Internally Sensed Optical Phased Arrays

Lyle Edward Roberts

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

Submitted
5th July, 2016

To my family,
and friends.

Declaration

This thesis is an account of research undertaken between February 2012 and March 2016 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.

Lyle Edward Roberts
July 2016

Acknowledgements

It is difficult to express my immense gratitude for my supervisors, Professor Daniel Shaddock, Dr Robert Ward, and Professor David McClelland. Without you, none of the work presented in this thesis would have been possible.

Daniel, it has been an incredible privilege to be your student. For five years you have been my supervisor, my mentor, but more importantly my *friend*. I have learnt so much from you, and I can't thank you enough for it.

Rob, you are an extraordinary scientist. You have always encouraged me to challenge myself, be creative, and to persevere no matter how challenging the situation. I owe you so much. I will forever be in awe of your encyclopaedic knowledge and prodigious intuition of optics, control systems, gravitational waves, and pretty much everything else, really. I have thoroughly enjoyed my time working with you, and I look forward to continuing to work with you into the future.

David, the success of the Centre for Gravitational Physics is a testament to your exuberant personality and remarkable academic prowess. I am sincerely grateful for your continued support and guidance over the years.

One of the best parts of working at the Centre for Gravitational Physics is being exposed to exceptional researchers such as Bram #Swagmolen, Jong Chow, Roland Fleddermann, and Terry McRae. Thank you for always being willing to help when I need it.

My time as a PhD student has been thoroughly enjoyable thanks to the incredible support and endless supply of entertainment provided by the graviteers. I thank Sam, Silvie, Jarrod, Georgia, Shasi, Paul, Thanh, Roland, Tarquin, Cathy, Chatters, Andrew Wade, David McManus, Perry, and Ben Courtney-Barrer for their friendship and encouragement. I am also grateful for the guidance that veteran graviteers Danielle Wuchenich, Andrew Sutton, Dave Bowman, and Timothy Lam have provided me over the years. I couldn't ask for a better team of people to work with.

In 2012, I was fortunate enough to participate in the NASA Ames Academy for Space Exploration. During my time there I forged friendships with some of the most incredible space cadets I will ever meet. Ryan, Michael, Anthony, Tyler, Steph, Alec, George, Alana, Doug and Christie, Mark, Arlen, Chris, Trish, Swarna, Laura, Takako, Ty, and Eleni; I am so grateful for your love and motivation over the years. Thank you for inspiring me to achieve greatness, and for introducing me to In-N-Out burger.

I also extend my sincerest gratitude to Kyle Hardman, Dr Carlos Kuhn, and Dr Nicholas Robins for their help setting up the new high-power optical phased array lab. Thank you for helping me construct the "clean tent" and letting me *borrow*[1] an optical bench.

I would like to thank my family for all their love and encouragement. I could not have done this without you. I thank my parents, my grandparents, and my brothers for always believing in me, and for raising me with such an enthusiasm for science and engineering.

Most of all, I thank my loving and patient wife, Nadine. You are the light in my life.

[1] You're not getting it back.

Abstract

The performance of existing ground based space debris laser ranging systems can be improved by directing more light onto space debris by coherently combining multiple lasers using an optical phased array (OPA). If the power delivered to target is sufficiently high then these systems may also provide the capability to remotely manoeuvre space debris via photon radiation pressure and/or ablation. By stabilising the relative output phase of multiple lasers, OPAs form a coherent optical wave-front in the far field. Since the phase of each laser can be controlled independently, they also have the ability to dynamically manipulate the distribution of optical power in the far field, potentially enabling them to compensate for atmospheric turbulence. This beam-forming functionality, combined with their inherent scalability and high power handling capabilities make OPAs a promising technology for future space debris laser ranging and manoeuvring systems.

In this thesis we describe the iterative development of a high-power compatible internally sensed OPA, which—in contrast to externally sensed OPAs that sense the output phase of each laser externally using free-space optics—relies on the small fraction of light that is reflected back into the fibre at the output of the OPA to stabilise its relative output phase. This allows internally sensed OPAs to be implemented entirely within fibre without any dependence on free-space optics at the output, offering potential advantages over externally sensed techniques when operating in the presence of shock and vibration.

A proof-of-concept experiment demonstrated the viability of internal sensing, but also highlighted a number of weaknesses that would affect its utility, specifically in supporting high optical powers greater than 100s of mW. An improved high-power compatible internally sensed OPA was designed to overcome these restrictions by isolating sensitive optical components from high optical powers using asymmetric fibre couplers. This concept was initially demonstrated experimentally using slave lasers offset phase-locked to a single master laser, and then again using fibre amplifiers in a master oscillator power amplifier configuration. The experimental demonstration of the fibre amplifier compatible OPA stabilised the relative output phase of three commercial 15 W fibre amplifiers, demonstrating a root-mean-squared output phase stability of $\lambda/194$, and the ability to steer the beam at up to 10 kHz.

The internally sensed OPA presented here requires the simultaneous measurement, and control of the phase of each emitter in the OPA. This is accomplished using digitally enhanced heterodyne interferometry and digitally implemented phasemeters, both of which rely heavily on high-speed digital signal processing resources provided by field-programmable gate-arrays.

Contents

List of Figures

List of Abbreviations

ADC	Analogue-to-Digital Converter
ANU	Australian National University
AOM	Acousto-Optic Modulator
BS	Beamsplitter
CBC	Coherent Beam Combining
CDF	Cumulative Density Function
CIC	Cascade Integrator-Comb
CPU	Central Processing Unit
CW	Continuous Wave
DAC	Digital-to-Analogue Converter
DAM	Debris Avoidance Manoeuvre
DEHI	Digitally Enhanced Heterodyne 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
G&H	Gooch & Housego
GPS	Global Positioning System
IQ	In-Phase and Quadrature
LFSR	Linear Feedback Shift Register
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
OOL	Out-of-Loop
OPA	Optical Phased Array
PBS	Polarising Beamsplitter
PCIe	Peripheral Component Interface Express
PDF	Probability Density Function
PLL	Phase-Locked Loop
PM	Polarisation Maintaining
PRN	Pseudo-Random Noise
PSD	Power Spectral Density
PZT	Piezoelectric Transducer
RF	Radio Frequency
RMS	Root Mean Square
ROM	Read Only Memory
RP	Return Path
RPSD	Root Power Spectral Density
SBC	Spectral Beam Combining
SBS	Stimulated Brillouin Scattering
SNR	Signal-to-Noise Ratio
SPGD	Stochastic Parallel Gradient Descent
SSN	Space Surveillance Network
TF	Transfer Function
VHDL	VHSIC Hardware Description Language
VHSIC	Very High-Speed Integrated Circuit
ZCD	Zero-Crossing Detector

Introduction

1.1 Orbital space debris

Orbital space debris[1] is a growing threat to the near-Earth space environment, which is currently home to 1,261 operational satellites, many of which we rely upon every day including Earth-observing weather satellites and the global positioning system [1]. Figure 1.1 shows the total number of man-made Earth orbiting objects currently tracked by the United States Space Surveillance Network (U.S. SSN), categorised by object type. The figure highlights the three dominant mechanisms that influence the total orbital debris population: i) the consistency with which new satellites are launched into orbit; ii) the occurrence of major debris-debris, debris-satellite, and satellite-satellite collisions; and iii) the natural orbital decay of objects due to atmospheric drag. As the debris population continues to rise, so too will the likelihood of subsequent collisions, resulting in more collisions and the generation of more debris. The danger of this kind of situation is that the density of space debris may eventually become so great that a single collision could trigger an irreversible runaway chain reaction of collisions in a scenario described by the Kessler syndrome [2]. An event of this type would be devastating, likely rendering specific high-density orbital bands in low-Earth orbit unusable for hundreds of years.

On January 11 2007, a Chinese anti-satellite missile test was performed to de-orbit the Chinese weather satellite Fengyun-1C, producing 2,259 trackable debris objects, and approximately 200,000 objects too small to be directly monitored [4]. Two years later on February 10 2009, an operational Iridium-33 communications satellite accidentally collided with a defunct Cosmos-2251 communications satellite, creating over 2,000 new debris objects in a region of relatively high spatial density, significantly increasing the likelihood of subsequent debris-debris and debris-satellite collisions [5, 6]. As of 2016, the ISS has performed two evasive manoeuvres to avoid the close approach of debris from the 2009 collision of Iridium-33 and Cosmos-2251, the first in late March 2011, and most recently on the 27$^{\text{th}}$ of October 2014[2]. The debris created by the intentional destruction of the Fengyun-1C weather satellite and accidental Iridium-Cosmos collision currently represent one third of all catalogued objects.

[1] Orbital space debris is classified as any object in Earth's orbit that serves no useful purpose, including defunct satellites, discarded rocket stages, and mission related debris (e.g., flecks of paint, explosive bolt remnants, and uncontained rocket fuel).

[2] In 2014 the International Space Station (ISS) was required to perform five pre-determined debris avoidance manoeuvres (DAMs), raising the total to 21 since its launch in 1998 [7]. If a potential collision is detected with insufficient time to perform a safe DAM, the crew retreats to a Soyuz space-craft in preparation for immediate evacuation in the event that the ISS is catastrophically damaged. Despite no trend in the number of DAMs the ISS is required to perform each year, it is expected to increase as the total number of objects in low-Earth orbit continues to rise.

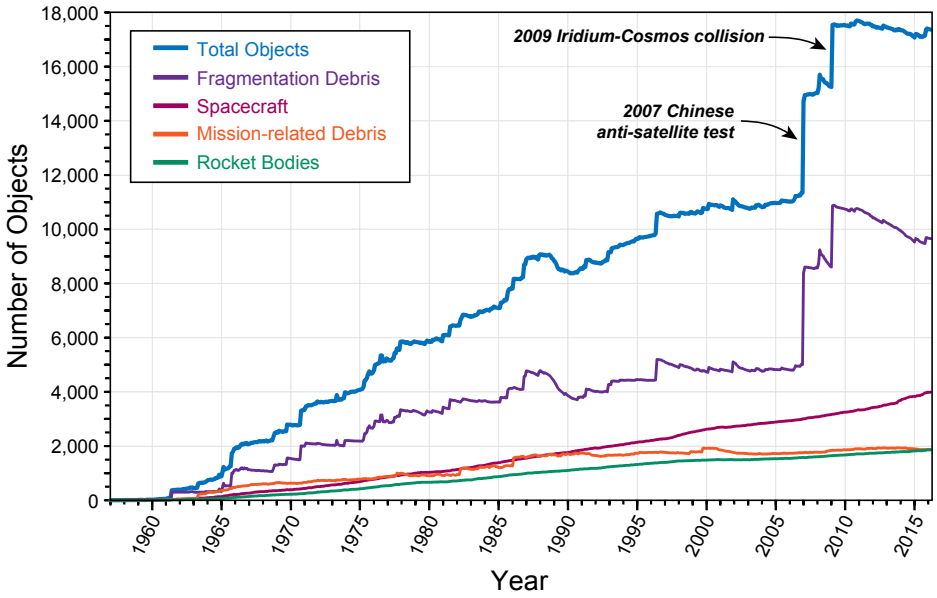

Figure 1.1: Summary of orbital debris objects in Earth orbit officially catalogued by the U.S. Space Surveillance Network. Fragmentation debris includes breakup debris and anomalous event debris. Mission-related debris includes all objects dispensed, separated, or released as part of a planned mission. Spacecraft include all operational and defunct satellites. Data retrieved from [3].

