

[26] E. Morrison, B. Meers, D. Robertson, and H. Ward, "Automatic alignment of optical interferometers," *Applied Optics*.

[27] E. Morrison, B. J. Meers, D. I. Robertson, and H. Ward, "Experimental demonstration of an automatic alignment system for optical interferometers," *Appl Opt*, vol. 33, no. 22, pp. 5037–40, 1994. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/20935884

[28] R. Fessenden, "Wireless signalling," *US Patent*, 1902.

[29] G. Heinzel, V. Wand, A. Garc´ıa, F. Guzm´an, F. Steier, C. Killow, D. Robertson, H. Ward, and C. Braxmaier, "Investigation of noise sources in the ltp interferometer," Report, 21 October 2005 2005.

[30] K. Meehan, "Characterizing the temperature sensitivity of the hartmann sensor," 2010.

[31] N. Patel and C. S. Narayanamurthy, "Measurement of optical misalignment using wavefront sensing," *Optical Engineering*, vol. 54, no. 10, 2015.

[32] A. F. Brooks, B. Abbott, M. A. Arain, G. Ciani, A. Cole, G. Grabeel, E. Gustafson, C. Guido, M. Heintze, A. Heptonstall, M. Jacobson, W. Kim, E. King, A. Lynch, S. O'Connor, D. Ottaway, K. Mailand, G. Mueller, J. Munch, V. Sannibale, Z. Shao, M. Smith, P. Veitch, T. Vo, C. Vorvick, and P. Willems, "Overview of advanced ligo adaptive optics," *Appl Opt*, vol. 55, no. 29, pp. 8256–8265, 2016. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/27828071

[33] P. Veitch, A. Brooks, W. Kim, C. Blair, H. Cao, G. Grabeel, T. Hardwick, M. Heintze, A. Heponstall, C. Ingram, J. Munch, D. Ottaway, and T. Vo, "Hartmann wavefront sensors for advanced ligo," in *Conference on Lasers and*

Electro-Optics, ser. OSA Technical Digest (online). Optical Society of America, Conference Proceedings, p. SW3M.5. [Online]. Available: http://www.osapublishing.org/abstract.cfm?URI=CLEO_SI-2018-SW3M.5

[34] B. C. Platt and R. Shack, "History and principles of shack-hartmann wavefront sensing," *J Refract Surg*, vol. 17, no. 5, pp. S573–7, 2001. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/11583233

[35] [Online]. Available: https://www.optics.arizona.edu/sites/optics.arizona.edu/files/pdf/Historical-Development-Shack-Hartman-Wavefront-Sensor.pdf

[36] 2018. [Online]. Available: http://www.bostonmicromachines.com/wavefront-sensor.html

[37] A. A. Goloborodko, M. M. Kotov, V. N. Kurashov, and D. V. Podanchuk, "Aberration measurements in speckle field by shack–hartmann wavefront sensor," *Optical Engineering*, vol. 55, no. 12, 2016. [Online]. Available: ⟨GotoISI⟩://WOS:000396372900012

[38] D. Blair, L. Ju, C. Zhao, L. Wen, H. Miao, R. Cai, J. Gao, X. Lin, D. Liu, L.-A. Wu, Z. Zhu, G. Hammond, H. J. Paik, V. Fafone, A. Rocchi, C. Blair, Y. Ma, J. Qin, and M. Page, "The next detectors for gravitational wave astronomy," *Science China Physics, Mechanics and Astronomy*, vol. 58, no. 12, p. 120405, 2015. [Online]. Available: https://doi.org/10.1007/s11433-015-5747-7