The population of orbital debris is, however, somewhat balanced by natural orbital decay due to atmospheric drag, which gradually degrades the velocity and therefore altitude of orbiting objects. Whilst drag will reduce the overall population of orbital debris over time, it is important to recognise that it will do so very slowly—certainly not any faster than the current rate of debris generation. Furthermore, because drag is proportional to the density of the atmosphere (which decreases with altitude), the effect it has on debris in mid- and high-Earth orbits is borderline negligible, which will remain in those orbits for thousands of years.

In order to preserve the near-Earth space environment, it is necessary to employ active mitigation strategies with the aim of reducing the likelihood of future collisions, and eventually removing debris from orbit altogether.

As of April 2016, the U.S. SSN monitors over 13,000 large debris objects (larger than 10 cm), forecasting their trajectories days in advance to predict the likelihood of potential conjunctions with operational satellites and other tracked debris [3]. This information allows satellite operators to assess the risk and, if necessary, manoeuvre their satellites to avoid collision[3]. This avoidance strategy is, however, only effective if at least one of the colliding objects has the ability to change its orbit. The majority of large objects in orbit do not have this capability [3].

Observations of debris too small to be observed by the U.S. SSN are conducted by various X- and Ku-band radars including Massachusetts Institute of Technology's Haystack radar [4], which track clouds of debris as opposed to individual objects. NASA's orbital debris program office, who are responsible for maintaining public orbital debris catalogues, estimate the current population of debris between 1 and 10 cm in diameter to be approx-

[3]These types of manoeuvres can be expensive, however, since satellite's have a limited supply of propellant used primarily for orbit maintenance. Evasive manoeuvres reduce a satellite's life-expectancy, and—as a consequence—their utility and commercial value.

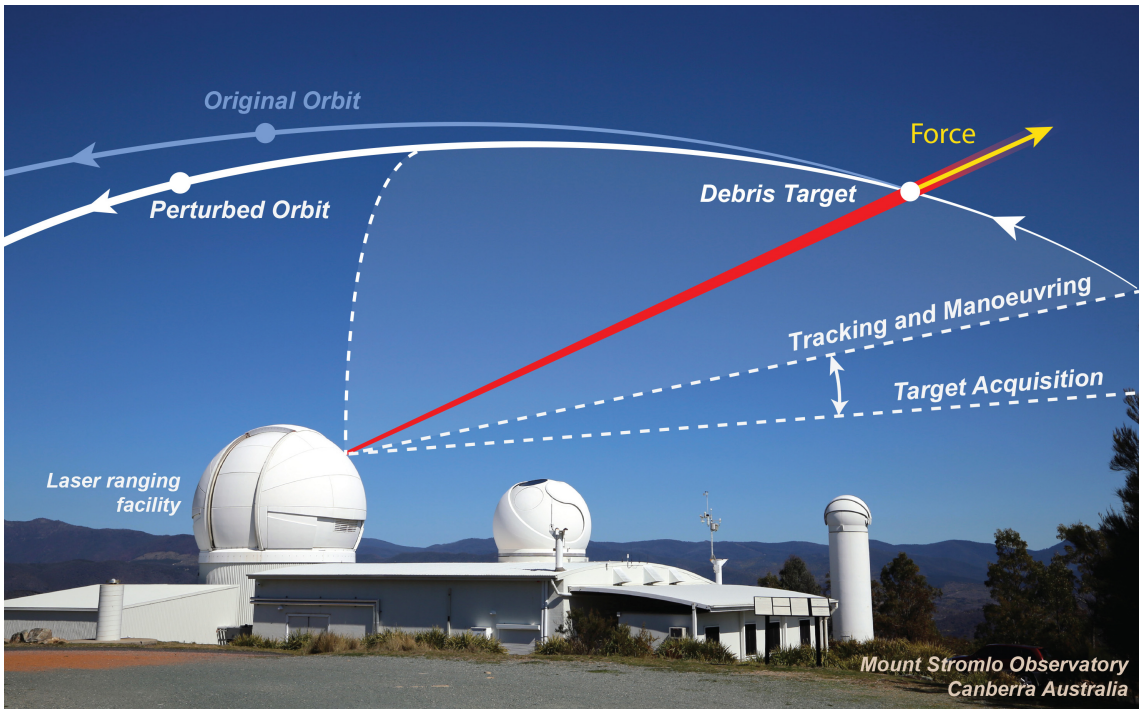

Figure 1.2: Remote orbital debris manoeuvring concept. Photograph taken by Lyle Roberts.

imately 500,000 [8]. The number of particles smaller than 1 cm is estimated to exceed 100 million. The threat posed by small debris is particularly high in low-Earth orbit due to the relatively high spatial density of space-craft and debris, and the fact that objects in low-Earth orbit travel at speeds up to 8 km/s, which means that collisions with even a small piece of debris involve considerable energy.

The world's capacity to manage space debris can be broadly separated into three categories: i) persistent identification and cataloguing of debris for pre-determined debris avoidance manoeuvres [9–11]; ii) passive mitigation to minimize mission-related debris produced by new satellite launches [12]; and iii) active mitigation using specific debris clearing satellites, or ground-based directed energy systems [13–15].

At present, pre-determined debris avoidance manoeuvres and passive mitigation are the only strategies employed to address the threat of orbital debris, despite a global effort to develop active mitigation technologies. Techniques to mechanically remove debris via rendezvous (e.g., electro-dynamic tethers and grapples [16]) are prohibitively expensive, and not effective at removing smaller debris [14]. There is also a risk associated with deploying intervening debris-clearing satellites into orbit, which are themselves at risk of collision.

A potentially cheaper, and lower risk option is to use high-power ground based lasers to remotely manoeuvre orbital debris via ablation [14]; and/or radiation pressure [13]. The concept for remotely manoeuvring orbital debris is shown in Figure 1.2, illustrating its potential to perturb an otherwise uncontrollable object's orbit. Ablation requires high-intensity pulsed lasers to evaporate material from the debris' surface to generate thrust in the form of a plasma jet. Radiation pressure is the result of photon momentum, and requires high average optical power which can be provided by continuous wave (CW) lasers. Ablation is not considered in this thesis, although details of its proposed application are presented in [14].

1.1.1 Remote manoeuvring via radiation pressure

The altitude of an orbiting satellite depends on its velocity v according to

$$A = \frac{GM_E}{v^2} - R_E$$

where R_E and M_E are the Earth's radius and mass respectively, and G is the universal gravitational constant. Perturbing a satellite's altitude therefore requires a change in its velocity Δv, which can be achieved by altering its momentum by exerting an external force on it. In this case the force is radiation pressure.

The momentum of an object with mass m and orbital velocity v is:

$$p_{\text{o}} = mv$$

The momentum of a large number of photons delivered by a laser is

$$p_\lambda = \frac{E}{c} = \frac{P\tau}{c}$$

where c is the speed of light in vacuum, and $E = P\tau$ is the total photon energy expressed as the total integrated photon power P over some period of illumination τ.

Consider the ideal scenario where photons engage an orbiting object head on. The change in the object's momentum is approximately

$$\Delta p_{\text{o}} = m\Delta v = \frac{P\tau}{c} \cdot R \cdot C_r$$

where C_r is the object's radiation pressure coefficient (related to the object's albedo), which is 0 for a perfectly transparent object, 1 for a perfectly absorbing object (i.e., a black body), and 2 for a perfectly reflecting material (e.g., a flat mirror facing the beam). The coefficient R describes the ratio of photons that hit the object's surface compared to the total number of photons delivered by the laser, which depends on the beam's shape and size as a function of propagation distance, the effective cross-sectional area of the orbiting object, and scintillation due to atmospheric turbulence. For this simple analysis it is important to recognise that R is proportional to the object's cross-sectional area.

The change in the object's velocity due to radiation pressure is therefore proportional to its cross-sectional area, and inversely proportional to its mass

$$\Delta v = \frac{P\tau}{mc} \cdot R \cdot C_r \tag{1.1}$$

which simply means that the orbit of objects with lower area to mass ratios are more difficult to manipulate. Equation (1.1) indicates that emphasis should be placed on increasing exposure time τ, the ratio of photons hitting the object R, and laser power P. The period of illumination τ can be increased by locating ground stations at high latitudes, or as close to the poles as possible to maximise engagement time on objects with polar orbits. The fraction of photons that hit an orbiting object is strongly influenced by the diffraction limit of the optical system used to illuminate it, and whether or not the telescope has the ability to compensate for atmospheric turbulence. The overall effectiveness of the integrated laser and telescope system also depends on the spatial quality of the laser, which influences the telescope's ability to focus the laser beam over long distances. High-power, single frequency mode lasers with high beam quality ($M^2 < 1.2$) are therefore preferred [13].

The analysis in [13] states that any ground-based laser manoeuvring system must be capable of effecting a Δv of at least 1 mm/s to overcome inherent ranging uncertainties (i.e., to provide any confidence that the object's orbit has actually been perturbed). The total estimated optical power required to change the velocity of a 5 kg object by $\Delta v = 1$ mm/s (assuming $C_r = 1$, $R = 0.1$, and a total exposure time of $\tau = 1000$ seconds) is 15 kW. It may soon be possible to achieve this using existing single mode commercial fibre lasers, which are currently capable of supporting up to 10 kW with good beam quality ($M^2 < 1.2$) [17]. Both IPG Photonics Corporation and Northrop Grumman have demonstrated 100 kW class lasers [18, 19], however the beam quality of these systems is currently unacceptably low ($M^2 \sim 3$ [19]) for applications requiring near diffraction limited propagation through the atmosphere.

1.2 Laser beam combining

One way to scale optical power beyond the capabilities of single laser sources is to combine the power of multiple lasers. Various laser beam combining techniques currently exist, and can be broadly separated into two categories: incoherent beam combining, and coherent beam combining (CBC). Incoherent beam combining refers to any configuration in which multiple lasers are combined without any attempt to stabilise their optical spectra or phases [20]. A good example of this is spectral beam combining (SBC) [21–30], in which multiple lasers operating at different wavelengths are combined using a wavelength sensitive beam combiner such as a prism or diffraction grating. An advantage of SBC is that it does not require strict temporal coherence of the combined beam, potentially simplifying its implementation. In 2011, SBC has demonstrated one of the highest total combined optical powers at 8.2 kW [27].

Coherent beam combining refers to the class of techniques in which a number of laser elements are combined with the same frequency and phase such that they interfere coherently [19, 31–62]. Two configurations used to perform CBC are tiled-aperture arrays (e.g., [42, 55]), in which the lasers are positioned next to each other; and filled-aperture implementations (e.g., [31]), where the interfering lasers are combined into a single beam with the same size and divergence. Both CBC implementations are illustrated in Figure 1.3 alongside spectral beam combining.

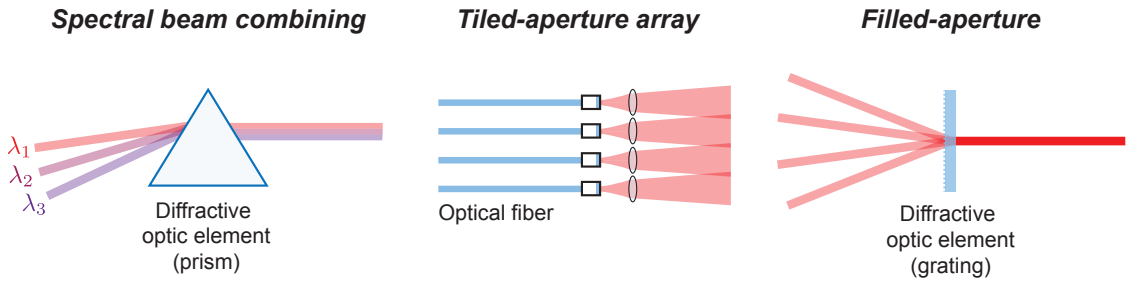

Figure 1.3: Illustration of different laser beam combining techniques. (Left) spectral beam combining involves combining multiple lasers operating at different wavelengths using a wavelength sensitive optical element. (Middle) tiled-aperture arrays are used to coherently interfere multiple spatially separate beams in the far field. (Right) filled-aperture implementation where multiple lasers are coherently interfered using a diffractive optic element.

In tiled-aperture implementations the beams interfere in the far field, producing far field

interference patterns where the peak intensity scales with the square of the number of emitters (N^2). In contrast, filled-aperture implementations interfere in the near-field where the peak far field intensity scales linearly with the number of emitters (N), effectively behaving like a single laser with increased optical power. The far field interference for tiled- and filled aperture implementations is illustrated in Figure 1.4 for three, four, and five element arrays.

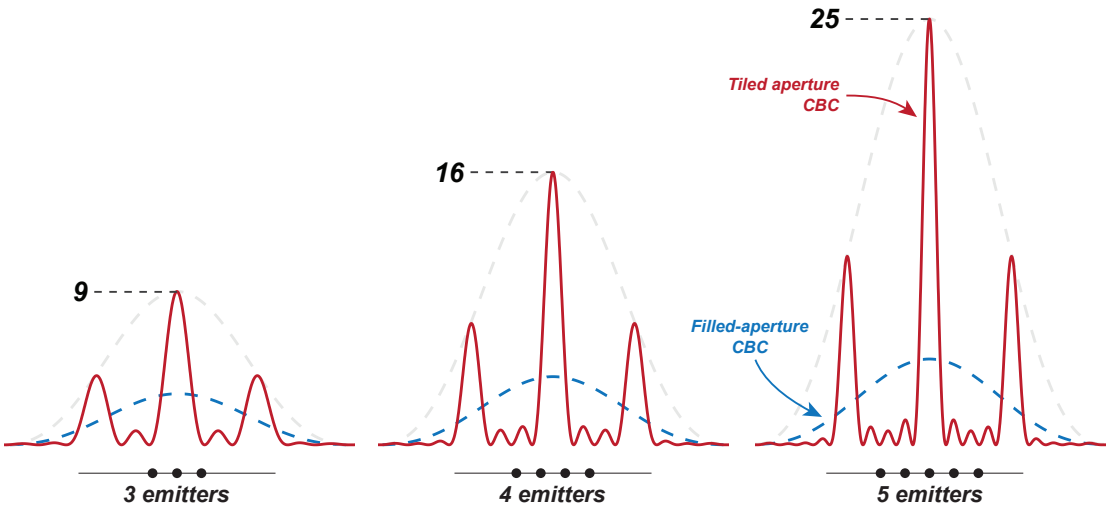

Figure 1.4: Coherent combination of multiple spatially separate optical fibre apertures results in quadratic scaling of the peak far field intensity, compared to linear scaling for incoherent combination.

Both CBC implementations require passive and/or active stabilisation of each laser's output phase. Passive stabilisation techniques rely on self phase-locking via passive coupling mechanisms (e.g., [49]), whereas active techniques rely on feedback control. The ultimate goal of CBC is to stabilise the relative output phase to within a fraction of a wavelength, which at optical frequencies (on the order of 1 μm) is extremely challenging.

1.2.1 Optical phased arrays

If the phase of each element in a tiled-array can be controlled independently, then it is possible to manipulate the distribution of optical power in the far field as shown in Figure 1.5, potentially enabling them to compensate for atmospheric turbulence (e.g., [36, 60]). With this functionality the tiled-aperture array becomes an optical analogue of the well established radar phased array [63], which are referred to in this thesis as optical phased arrays (OPAs).

This beam-forming functionality, combined with their scalability and high power handling capabilities present optical phased arrays as a promising technology for future space debris laser ranging and manoeuvring applications that require high intensity laser sources to be propagated through the atmosphere.

1.2.2 Performance characterisation

A key metric used to characterise the performance of CBC is root mean square (RMS) output phase stability, typically specified as a fraction of a wavelength (e.g., $\Phi_{\mathrm{RMS}} = \lambda/K$).

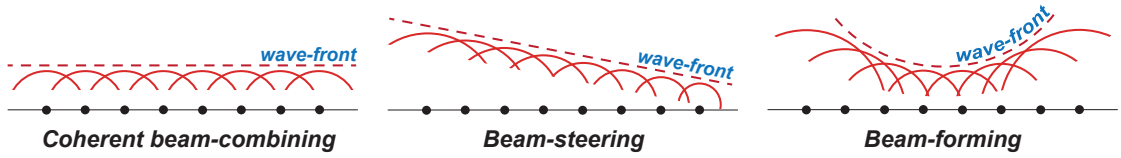

Figure 1.5: Optical phased arrays interfere multiple coherent lasers sources to manipulate the distribution of optical power in the far field. This allows them to perform beam-steering and forming by adjusting the relative output phase of individual emitters in the array.

It describes: i) the overall coherence of the optical wave-front, and ii) how effective the control system is at stabilising the relative output phase of the OPA. Phase stability at the output of the OPA is important because it influences the distribution of optical power in the far field. Increases in Φ_{RMS} lead to larger fluctuations in intensity at a fixed point in the far field, reducing the efficiency with which power is delivered to a specific location. This effect is shown in Figure 1.6, which was generated using a Monte Carlo simulation assuming the underlying RMS output phase error is normally distributed with standard deviation $\Phi_{\text{RMS}} = \lambda/20$, $\lambda/30$, and $\lambda/40$. The probability density functions (PDFs) of intensity on the right have log-normal distributions. RMS output phase stability can measured using heterodyne detection, or inferred from the characteristics of the interfered beam (e.g., fringe contrast, motion of the interference pattern).

To be effective at manoeuvring orbital debris, it was decided that an OPA must be able to deliver at least 90% intensity to a specific location in the far field no less than 90% of the time [64]. This requirement was subsequently translated into a maximum tolerable RMS output phase error using Monte Carlo simulations to generate cumulative density functions (CDFs) of intensity for different values of Φ_{RMS}. These simulations assumed a 100 element (10-by-10) square tiled array. The CDFs produced by these simulations are shown in Figure 1.7.

The intersection of the vertical and horizontal dashed red lines indicates the point at which there is a 10% likelihood of there being *less* than 90% intensity on target, which also means there is a 90% chance of there being *more* than 90% intensity. Any CDF that passes through the shaded region at the bottom right is therefore compliant. From these simulations the required RMS output phase error for a 100 element (10-by-10) array was found to be $\Phi_{\text{RMS}} \leq \lambda/24$.[4]

A detailed analysis of CBC performance is presented in [65].

1.2.3 External sensing

Almost all actively stabilised CBC technologies measure the output wave-front phase of the OPA by sensing the outgoing beam externally. Some of the most common external sensing techniques include optical heterodyne detection (e.g., [34, 38, 50]), stochastic parallel gradient descent (SPGD) (e.g., [54, 55]), and multiple-frequency dithering (e.g., [41, 42, 57, 61]).

[4]The topology of the OPA is specified in the requirement because it has a very slight dependence on number of emitters in the array. As the number of emitters in each dimension of the array increases, the width of the central interference lobe in those dimensions decreases (this can be seen in Figure 1.4). This narrowing of the central interference lobe increases the effect Φ_{RMS} has on the intensity in the far field. The required RMS output phase error for a 16 element (4-by-4) square tiled array is $\Phi_{\text{RMS}} \leq \lambda/23$, which is a slight relaxation of the $\lambda/24$ requirement for a 10-by-10 array.

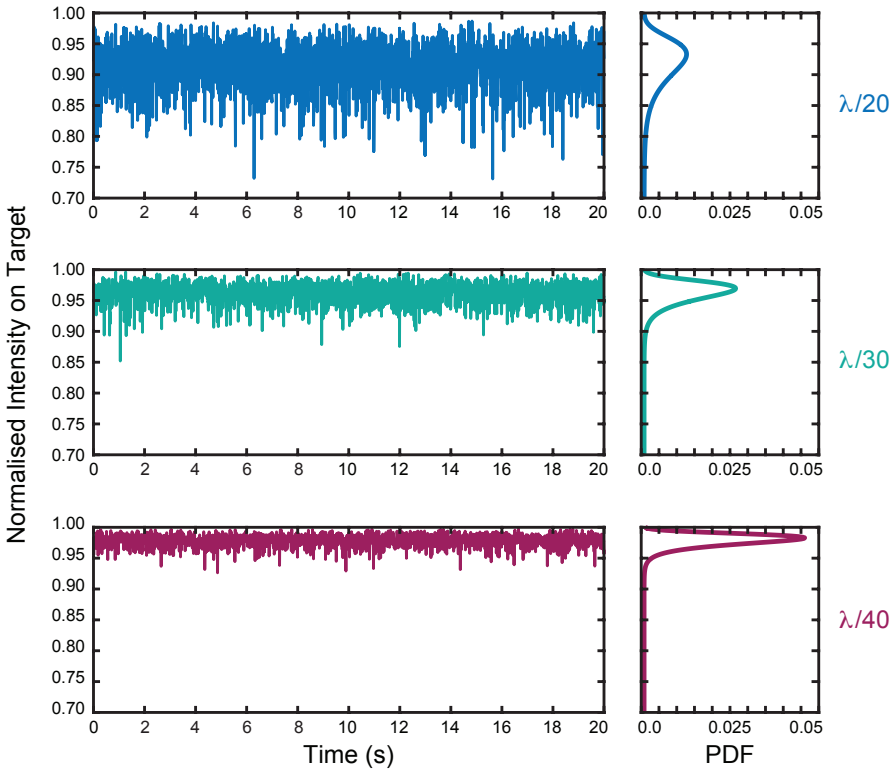

Figure 1.6: Monte Carlo simulation of normalised intensity for $\Phi_{\mathrm{RMS}} = \lambda/20$, $\lambda/30$, and $\lambda/40$. The probability density functions on the right highlight the increasing consistency with which power is delivered to a point in the far field as Φ_{RMS} decreases.

Optical heterodyne detection involves interfering the light from each emitter with a frequency shifted reference beam at an array of photodetectors. The heterodyne beat-note produced by each photodetector contains the phase information of its emitter and the reference beam. Since phase of the reference beam is common to each measurement, it can be eliminated by subtracting the measured phase of two emitters to produce a differential phase error signal for feedback control. Optical heterodyne detection's strength is its potential to precisely measure the true output phase of each emitter. Its weakness is the inherent requirement to precisely align the interfering beams onto N photodetectors for an N element array, making it susceptible to shock and vibration that may misalign the free-space optics. It may be possible to overcome this weakness using a monolithic waveguide structure, however this concept has not yet been demonstrated. Optical heterodyne detection has been used to coherently combine two beams with a total output power of 1.4 kW with an RMS output phase stability of $\lambda/80$ [50].