[39] D. Blair, L. Ju, C. Zhao, L. Wen, Q. Chu, Q. Fang, R. Cai, J. Gao, X. Lin, D. Liu, L.-A. Wu, Z. Zhu, D. H. Reitze, K. Arai, F. Zhang, R. Flaminio, X. Zhu, G. Hobbs, R. N. Manchester, R. M. Shannon, C. Baccigalupi, W. Gao, P. Xu, X. Bian, Z. Cao, Z. Chang, P. Dong, X. Gong, S. Huang, P. Ju, Z. Luo, L. Qiang, W. Tang, X. Wan, Y. Wang, S. Xu, Y. Zang, H. Zhang, Y.-K. Lau, and W.-T. Ni, "Gravitational wave astronomy: the current status," *Science China Physics, Mechanics and Astronomy*, vol. 58, no. 12, p. 120402, 2015. [Online]. Available: ⟨GotoISI⟩://WOS:000365817400003

[40] M. Evans, S. Gras, P. Fritschel, J. Miller, L. Barsotti, D. Martynov, A. Brooks, D. Coyne, R. Abbott, R. X. Adhikari, K. Arai, R. Bork, B. Kells, J. Rollins, N. Smith-Lefebvre, G. Vajente, H. Yamamoto, C. Adams, S. Aston, J. Betzweiser, V. Frolov, A. Mullavey, A. Pele, J. Romie, M. Thomas, K. Thorne, S. Dwyer, K. Izumi, K. Kawabe, D. Sigg, R. Derosa, A. Effler, K. Kokeyama, S. Ballmer, T. J. Massinger, A. Staley, M. Heinze, C. Mueller, H. Grote, R. Ward, E. King, D. Blair, L. Ju, and C. Zhao, "Observation of parametric instability in advanced ligo," *Phys Rev Lett*, vol. 114, no. 16, p. 161102, 2015. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/25955042

[41] L. Ju, M. Aoun, P. Barriga, D. G. Blair, A. Brooks, R. Burman, R. Burston, X. T. Chin, E. J. Chin, C. Y. Lee, D. Coward, B. Cusack, G. d. Vine, J. Degallaix, J. C. Dumas, F. Garoi, S. Gras, M. Gray, D. Hosken, E. Howell, J. S. Jacob, T. L. Kelly, B. Lee, K. T. Lee, T. Lun, D. McClelland, C. Mow-Lowry, D. Mudge, J. Munch, S. Schediwy, S. Scott, A. Searle, B. Sheard, B. Slagmolen, P. Veitch, J. Winterflood, A. Woolley, Z. Yan, and C. Zhao, "Aciga's high optical power test facility," *Classical and Quantum Gravity*, vol. 21, no. 5, pp. S887–S893, 2004. [Online]. Available: ⟨GotoISI⟩://WOS:000220344800078

[42] M. Evans, S. Gras, P. Fritschel, J. Miller, L. Barsotti, D. Martynov, A. Brooks, D. Coyne, R. Abbott, R. X. Adhikari, K. Arai, R. Bork, B. Kells, J. Rollins,

N. Smith-Lefebvre, G. Vajente, H. Yamamoto, C. Adams, S. Aston, J. Betzweiser, V. Frolov, A. Mullavey, A. Pele, J. Romie, M. Thomas, K. Thorne, S. Dwyer, K. Izumi, K. Kawabe, D. Sigg, R. Derosa, A. Effler, K. Kokeyama, S. Ballmer, T. J. Massinger, A. Staley, M. Heinze, C. Mueller, H. Grote, R. Ward, E. King, D. Blair, L. Ju, and C. Zhao, "Observation of parametric instability in advanced ligo," *Physical Review Letters*, vol. 114, no. 16, p. 161102, 2015. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.114.161102

[43] D. A. Shaddock, "Digitally enhanced heterodyne interferometry," *Optics Letters*, vol. 32, no. 22, pp. 3355–3357, 2007. [Online]. Available: http://ol.osa.org/abstract.cfm?URI=ol-32-22-3355

[44] S. Raj, "Matched template signal processing of continuous wave laser for ranging to space debris," Thesis, 2018.