Stochastic parallel gradient descent is an intensity-based technique that uses a hill climbing algorithm to maximise far field intensity. A key advantage of SPGD is that it requires a single photodetector. It's disadvantage is that its control bandwidth scales inversely with the number of emitters, potentially affecting its ability to scale. SPGD has been demonstrated to coherently combine a total output power of 4 kW using an eight element array with an RMS output phase stability of $\lambda/40$ [55].

Multiple-frequency dithering works by modulating (or dithering) the phase of each element in the array at a unique radio frequency (RF), which (when interfered at a single external photodetector) allows phase errors between pairs of channels to be isolated via coherent demodulation. An example of multiple-frequency dithering is locking of optical coherence

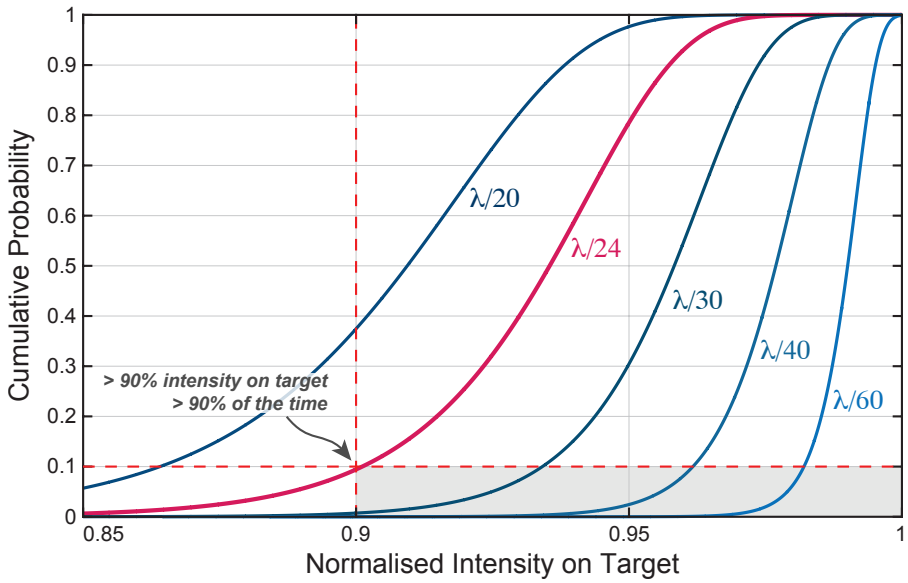

Figure 1.7: Simulated cumulative distribution functions (CDF) of intensity on target for different RMS output phase errors over a 100 element (10-by-10) array. Any CDF that passes through the shaded box at the bottom right of the figure satisfies the requirement to maintain at least 90% intensity no less than 90% of the time. The magenta curve identifies the minimum required RMS output phase stability at $\Phi_{\mathrm{RMS}} \leq \lambda/24$.

by single-detector electronic-frequency tagging (LOCSET) [42]. LOCSET's advantage is its relative simplicity and scalability, requiring only a single photodetector at the output of the array. A crucial limitation of LOCSET is that the relative phase between emitters cannot be locked to an arbitrary set point, potentially preventing it from being used to perform beam-steering and -forming.

LOCSET has been the subject of intense research and development for the previous decade, and is now a standard technique employed by United States Air Force Research Laboratory due to its robustness and versatility. In 2009 it was used to coherently combine five 145 W amplifiers with an RMS output phase of $\lambda/60$ [48], establishing its capacity to support high total output power. In 2012 it was used to coherently combine 32 low power (< 1 mW) beams with an RMS output phase error of $\lambda/71$, demonstrating its potential scalability [57].

1.2.4 Internal sensing

In contrast to external sensing, we present a technique that does not require free-space optics to measure the output phase of the beam, instead relying on the small fraction of light that is reflected back into the fibre at the OPA's glass-air interface to infer the relative phase of each emitter. This internal sensing technique infers the differential phase between uncommon paths by measuring the phase of the back-reflected light that double-passes each fibre. The relative phase at the output of the array is then stabilized by actively controlling the path lengths of each fibre.

The output phase of each emitter is isolated and measured using digitally enhanced heterodyne interferometry (DEHI) [66], which employs spread-spectrum modulation techniques to discriminate individual reflections at a single detector, without sacrificing the

sub-wavelength sensitivity of conventional heterodyne interferometry. DEHI works by exploiting the correlation properties of pseudo-random noise (PRN) codes that are phase modulated onto the carrier electric field using an electro-optic modulator (EOM). Signals belonging to the various reflections in the optical system are then isolated by digitally demodulating the detected signal using correspondingly delay matched versions of the PRN code, suppressing contributions from other reflections. The phase of each demodulated signal is measured using digitally implemented phasemeters to generate feedback control for relative path-length stabilisation. Field-programmable gate-arrays (FPGAs) are utilized extensively in the implementation of the OPA's control system.

Throughout this thesis, we define internal sensing to be any technique for beam combining that does not require external sampling optics to sense the OPA's relative output phase for stabilisation.

1.3 Thesis Overview

The structure of this thesis reflects the evolutionary development of the internally sensed OPA.

Chapters 2 and 3 introduce **digitally enhanced heterodyne interferometry** and digitally implemented **phasemeters**, providing a baseline understanding of the two core technologies that enable internal sensing. The advantage of DEHI is its ability to discriminate individual reflections at a single photodetector, without sacrificing the sub-wavelength sensitivity of conventional heterodyne interferometry. Digitally implemented phasemeters are sensors that enable the precise measurement of differential optical path lengths in the OPA.

The first stage of the internally sensed OPA's development was to establish a **proof-of-concept** experiment to assess the viability of internal sensing. Details of this characterization experiment, experimental results, and a discussion of key challenges associated with its implementation are presented in Chapter 4. Whilst the proof-of-concept OPA demonstrated the viability of internal sensing, it also highlighted two crucial weaknesses that would limit its usefulness in most applications. Of particular concern was its limited power handling capabilities.

Chapter 5 introduces a revised OPA architecture that overcomes the low power handling capabilities of the proof-of-concept system using dedicated offset phase-locked slave lasers and asymmetric fibre couplers. The **offset-phase locked OPA** was implemented experimentally, demonstrating its ability to coherently combine three independent slave lasers with an RMS output phase stability of $\lambda/104$.

The final OPA design presented in Chapter 6 improves upon the offset phase-locked OPA by using fibre amplifiers to deliver up to 15 W of optical power to each emitter. The **amplifier compatible OPA** was implemented experimentally, demonstrating an output phase stability of $\lambda/194$, and the ability to steer the coherently combined beam at up to 10 kHz. This chapter concentrates on the technical details surrounding the experimental implementation of the amplifier compatible OPA.

Conclusions drawn from the development of the internally sensed OPA are presented in Chapter 7, including a discussion of opportunities for future development.

1.4 Other applications

The work presented in this thesis has potential applications beyond remote orbital debris manoeuvring. For example, the control system developed to stabilise the relative output phase of the internally sensed optical phased array could be used to combine semiconductor lasers [47], whereas the OPA itself could be used to compensate for atmospheric turbulence in free-space optical communications systems [36, 60], or to perform segmented wave-front sensing.

1.5 Publications

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

L.E. Roberts, R.L. Ward, A.J. Sutton, R. Fleddermann, G. de Vine, D.M.Wuchenich, E.A. Malikides, D.E. McClelland, and D.A. Shaddock, *"Coherent beam combining using a 2D internally sensed optical phased array,"* Appl. Opt. **53**(22), 4881—4885 (2014).

L.E. Roberts, R.L. Ward, S.P. Francis, P.G. Sibley, R. Fleddermann, A.J. Sutton, C. Smith, D.E. McClelland, and D.A. Shaddock, *"High power compatible internally sensed optical phased array,"* Opt. Express **24**(12), 13467—13479 (2016).

Digitally enhanced heterodyne interferometry

2.1 Introduction

Digitally enhanced heterodyne interferometry is an advanced form of optical metrology that employs spread-spectrum modulation techniques to discriminate multiple interferometric signals at a single photodetector (PD), without sacrificing the sub-wavelength sensitivity of conventional heterodyne interferometry [66, 67]. DEHI works by exploiting the correlation properties of pseudo-random noise codes that are phase modulated onto the carrier electric field using an electro-optic modulator. Signals belonging to the various interferometric signals in the optical system are then isolated by digitally demodulating the detected signal using correspondingly delay matched versions of the PRN code, suppressing contributions from other reflections. DEHI shifts measurement complexity away from the optical system, instead relying on the efficient utilisation of digital signal processing (DSP) resources. It is also readily scalable and well suited to the computational power of field-programmable gate-arrays.

2.2 Interferometry

Interferometry is a technique used to extract information from electro-magnetic waves by measuring the intensity of their interfered electric fields. A simple homodyne Mach-Zehnder interferometer is shown in Figure 2.1, where light emitted from a laser is separated into two paths at a free-space beam-splitter (BS_1), followed by their recombination some time later at a second beam-splitter (BS_2). The amplitude reflectivity and transmissivity of the beamsplitters are r and it respectively. The convention that electric fields experience a $\pi/2$ phase shift upon transmission through beam-splitters is used throughout this thesis, which is sufficient to ensure conservation of energy.

The recombined fields will interfere constructively or destructively depending on the relative path lengths of the two arms, L_1 and L_2. By measuring the power of the recombined fields, $P = E^*E$, it is possible to infer the relative optical path length difference of the two arms with sub-wavelength precision.