[45] T. Helleseth, "Sequences with good correlations and some open problems," *Electronic Notes in Discrete Mathematics*, vol. 6, pp. 507–517, 2001. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1571065304002033

[46] T. Helleseth and A. Kholosha, "Crosscorrelation of m-sequences, exponential sums, bent functions and jacobsthal sums," *Cryptography and Communications*, vol. 3, no. 4, pp. 281–291, 2011. [Online]. Available: ⟨GotoISI⟩://WOS:000208850600008

[47] "Linear feedback shift registers," 2013. [Online]. Available: http://in.ncu.edu.tw/ncume_ee/digilogi/prbs.htm

[48] A. Canteaut, P. Charpin, and H. Dobbertin, "Binary m-sequences with three-valued crosscorrelation: a proof of welch's conjecture," *IEEE Transactions on Information Theory*, vol. 46, no. 1, pp. 4–8, 2000. [Online]. Available: ⟨GotoISI⟩://WOS:000084870900002

[49] J. Seberry and M. Yamada, *Hadamard matrices, Sequences, and Block Designs*, ser. Contemporary Design Theory – A Collection of Surveys. John Wiley and Sons, 1992.

[50] L. D. Baumert and M. Hall Jr, "A new construction for hadamard matrices," *Bulletin of the American Mathematical Society*, vol. 71, no. 1, pp. 169–171, 1965. [Online]. Available: ⟨GotoISI⟩://WOS:A19656313300016

[51] N. I. Gillespie and C. E. Praeger, "Uniqueness of certain completely regular hadamard codes," *Journal of Combinatorial Theory, Series A*, vol. 120, no. 7, pp. 1394–1400, 2013. [Online]. Available: ⟨GotoISI⟩://WOS:000323868200003

[52] D. H. Smith, R. P. Ward, and S. Perkins, "Gold codes, hadamard partitions and the security of cdma systems," *Designs, Codes and Cryptography*, vol. 51, no. 3, pp. 231–243, 2008. [Online]. Available: https://doi.org/10.1007/s10623-008-9257-8

[53] W. C. Y. Lee, "Overview of cellular cdma," *IEEE Transactions on Vehicular Technology*, vol. 40, no. 2, pp. 291–302, 1991. [Online]. Available: ⟨GotoISI⟩://WOS:A1991FN86500003

[54] P. Altin, "Wavefront sensing with a single-element detector," 2016.

[55] L. Barsotti, M. Evans, and P. Fritschel, "Alignment sensing and control in advanced ligo," *Classical and Quantum Gravity*, vol. 27, no. 8, p. 084026, 2010. [Online]. Available: ⟨GotoISI⟩://WOS:000282154300027

[56] E. D. Black, "Notes on the pound-drever-hall technique," 2014.

[57] C. Bond, D. Brown, A. Freise, and K. Strain, "Interferometer techniques for gravitational-wave detection," *Living Reviews in Relativity*, vol. 13, no. 1, pp. 1–207, 2015.

[58] R. Ward, "Pdh locking," 2017.

[59] A. E. Siegman, "Laser beams and resonators: Beyond the 1960s," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 6, no. 6, pp. 1389–1399, 2000. [Online]. Available: ⟨GotoISI⟩://WOS:000167275600063

[60] O. O. Courses, "Optique pour l'ingénieur (optics for the engineer)," 2009.

[61] H. Wang, M. Alvarez, C. Collins, D. Brown, M. Wang, C. Mow-Lowry, and A. Freise, "Experimental optical characterisation of geometrically near-unstable cavity," University of Birmingham, Report, 2 October 2017 2017.

[62] H. Wang, M. Dovale-Álvarez, C. Collins, D. D. Brown, M. Wang, C. M. Mow-Lowry, S. Han, and A. Freise, "Feasibility of near-unstable cavities for future gravitational wave detectors," *Physical Review D*, vol. 97, no. 2, p. 022001, 2018. [Online]. Available: ⟨GotoISI⟩://WOS:000423126100002

[63] C. J. Hood, H. J. Kimble, and J. Ye, "Characterization of high-finesse mirrors: Loss, phase shifts, and mode structure in an optical cavity," *Physical Review A*, vol. 64, no. 3, pp. 1–7, 2001. [Online]. Available: ⟨GotoISI⟩://WOS:000170978600110

[64] M. Suter and P. Dietiker, "Calculation of the finesse of an ideal fabry-perot resonator," *Appl Opt*, vol. 53, no. 30, pp. 7004–10, 2014. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/25402787

[65] C. Mateo-Segura, A. P. Feresidis, and G. Goussetis, "Analysis of broadband highly-directive fabry-perot cavity leaky-wave antennas with two periodic layers," in *2010 IEEE Antennas and Propagation Society International Symposium*, Conference Proceedings, pp. 1–4.

[66] B. Klein, "Cavity q and finesse," 2018. [Online]. Available: http://bklein.ece.gatech.edu/laser-photonics/cavity-q-finesse-and-photon-lifetime/

[67] A. K. Ghatak, *Optical Electronics*. New York: New York Cambridge University Press, 1989.

[68] D. Jena, "Lasers and optoelectronics lecture 10," 2017.

[69] N. Ismail, C. C. Kores, D. Geskus, and M. Pollnau, "Fabry-perot resonator: spectral line shapes, generic and related airy distributions, linewidths, finesses, and performance at low or frequency-dependent reflectivity," *Opt Express*, vol. 24, no. 15, pp. 16 366–89, 2016. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/27464090

[70] J. Stone and D. Marcuse, "Ultrahigh finesse fiber fabry-perot interferometers," *Journal of Lightwave Technology*, vol. 4, no. 4, pp. 382–385, 1986. [Online]. Available: ⟨GotoISI⟩://WOS:A1986A815500002

[71] E. D. Black, "An introduction to pound–drever–hall laser frequency stabilization," *American Journal of Physics*, vol. 69, no. 1, pp. 79–87, 2001. [Online]. Available: ⟨GotoISI⟩://WOS:000166018200011

[72] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator," *Applied Physics B Photophysics and Laser Chemistry*, vol. 31, no. 2, pp. 97–105, 1983. [Online]. Available: ⟨GotoISI⟩://WOS:A1983QU19000008

[73] M. Nickerson, "A review of pound-drever-hall laser frequency locking."

[74] D. A. Shaddock, M. B. Gray, and D. E. McClelland, "Frequency locking a laser to an optical cavity by use of spatial mode interference," *Optics Letters*, vol. 24, no. 21, pp. 1499–1501, 1999. [Online]. Available: ⟨GotoISI⟩://WOS:000083598000019

[75] T. W. Hansch and B. Couillaud, "Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity," *Optics Communications*, vol. 35, no. 3, pp. 441–444, 1980. [Online]. Available: ⟨GotoISI⟩://WOS:A1980KW09600032

[76] J. E. Bateman, R. L. D. Murray, M. Himsworth, H. Ohadi, A. Xuereb, and T. Freegarde, "Hänsch–couillaud locking of mach–zehnder interferometer for carrier removal from a phase-modulated optical spectrum," *Journal of the Optical Society of America B*, vol. 27, no. 8, pp. 1530–1533, 2010. [Online]. Available: ⟨GotoISI⟩://WOS:000280665300005

[77] R. Arieli, *The Laser Adventure*. Rehovot, ISRAEL: Weizmann Institute of Science, 1997.