The propagation lengths of the separated electric fields can be expressed as a common

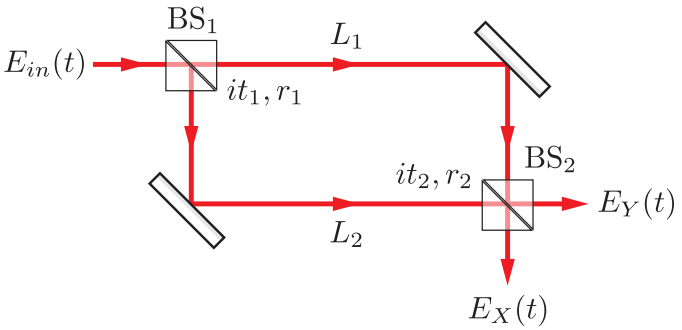

Figure 2.1: Homodyne Mach-Zehnder interferometer.

arm length L with a differential arm length component $\Delta L/2$:

$$L_1 = L + \frac{\Delta L}{2}$$
$$L_2 = L - \frac{\Delta L}{2}$$

The relative optical path length difference of the two arms is thus:

$$L_1 - L_2 = \left(L + \frac{\Delta L}{2}\right) - \left(L - \frac{\Delta L}{2}\right) = \Delta L$$

As the separated electric fields propagate along each arm of the interferometer they accumulate phase inversely proportional to the wavelength of the electric field, λ

$$\phi_1 = \left(\frac{2\pi}{\lambda}\right) L_1 = kL_1$$
$$\phi_2 = \left(\frac{2\pi}{\lambda}\right) L_2 = kL_2$$

where $k = 2\pi/\lambda$ is the wave number. The separated electric fields just prior to recombination at the second beamsplitter BS_2 are

$$E_1(t) = it_1 E_{in} e^{i(\omega t + \phi_1)}$$
$$E_2(t) = r_1 E_{in} e^{i(\omega t + \phi_2)}$$

which means the recombined fields exiting BS_2 are:

$$E_X(t) = -t_1 t_2 E_{in} e^{i(\omega t + \phi_1)} + r_1 r_2 E_{in} e^{i(\omega t + \phi_2)}$$
$$E_Y(t) = i r_1 t_2 E_{in} e^{i(\omega t + \phi_1)} + i r_2 t_1 E_{in} e^{i(\omega t + \phi_2)}$$

The optical power at the two outputs of BS_2 are thus

$$\begin{aligned} P_X(t) &= E_{in}^2 \left(t_1^2 t_2^2 + r_1^2 r_2^2 + r_1 r_2 t_1 t_2 \left(e^{i(\phi_1 - \phi_2)} - e^{-i(\phi_1 - \phi_2)}\right)\right) \\ &= E_{in}^2 \left(t_1^2 t_2^2 + r_1^2 r_2^2 - 2 r_1 r_2 t_1 t_2 \cos(k\Delta L)\right) \end{aligned} \tag{2.1}$$

and:

$$\begin{aligned} P_Y(t) &= E_{in}^2 \left(r_1^2 t_2^2 + t_1^2 r_2^2 + r_1 r_2 t_1 t_2 \left(e^{i(\phi_1 - \phi_2)} + e^{-i(\phi_1 - \phi_2)}\right)\right) \\ &= E_{in}^2 \left(r_1^2 t_2^2 + t_1^2 r_2^2 + 2 r_1 r_2 t_1 t_2 \cos(k\Delta L)\right) \end{aligned} \tag{2.2}$$

If $r_1 = r_2 = t_1 = t_2$, equations (2.1) and (2.2) simplify to:

$$P_X(t) = \frac{E_{in}^2}{2}\big(1 - \cos(k\Delta L)\big)$$
$$P_Y(t) = \frac{E_{in}^2}{2}\big(1 + \cos(k\Delta L)\big)$$

The intensity of the interfered fields is therefore sensitive to changes in differential optical path length, ΔL.

2.2.1 Heterodyne interferometry

Heterodyne interferometry involves interfering an electric field with a frequency shifted local oscillator (LO) to produce a beat-note at their difference frequency. The local oscillator is often used as a reference, while the unshifted beam 'interrogates' the optical system. From this point we refer to the unshifted light as the interrogating beam.

A simple heterodyne Mach-Zehnder interferometer is shown in Figure 2.2, where an acousto-optic modulator (AOM) [68] is used to shift the frequency of the field in the lower arm of the interferometer by $+f_h$.

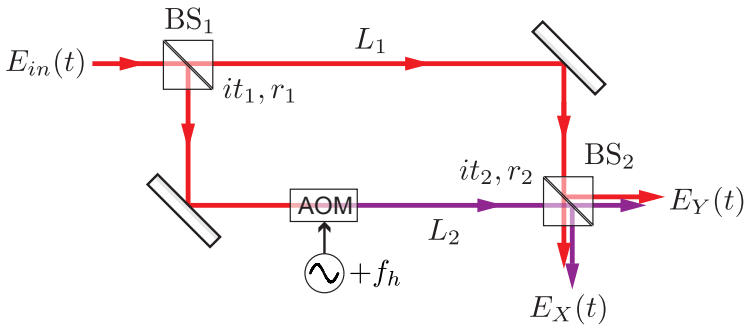

Figure 2.2: Heterodyne Mach-Zehnder interferometer.

Assuming BS_1 and BS_2 have equal transmissivity and reflectivity, the optical power at the two outputs of BS_2 is

$$P_X(t) = \frac{E_{in}^2}{2}\big(1 - \cos(2\pi f_h t + \delta\phi)\big) \tag{2.3}$$
$$P_Y(t) = \frac{E_{in}^2}{2}\big(1 + \cos(2\pi f_h t + \delta\phi)\big) \tag{2.4}$$

where $\delta\phi = k\Delta L$. Changes in relative optical path length ΔL are thus captured in the *phase* of the heterodyne beat-note oscillating at f_h, which can be measured with exceptional precision using a phasemeter (e.g., [69]).

2.3 Pseudo-random noise codes

Pseudo random noise codes are deterministic and periodic binary sequences that resemble random noise. PRN codes have finite length, and are made up of a pseudo random sequence of 1's and 0's called 'chips'. An example PRN code is shown in Figure 2.3, identifying an individual chip and the code's period.

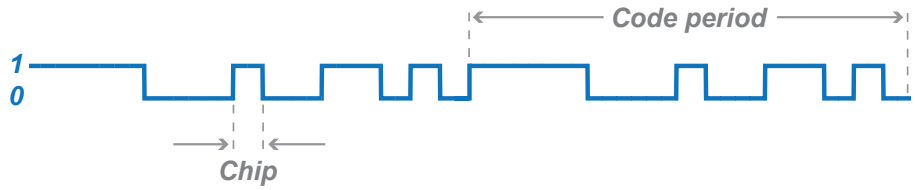

Figure 2.3: Example 15 chip PRN sequence.

The choice of PRN for a particular application depends on their cross- and auto-correlation properties. Cross-correlation is a measure of how similar one signal is to a different signal as a function of delay, τ. Auto-correlation is a measure how similar a signal is to a delayed version of *itself*. Weak cross-correlation is desirable in systems requiring multiple PRN sequences, for example the global positioning system. In general, DEHI only requires a single PRN sequence, motivating the selection of codes with strong auto-correlation.

One class of PRN called maximal-length sequences (m-sequences) have auto-correlation functions (averaged over a single code period) defined by

$$A(\tau) = \begin{cases} 2^N - 1 & \text{for } \tau = 0, L, 2L, \cdots \\ -1 & \text{for any other } \tau \end{cases} \tag{2.5}$$

where $L = 2^N - 1$ represents the length of the sequence measured in chips, and τ is a discrete time delay measured in chips. This auto-correlation function is illustrated in Figure 2.4.

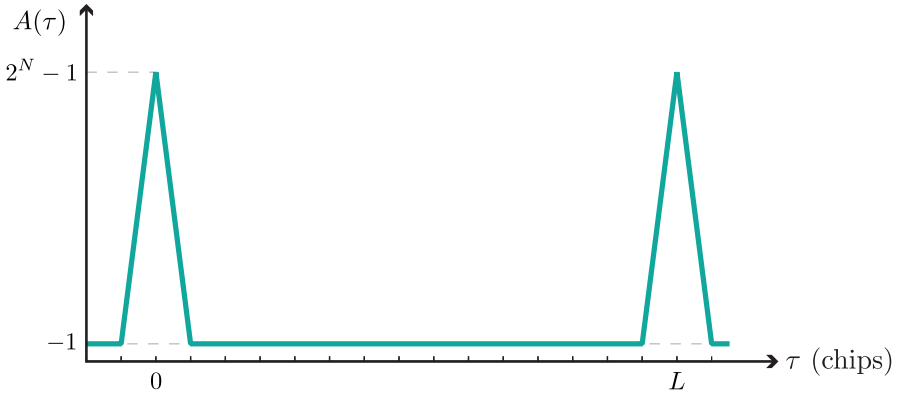

Figure 2.4: Auto-correlation profile of an m-sequence. Auto-correlation peaks occur when the two codes have a relative delay of $\tau = 0$. For all other delays the auto-correlation is -1.

2.3.1 Generating maximal length sequences

Maximal-length sequences are created efficiently using linear feedback shift registers (LFSRs) [70, 71]. The properties of an m-sequence depend on the LFSR's length (number of bits), and specific feedback 'tap' configuration. An N-bit LFSR can produce a PRN sequence with maximum length $2^N - 1$, if and only if its arrangement of feedback taps is described by the binary coefficients of a primitive polynomial of order N. A 4-bit LFSR with tap arrangement defined by the primitive polynomial $x^4 + x^3 + 1$ is shown in Figure 2.5. Subsequent 'chips' in the code are determined by performing the modulo-2

(binary) addition[1] of the LFSR's third and fourth elements ($x^4 + x^3$). The state of the LFSR is advanced by shifting its contents to the right by one bit, and simultaneously writing the new chip into its leftmost bit.

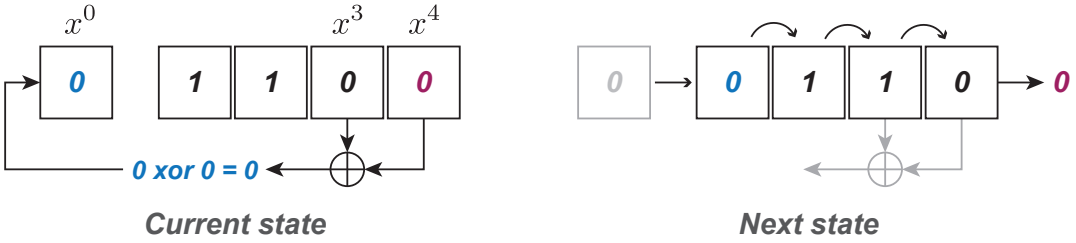

Figure 2.5: 4-bit linear feedback shift register with feedback tap arrangement defined by the primitive polynomial $x^4 + x^3 + 1$.

The transition from the current state of a LFSR to the next can be described mathematically using linear algebra and binary arithmetic

$$\mathbf{S}_{\text{new}} = \mathbf{M}\mathbf{S}_{\text{old}} \tag{2.6}$$

where $\mathbf{M}$ represents an $m \times m$ matrix that maps one state to the next, $\mathbf{S}_{\text{old}}$ is an $m \times 1$ vector representing the current state of the LFSR, and $\mathbf{S}_{\text{new}}$ is an $m \times 1$ vector representing the new state. The contents of the matrix $\mathbf{M}$ are defined by an identity matrix, and the binary coefficients of a primitive polynomial. Consider the generic state transition matrix

$$\mathbf{M} = \begin{bmatrix} t_1 & \cdots & t_{m-1} & t_m \\ 1 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{2.7}$$

where the coefficients t_1 to t_m represent the binary coefficients of an order m primitive polynomial (i.e., $t_m x^m + t_{m-1}x^{m-1} + \ldots + t_1 x^1$). The zero'th coefficient ($x^0 = 1$) does not correspond to a feedback tap, and instead represents the newly generated chip. If the current state of the LFSR is an $m \times 1$ column vector with elements a_1 to a_m, then the next state will be:

$$\mathbf{S}_{\text{new}} = \begin{bmatrix} t_1 & \cdots & t_{m-1} & t_m \\ 1 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_{m-1} \\ a_m \end{bmatrix} = \begin{bmatrix} t_1 a_1 \oplus \cdots \oplus t_{m-1} a_{m-1} \oplus t_m a_m \\ \vdots \\ a_{m-2} \\ a_{m-1} \end{bmatrix} \tag{2.8}$$

The result of $\mathbf{S}_{\text{new}}$ defines the LFSR's next state logic for a specific feedback tap arrangement defined in $\mathbf{M}$. The next state logic for the 4-bit LFSR in Figure 2.5 is:

$$\mathbf{S}_{\text{new}} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} = \begin{bmatrix} a_3 \oplus a_4 \\ a_1 \\ a_2 \\ a_3 \end{bmatrix}$$

[1] Binary addition is performed in an FPGA using asynchronous exclusive-or (XOR) gates.

The next state logic defined in equation (2.8) can be extended to generate multiple chips in the PRN sequence at a time [72], which is necessary in order to advance or delay one code relative to another (e.g., when scanning through code delays during demodulation).

If the initial and next state of the LFSR are $\mathbf{S}_0$ and $\mathbf{S}_1 = \mathbf{M}\mathbf{S}_0$ respectively, then $\mathbf{S}_2$ is:

$$\mathbf{S}_2 = \mathbf{M}\mathbf{S}_1 = \mathbf{M}^\mathbf{2}\mathbf{S}_0$$

The LFSR can thus be advanced k times during a single transition by multiplying the initial state by $\mathbf{M}^\mathbf{k}$:

$$\mathbf{S}_\mathrm{k} = \mathbf{M}^\mathbf{k}\mathbf{S}_0 \tag{2.9}$$

The next state logic required to advance the state of the 4-bit LFSR shown in Figure 2.5 by two steps is:

$${}^2\mathbf{S}_\mathrm{new} = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}^2 \begin{bmatrix} a_3 \\ a_2 \\ a_1 \\ a_0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} a_3 \\ a_2 \\ a_1 \\ a_0 \end{bmatrix} = \begin{bmatrix} a_2 \oplus a_1 \\ a_1 \oplus a_0 \\ a_3 \\ a_2 \end{bmatrix} \tag{2.10}$$

The k-channel parallel m-sequence is generated by reading the k least significant bits of the LFSR, identified in equation (2.10) by the k elements, a_3 and a_2. An N-bit LFSR can be advanced by at most N steps in a single clock cycle.

2.3.2 Spectral characteristics of PRN

The spectral characteristics of binary PRN depends on the chip frequency f_{chip}, which is the rate at which new chips are modulated; and the code repetition frequency, which describes the rate at which the sequence repeats itself:

$$f_{code} = \frac{f_{chip}}{L}$$

Simulated amplitude spectra for 3, 4, and 5-bit binary m-sequences comprised of 1's and -1's (normalised to the peak amplitude of the 3-bit code) modulated at $f_{chip} = 5$ MHz are shown in Figure 2.6. All three spectra consist of a sinc^2 envelope with nulls at integer multiples of the PRN chip frequency, and harmonics located at integer multiples of the code repetition frequency. The effective power in the harmonics decreases with increasing code length as power is distributed between more of them. The harmonic at 0 Hz is the result of a DC offset caused by the unequal number of 1's and -1's in the m-sequence.

2.4 PRN modulation

The principal difference between conventional heterodyne interferometry and DEHI is the phase modulation of the interrogating field with PRN. A simple digitally enhanced heterodyne Mach-Zehnder interferometer is shown in Figure 2.7, which includes an electro-optic modulator in the interrogating beam to modulate its phase with a PRN code.

The electric field of the phase modulated beam when it arrives at BS_2 is

$$E_{PRN}(t) = itE_{in}e^{i(2\pi ft+\phi+\beta c(t-\tau))}$$

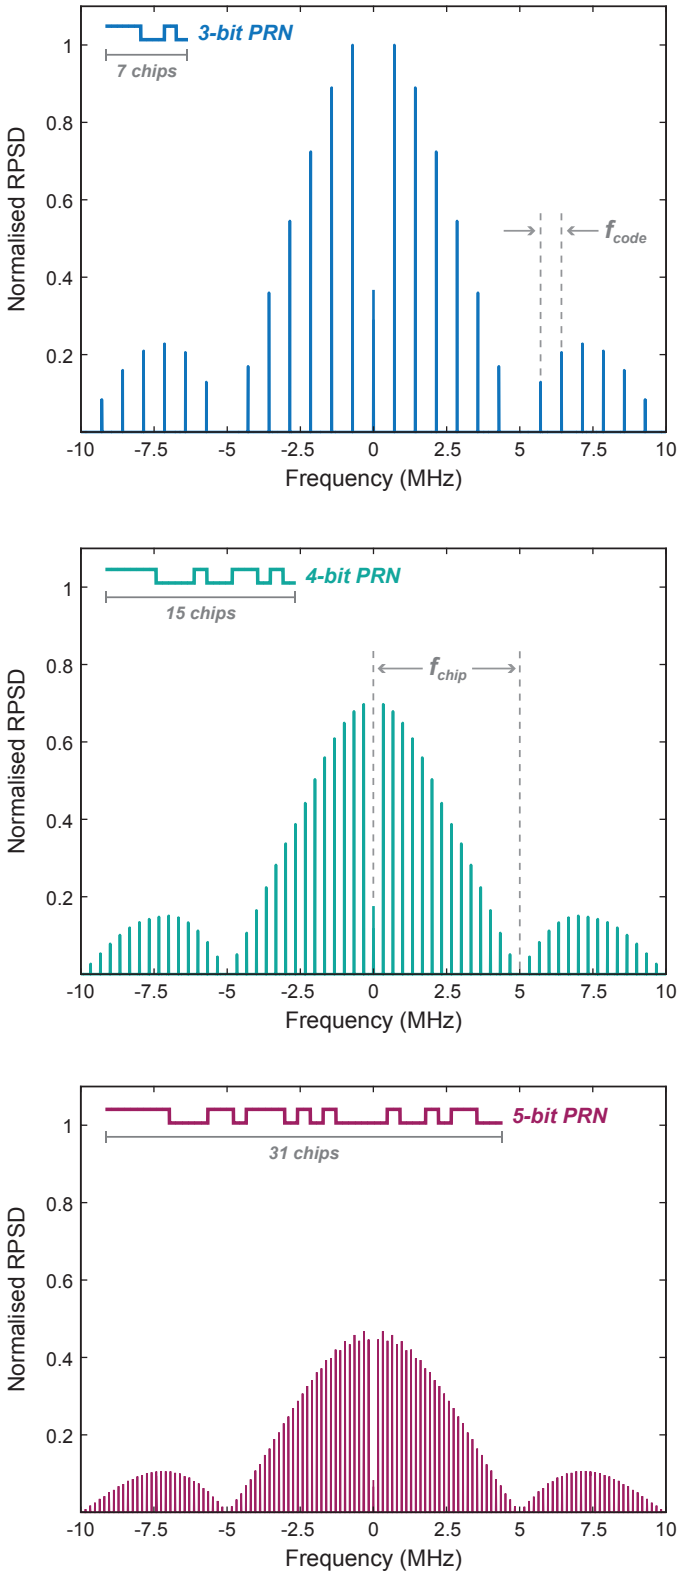

Figure 2.6: Simulated spectral distributions of 3, 4, and 5-bit binary PRN codes with a 5 MHz chip frequency, normalised to the peak amplitude of the 3-bit code. The PRN spectra consists of a sinc^2 envelope with nulls at integer multiples of the PRN chip frequency, and harmonics located at integer multiples of the code repetition frequency.

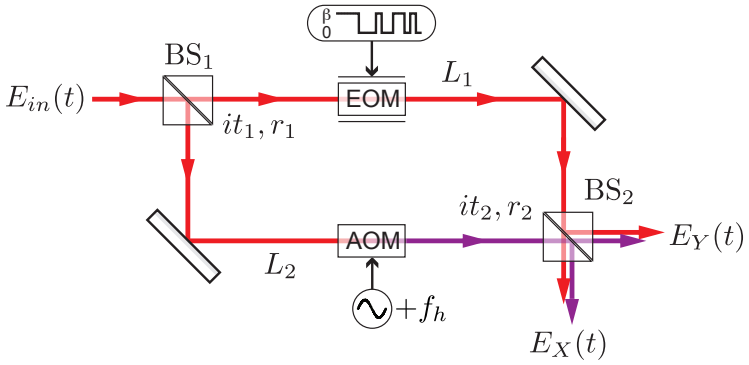

Figure 2.7: Digitally enhanced heterodyne Mach-Zehnder interferometer.

where ϕ represents phase, $c(t) \in [0, 1]$ is the modulating PRN code with modulation depth β, and τ is the propagation delay from the EOM to BS_2. Assuming BS_1 and BS_2 have equal transmissivity and reflectivity, the optical power at the output Y is

$$P_Y(t) = \frac{E_{in}^2}{2}\big(1 + \cos(2\pi f_h t + \delta\phi - \beta c(t-\tau))\big) \tag{2.11}$$

where f_h represents the heterodyne frequency shift applied to the local oscillator by the AOM, and $\delta\phi$ is the phase difference between the two arms of the interferometer. When the modulation depth $\beta = \pi$, we can use the identity $\cos(\theta \pm \pi) = -\cos(\theta)$ to simplify equation (2.11) using the equation $p(t-\tau) = 1 - 2c(t-\tau)$ to map $c(t) \in [0, 1]$ to $p(t) \in [1, -1]$:

$$P_Y(t) = \frac{E_{in}^2}{2}\big(1 + p(t-\tau)\cos(2\pi f_h t + \delta\phi)\big) \tag{2.12}$$

This allows us to represent the AC component of the PRN modulated signal as the multiplication of the PRN code $p(t) \in [1, -1]$ and un-modulated carrier, $\cos(2\pi f_h t + \delta\phi)$.

Multiplication in the time domain is equivalent to convolution in the frequency domain. The spectrum of a sinusoid is an impulse centred at its frequency of oscillation, f_h. The spectrum of binary PRN is shown in Figure 2.6, which has a sinc^2 envelope with nulls at integer multiples of the chip frequency f_{chip}, and harmonics located at integer multiples of the code repetition frequency. The spectrum of the PRN modulated tone is therefore the spectrum of the PRN sequence centred at the impulse frequency f_h, as shown in Figure 2.8.

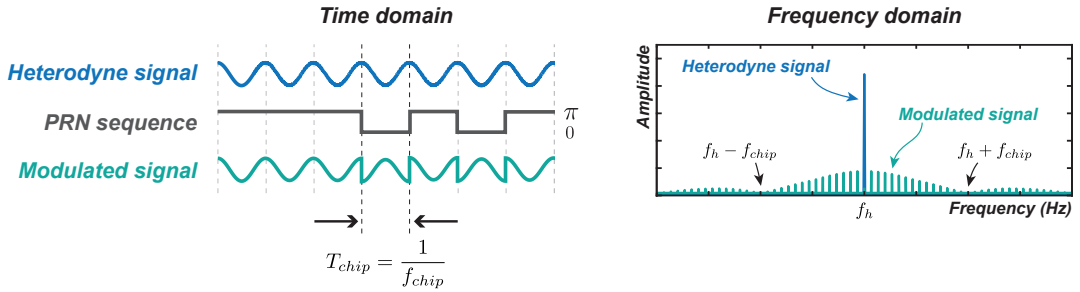

Figure 2.8: Time and frequency domain illustration of the PRN phase modulated heterodyne signal. The frequency spectrum of the PRN modulated signal is that of the PRN sequence centred at the signal frequency, f_h.

2.5 PRN demodulation

Conventional heterodyne detection responds to the vector sum of all interferometric signals arriving at the photodetector. Phase information belonging to each signal is lost because there is no way to distinguish them. When the interrogating beam is modulated with PRN, different signals originating within the interferometer can be isolated based on their unique electro-optic propagation delays by demodulating (or *decoding*) the photodetector output with appropriately delayed copies of the modulating PRN code. The decoded signal can then be time-averaged to exploit the auto-correlation properties of PRN.