[78] M. A. Bandres and J. C. Gutiérrez-Vega, "Ince–gaussian beams," *Optics Letters*, vol. 29, no. 2, pp. 144–146, 2004. [Online]. Available: http://ol.osa.org/abstract.cfm?URI=ol-29-2-144

[79] S. Saghafi and C. J. R. Sheppard, "Near field and far field of elegant hermite-gaussian and laguerre-gaussian modes," *Journal of Modern Optics*, vol. 45, no. 10, pp. 1999–2009, 1998. [Online]. Available: https://doi.org/10.1080/09500349808231738

[80] I. Kimel and L. R. Elias, "Relations between hermite and laguerre gaussian modes," *IEEE Journal of Quantum Electronics*, vol. 29, no. 9, pp. 2562–2567, 1993. [Online]. Available: ⟨GotoISI⟩://WOS:A1993MM21500017

[81] M. VanValkenburgh, "Laguerre–gaussian modes and the wigner transform," *Journal of Modern Optics*, vol. 55, no. 21, pp. 3537–3549, 2008. [Online]. Available: ⟨GotoISI⟩://WOS:000261803400006

[82] D. Z. Anderson, "Alignment of resonant optical cavities," *Appl Opt*, vol. 23, no. 17, p. 2944, 1984. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/18213100

[83] G. Heinzel, J. Bogenstahl, C. Braxmaier, K. Danzmann, A. Garcia, F. Guzman, J. Hough, D. Hoyland, O. Jennrich, C. Killow, D. Robertson, Z. Sodnik, F. Steier, H. Ward, and V. Wand, "Interferometry for the lisa technology package ltp: an update," *Journal of Physics: Conference Series*, vol. 32, no. 1, p. 132, 2006. [Online]. Available: http://stacks.iop.org/1742-6596/32/i=1/a=021

[84] G. Heinzel, C. Braxmaier, R. Schilling, A. Rüdiger, D. Robertson, M. t. Plate, V. Wand, K. Arai, U. Johann, and K. Danzmann, "Interferometry for the lisa technology package (ltp) aboard smart-2," *Classical and Quantum Gravity*, vol. 20,

no. 10, p. S153, 2003. [Online]. Available: http://stacks.iop.org/0264-9381/20/i=10/a=318

[85] Y. Wang, "On inter-satellite laser ranging, clock synchronization and gravitational wave data analysis," Thesis.

[86] K. Danzmann, "Laser interferometer space antenna," Albert Einsteint Institutue Hannover, Report, 2017.

[87] X. Yu, S. Gillmer, and J. D. Ellis, "Beam geometry, alignment, and wavefront aberration effects on interferometric differential wavefront sensing," *Measurement Science and Technology*, vol. 26, no. 125203, 2015.

[88] J. Pan, C. Che, Y. Zhu, and M. Wang, "Satellite jitter estimation and validation using parallax images," *Sensors*, vol. 17, no. 1, p. 83, 2017. [Online]. Available: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC5298656/

[89] D. Schutze, G. Stede, V. Muller, O. Gerberding, T. Bandikova, B. S. Sheard, G. Heinzel, and K. Danzmann, "Laser beam steering for grace follow-on intersatellite interferometry," *Optics Express*, vol. 22, no. 20, 2014.

[90] A. F. Brooks, D. Hosken, J. Munch, P. J. Veitch, Z. Yan, C. Zhao, Y. Fan, L. Ju, D. Blair, P. Willems, B. Slagmolen, and J. Degallaix, "Direct measurement of absorption-induced wavefront distortion in high optical power systems," *Appl Opt*, vol. 48, no. 2, pp. 355–64, 2009. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/19137047

[91] A. Brooks, "Hartmann wavefront sensors for advanced gravitational wave interferometers," Thesis, 2007.