A crucial property of PRN demodulation is that the multiplication of two identical and delay matched codes results in $p(t-\tau_i)\cdot p(t-\tau_i)=1$, and $p(t-\tau_i)\cdot p(t-\tau_j)$ produces another random sequence (which is *not* necessarily an m-sequence). Signals with different electro-optic propagation delays relative to the demodulating sequence therefore remain as broadband noise. This process is illustrated in Figure 2.9.

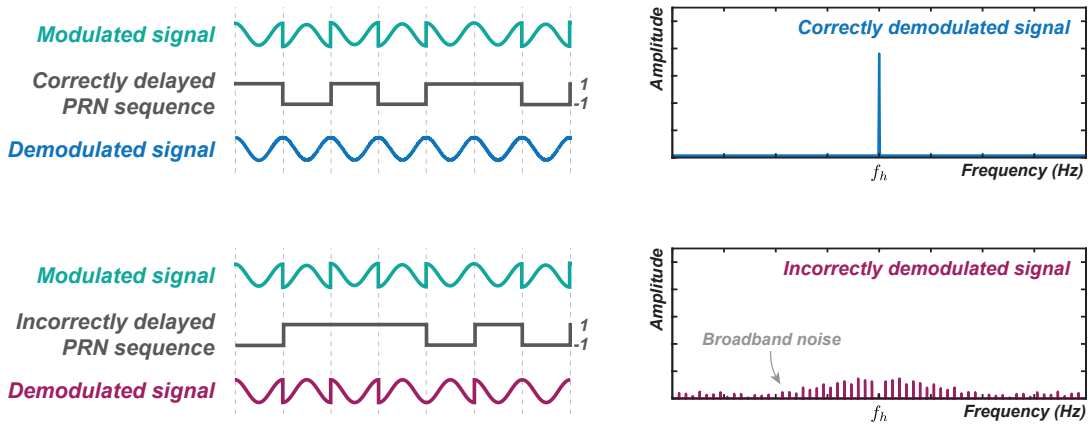

Figure 2.9: Time and frequency domain illustrations of correctly (top) and incorrectly (bottom) demodulated signals.

2.5.1 Multiplexing

A key advantage of DEHI is its ability to interrogate individual interferometric signals at a single photodetector based on their specific electro-optic delays, whilst suppressing spurious interference that would otherwise limit measurement sensitivity. This multiplexing capability depends on the spatial resolution—or optical 'range-gate'—of the modulating PRN sequence.

The range-gate of a modulating PRN sequence is the length of an individual PRN chip f_{chip} propagating at the speed of light c through a medium of refractive index n

$$d_r \geq \frac{c}{n f_{chip}} \tag{2.13}$$

and describes the minimum tolerable optical path length separation of two signals such that they can be individually resolved at the detector. If two signals originate from within the same range-gate, then they will be unable to be resolved as they will both be partially recovered during demodulation, limiting the achievable measurement sensitivity due to cross-talk.

Figure 2.10 shows a multiplexed digitally enhanced heterodyne Mach-Zehnder interferometer implemented in optical fibre where one of the arms has been split into three paths with different optical path lengths. Light propagating through each path will experience a unique time delay τ_k proportional to its total optical path length L_k

$$\tau_k = \frac{nL_k}{c} \tag{2.14}$$

where c is the speed of light in vacuum, and n is the refractive index of the optical fibre.

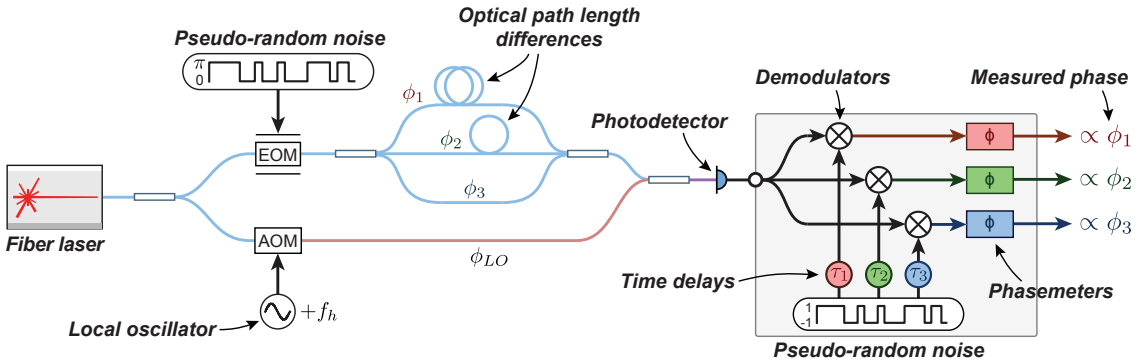

Figure 2.10: Multiplexed digitally enhanced heterodyne Mach-Zehnder interferometer.

The electric field at the photodetector consists of the sum of the three PRN modulated signals, interfered against the local oscillator field. Ignoring DC terms, the measured interference signal $s(t)$ is

$$\begin{aligned} s(t) = E_0^2 \Big[& \alpha_1 \cos(2\pi f_h t + \phi_{LO} - \phi_1 - \pi c(t - \tau_1)) \\ & + \alpha_2 \cos(2\pi f_h t + \phi_{LO} - \phi_2 - \pi c(t - \tau_2)) \\ & + \alpha_3 \cos(2\pi f_h t + \phi_{LO} - \phi_3 - \pi c(t - \tau_3)) \Big] \end{aligned} \tag{2.15}$$

where E_0^2 is the total optical power delivered by the fibre laser, α_k is an attenuation coefficient, f_h is the heterodyne beat note frequency, ϕ_k and ϕ_{LO} represent the phase of the individual paths and local oscillator respectively, and $c(t-\tau_k) \in [0, 1]$ is the PRN code delayed by τ_k. The identity $\cos(\theta \pm \pi) = -\cos(\theta)$ allows us to simplify equation (2.15), again using the equation $p(t - \tau) = 1 - 2c(t - \tau)$ to map $c(t) \in [0, 1]$ to $p(t) \in [1, -1]$:

$$\begin{aligned} s(t) = E_0^2 \Big[& \alpha_1 \, p(t - \tau_1) \cos(2\pi f_h t + \phi_{LO} - \phi_1) \\ & + \alpha_2 \, p(t - \tau_2) \cos(2\pi f_h t + \phi_{LO} - \phi_2) \\ & + \alpha_3 \, p(t - \tau_3) \cos(2\pi f_h t + \phi_{LO} - \phi_3) \Big] \end{aligned}$$

Assuming the differential optical path lengths of the three arms are greater than the modulating PRN code's range-gate, we can isolate the the first signal by demodulating $s(t)$ with the appropriately delayed PRN code $p(t - \tau_1)$:

$$\begin{aligned} d_1(t) &= s(t) \cdot p(t - \tau_1) \\ &= E_0^2 \Big[\alpha_1 \cos(2\pi f_h t + \phi_{LO} - \phi_1) \\ &\quad + \alpha_2 \, p(t - \tau_1) \cdot p(t - \tau_2) \cos(2\pi f_h t + \phi_{LO} - \phi_2) \\ &\quad + \alpha_3 \, p(t - \tau_1) \cdot p(t - \tau_3) \cos(2\pi f_h t + \phi_{LO} - \phi_3) \Big] \end{aligned}$$

The signal-to-noise ratio of the demodulated signal $d_1(t)$ can then be improved by averaging over the full length of the PRN code (or at least some fraction of it) to benefit from the auto-correlation properties of m-sequences described in Section 2.3. Averaging over the full code length L yields

$$\begin{aligned}\sum^{L} d_1(t) \approx E_0^2 \big[&\alpha_1 L \cos(2\pi f_h t + \phi_{LO} - \phi_1) \\ &- \alpha_2 \cos(2\pi f_h t + \phi_{LO} - \phi_2) \\ &- \alpha_3 \cos(2\pi f_h t + \phi_{LO} - \phi_3)\big]\end{aligned} \tag{2.16}$$

which shows that the time-averaged correctly demodulated signal is approximately L times greater in amplitude than the incorrectly demodulated signals.

2.6 Field-programmable gate-arrays

Field-programmable gate-arrays are integrated circuits that contain arrays of configurable logic blocks that can be wired together via programmable interconnects [73]. Individual logic blocks can be configured to perform basic combinational logic operations (e.g., AND and XOR operations) and store information using registers. Multiple logic blocks can be configured in series and/or parallel to perform complex digital signal processing algorithms.

FPGAs have inherently parallel and deterministic architectures, making them well suited to low-latency, high-throughput applications that would otherwise not be implementable using a central processing unit (CPU). That being said, the relatively high complexity of some operations like floating point division are better suited to dedicated hardware, such as the dedicated floating point units found in modern CPUs.

FPGAs are configured using code written in a hardware description language (e.g., VHDL, Verilog [74]) that is translated into hardware level logic by dedicated software. The FPGA implemented DSP discussed in this thesis was programmed using LabVIEW to target a National Instruments (NI) FPGA development platform.

2.6.1 Integer arithmetic

Numbers in the hardware description language are typically represented as unsigned or signed N-bit integers. Unsigned numbers can represent any positive integer between 0 and 2^N-1, whereas signed N-bit numbers can represent both positive and negative integers between -2^{N-1} and $2^{N-1} - 1$.

Integer addition, multiplication, and subtraction are well supported arithmetic operations within the FPGA. Division, however, is not well supported due to inefficient consumption of logic resources [74]. For this reason, division in the FPGA is almost exclusively limited to inverse powers of two, 2^{-X}, which can be performed efficiently by bit-shifting the number to the right by X bits. Division by non-powers of two can be performed using a combination of bit shifts and multipliers. For example, division by five can be approximated by multiplying the number by 205, and then bit shifting it to the right by 10 (i.e., dividing it by 1024), where the resulting fractional error disappears since the FPGA rounds down to the nearest integer.

Phasemeters

3.1 Introduction

The term *phasemeter* is used to classify any device that measures the phase of an oscillating signal. Phase can be measured in various ways, and three of the most common approaches are zero-crossing detection, which infer phase by measuring the time interval between zero-crossings of two periodic signals; in-phase (I) and quadrature (Q) demodulation, which measures phase by calculating the inverse tangent of the signals' orthogonal I and Q components; and phase-locked loops (PLLs), which use feedback to continuously update the phase of a local model of the periodic signal, thus providing a measurement of the phase itself.

3.1.1 Zero-crossing detector

Zero-crossing detectors (ZCD) infer the phase difference between two signals by measuring the interval of time Δt between relative zero-crossings between signals, which is proportional to their phase difference $\Delta\phi$ as illustrated in Figure 3.1. For zero-crossings in a single direction (e.g., negative to positive), phase is calculated using the equation

$$\Delta\phi = \left(\frac{\Delta t}{T}\right)2\pi$$

where T is the measured period of oscillation.

Because the ZCD is only sensitive to zero-crossings, its effective sampling frequency is equal to the frequency f_h of the input signal (or twice f_h if both zero-crossings are detected in a single cycle). This means the ZCD ignores all phase information between successive zero-crossings, and that any noise at integer multiples of f_h will alias into the measurement as shown in Figure 3.2.