[92] A. Brooks, P. Veitch, J. Munch, and T.-L. Kelly, "An off-axis hartmann sensor for the measurement of absorption-induced wavefront distortion in advanced gravitational wave interferometers," *General Relativity and Gravitation*, vol. 37, no. 9, pp. 1575–1580, 2005. [Online]. Available: ⟨GotoISI⟩://WOS:000232616300009

[93] B. S. Sheard, G. Heinzel, K. Danzmann, D. A. Shaddock, W. M. Klipstein, and W. M. Folkner, "Intersatellite laser ranging instrument for the grace follow-on mission," *Journal of Geodesy*, vol. 86, no. 12, pp. 1083–1095, 2012. [Online]. Available: ⟨GotoISI⟩://WOS:000310964800001

[94] D. Marina, C. Felipe Guzmán, S. Benjamin, H. Gerhard, and D. Karsten, "Laser interferometer for spaceborne mapping of the earth's gravity field," *Journal of Physics: Conference Series*, vol. 154, no. 1, p. 012023, 2009. [Online]. Available: http://stacks.iop.org/1742-6596/154/i=1/a=012023

[95] G. Heinzel, A. Rudiger, R. Schilling, K. Strain, W. Winkler, J. Mizuno, and D. K, "Automatic beam alignment in the garching 30-m prototype of a laser-interferometric gravitational wave detector," *Optics Communications*, vol. 160, pp. 321–334, 1999.

[96] D. Robertson, C. Killow, H. Ward, J. Hough, G. Heinzel, A. Garcia, V. Wand, U. Johann, and C. Braxmaier, "Ltp interferometer—noise sources and performance," *Classical and Quantum Gravity*, vol. 22, no. 10, pp. S155–S163, 2005. [Online]. Available: ⟨GotoISI⟩://WOS:000230230100005

[97] B. Micromachines, "Multi-dm deformable mirror system - user manual v3.7," 2014.

[98] K. Lutz, "Multi-dm specification sheet," 2014.

[99] Boston-MicroMachines, "X-driver," 2017.

[100] 2017. [Online]. Available: http://www.ntp.org/ntpfaq/NTP-s-sw-clocks-quality.htm

[101] A. Romann, D. S. Rabeling, G. de Vine, J. H. Chow, M. B. Gray, and D. E. McClelland, "Automatic alignment of a rigid spacer cavity," *General Relativity and Gravitation*, vol. 37, no. 9, pp. 1591–1599, 2005. [Online]. Available: ⟨GotoISI⟩://WOS:000232616300011

[102] T. C. Briles, D. C. Yost, A. Cingoz, J. Ye, and T. R. Schibli, "Simple piezoelectric-actuated mirror with 180 khz servo bandwidth," *Opt Express*, vol. 18, no. 10, pp. 9739–46, 2010. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/20588824

[103] 2018. [Online]. Available: https://www.edmundoptics.com/resources/application-notes/imaging/distortion/

[104] T. U. o. Denmark, "Higher order modes," *Department of Photnics Engineering*, 2018. [Online]. Available: http://www.fotonik.dtu.dk/english/Research/NonlinearOpticsBiophotonics/FOD/Research-Topics/Higher-order-Modes

[105] Y. Baykal, "Higher order mode laser beam scintillations in oceanic medium," *Waves in Random and Complex Media*, vol. 26, no. 1, pp. 21–29, 2015. [Online]. Available: https://doi.org/10.1080/17455030.2015.1099760

[106] X. Peng, K. Kim, M. Mielke, T. Booth, J. W. Nicholson, J. M. Fini, X. Liu, A. Desantolo, P. S. Westbrook, R. S. Windeler, E. M. Monberg, F. V. Dimarcello, C. Headley, and D. J. Digiovanni, "Higher-order mode fiber enables high energy chirped-pulse amplification," *Opt Express*, vol. 21, no. 26, pp. 32 411–6, 2013. [Online]. Available: https://www.ncbi.nlm.nih.gov/pubmed/24514834

[107] L. S. Collaboration, "Instrument science white paper," LIGO Scientific Collaboration, Report, 2017.

[108] M. Feinberg, "Power limitations on slms - email:19 oct 2017," 2017.

[109] N. Smith-Lefebvre and N. Mavalvala, "Modematching feedback control for interferometers with an output mode cleaner,," *LIGO SCientific Collaboration*, vol. LIGO-P1200034, 2012.

[110] L. Roberts, "Internally sensed optical phased arrays," Thesis, 2016.