Zero-crossing detectors are simple to implement and are accurate enough for many applications requiring rudimentary phase measurements. Their use is limited, however, when either of the two input signals has a low signal-to-noise ratio (SNR) and/or more than a single harmonic, since these conditions can both trigger spurious zero-crossings.

3.1.2 In-phase and quadrature demodulation

In-phase and quadrature (IQ) demodulation measures a signal's phase by calculating the arctangent of its in-phase and quadrature components. Consider the input signal

$$s_{in}(t) = A\sin(\omega_{in}t + \phi_{in}) \tag{3.1}$$

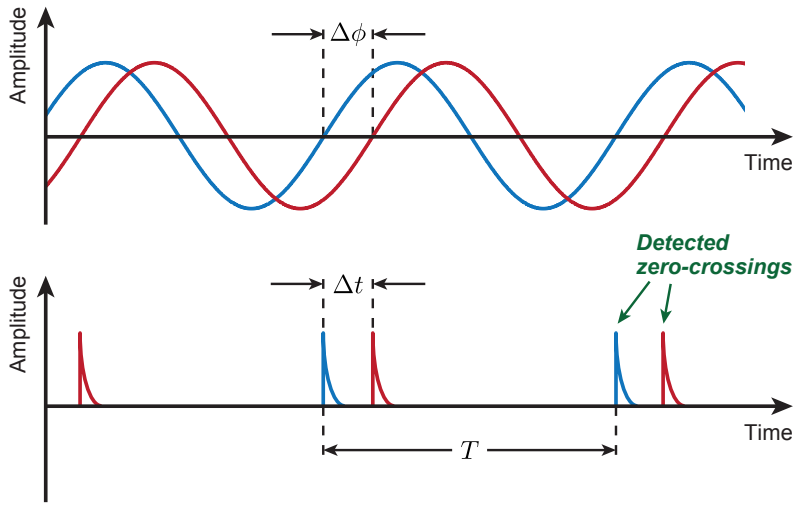

Figure 3.1: Zero-crossing detector phasemeter concept. The time interval between zero-crossings of two periodic signals is proportional to their phase difference.

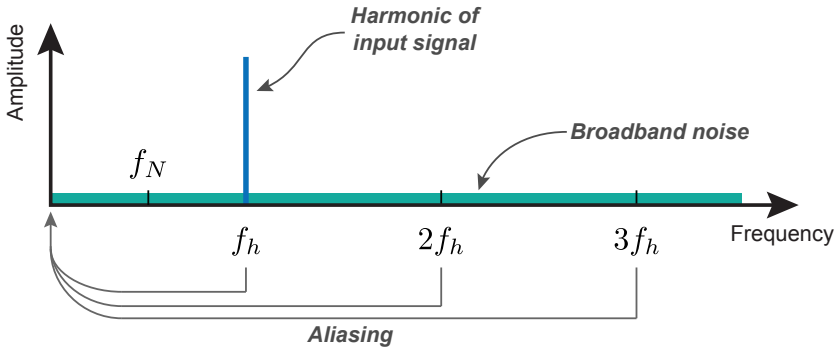

Figure 3.2: Illustration of how noise aliases into the measurement of a zero-crossing detector. Because ZCDs are only sensitive to zero-crossings, they can be thought of as effectively re-sampling the input signal at its frequency, f_h. This causes noise at integer multiples of f_h to alias to DC where it is indistinguishable from the signal's true phase.

with amplitude A, frequency ω_{in}, and phase ϕ_{in}. The I and Q components of $s_{in}(t)$ can be resolved via demodulation with orthogonal local oscillators

$$s_1(t) = B\cos(\omega_{in}t + \phi_{LO})$$
$$s_2(t) = B\sin(\omega_{in}t + \phi_{LO})$$

which have amplitude B, frequency ω_{in}, and common local oscillator phase ϕ_{LO}. Multiplying $s_{in}(t)$ separately with $s_1(t)$ and $s_2(t)$ gives us

$$\begin{aligned} s_{M,1}(t) &= s_{in}(t)s_1(t) \\ &= A\sin(\omega_{in}t + \phi_{in})B\cos(\omega_{in}t + \phi_{LO}) \\ &= \frac{AB}{2}\Big(\sin(\Delta\phi) + \sin(2\omega_{in}t + \phi_{in} + \phi_{LO})\Big) \end{aligned} \tag{3.2}$$

and

$$\begin{aligned} s_{M,2}(t) &= s_{in}(t)s_2(t) \\ &= A\sin(\omega_{in}t + \phi_{in})B\sin(\omega_{in}t + \phi_{LO}) \\ &= \frac{AB}{2}\Big(\cos(\Delta\phi) - \sin(2\omega_{in}t + \phi_{in} + \phi_{LO})\Big) \end{aligned} \tag{3.3}$$

where $\Delta\phi = \phi_{in} - \phi_{LO}$. Equations (3.2) and (3.3) contain the Q and I components $\sin(\Delta\phi)$ and $\cos(\Delta\phi)$ needed to calculate $\Delta\phi$, as well as non-linear second harmonics generated by the multiplication process: $\sin(2\omega_{in}t + \phi_{in} + \phi_{LO})$. Before the I and Q components can be isolated, the second harmonic component must first be removed using a low-pass filter. After filtering, equations (3.2) and (3.3) become the I and Q components of $s_{in}(t)$:

$$s_I(t) = \frac{AB}{2}\cos(\Delta\phi) \tag{3.4}$$

$$s_Q(t) = \frac{AB}{2}\sin(\Delta\phi) \tag{3.5}$$

The phase difference $\Delta\phi$ is recovered by calculating the inverse tangent of $s_Q(t)$ and $s_I(t)$:

$$\begin{aligned} \Delta\phi &= \tan^{-1}\left(\frac{s_Q(t)}{s_I(t)}\right) \\ &= \tan^{-1}\left(\frac{\frac{AB}{2}\sin(\Delta\phi)}{\frac{AB}{2}\cos(\Delta\phi)}\right) \\ &= \tan^{-1}\left(\frac{\sin(\Delta\phi)}{\cos(\Delta\phi)}\right) \end{aligned} \tag{3.6}$$

A consequence of I and Q demodulation is the discontinuity in the calculated phase $\Delta\phi$, bounded by the asymptotic limits of the inverse tangent function at $\pm\pi$. Whenever $\Delta\phi$ exceeds these limits it will immediately 'wrap' by $\mp 2\pi$ to the opposite boundary. The calculated phase must therefore be unwrapped in order to reconstruct a continuous measurement of $\Delta\phi$.

The amplitude of the input signal can be measured using Pythagoras' theorem, and accounting for the multiplication gain by scaling the result by $\frac{2}{B}$:

$$\begin{aligned} A_{\text{meas}} &= \frac{2}{B}\sqrt{s_I^2(t) + s_Q^2(t)} \\ &= \frac{2}{B}\frac{AB}{2}\sqrt{\sin^2(\Delta\phi) + \cos^2(\Delta\phi)} \\ &= A \end{aligned} \tag{3.7}$$

Whilst I and Q demodulation phasemeters are highly effective at measuring phase, they are not easily implemented in a field programmable gate array. Division at high numerical precision requires a substantial amount of FPGA resources, and in practice it is more efficient to perform division and any subsequent calculations using a central processing unit instead. A disadvantage of using a CPU is that the calculation of phase cannot be performed in hard real-time, which may have some implications for applications requiring low-latency feedback of phase information.

3.1.3 Phase-locked loop

Phase-locked loops measure phase by generating a local oscillator whose phase is proportional to that of the input signal. The phase of the local oscillator is continuously updated using a feedback loop that detects and corrects any phase error between the two signals.

The two primary functions of the PLL are to: 1) detect the phase error between the LO model and the input signal; and 2) update the LO model. The phase error is detected using a frequency mixer. The LO model is updated by adjusting the phase of a locally generated signal proportional to the detected phase error. The functional architecture of a PLL is illustrated Figure 3.3.

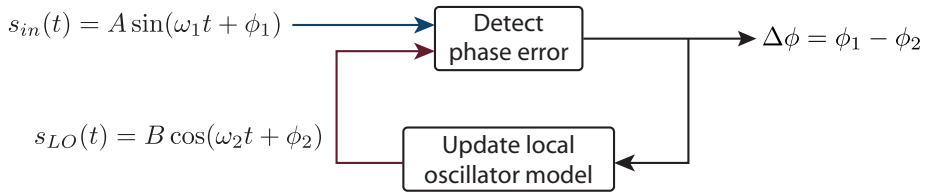

Figure 3.3: Functional architecture of a phase-locked loop.

Consider the input signal $s_{in}(t)$ and local oscillator $s_{LO}(t)$ with amplitudes A and B, angular frequencies ω_1 and ω_2, and phases ϕ_1 and ϕ_2:

$$s_{in}(t) = A\sin(\omega_1 t + \phi_1) \tag{3.8}$$

$$s_{LO}(t) = B\cos(\omega_2 t + \phi_2) \tag{3.9}$$

The multiplication of $s_{in}(t)$ and $s_{LO}(t)$ gives us

$$\begin{aligned} s_M(t) &= s_{in}(t)s_{LO}(t) \\ &= A\sin(\omega_1 t + \phi_1)B\cos(\omega_2 t + \phi_2) \\ &= \frac{AB}{2}\Big(\sin\big((\omega_1 - \omega_2)t + \phi_1 - \phi_2\big) + \sin\big((\omega_1 + \omega_2)t + \phi_1 + \phi_2\big)\Big) \\ &= \frac{AB}{2}\Big(\sin\big((\omega_1 - \omega_2)t + \Delta\phi\big) + \sin\big((\omega_1 + \omega_2)t + \phi_1 + \phi_2\big)\Big) \end{aligned} \tag{3.10}$$

where $\Delta\phi = \phi_1 - \phi_2$. Assuming $\omega_1 \approx \omega_2 \approx \omega$, equation (3.10) becomes:

$$s_M(t) = \frac{AB}{2}\Big(\sin\big(\Delta\phi\big) + \sin\big(2\omega t + \phi_1 + \phi_2\big)\Big)$$

When the PLL is locked, the relative phase difference between $s_1(t)$ and $s_2(t)$ is very small ($\Delta\phi \ll 1$), allowing us to simplify $s_M(t)$ further using the small angle approximation for sine ($\sin(\theta) \approx \theta$ for $\theta \ll 1$):

$$s_M(t) = \frac{AB}{2}\Big(\Delta\phi + \sin\big(2\omega t + \phi_1 + \phi_2\big)\Big) \tag{3.11}$$

Equation (3.11) states that when the PLL is locked, the mixed signal is proportional to the sum of the phase error $\Delta\phi$ and the second harmonic $\sin\big(2\omega t + \phi_1 + \phi_2\big)$ generated by the mixing process. The second harmonic behaves non-linearly and can be suppressed using a low-pass filter, yielding:

$$s_F(t) = \frac{AB}{2}\Delta\phi = \frac{AB}{2}(\phi_1 - \phi_2) \tag{3.12}$$

The filtered signal $s_F(t)$ is thus proportional to the relative phase difference $\Delta\phi$ of the two signals $s_{in}(t)$ and $s_{LO}(t)$, scaled by the frequency independent mixer gain $M = \frac{AB}{2}$.

Phase-locked loops work in a similar way to IQ demodulation, except there is no need to calculate the inverse tangent of the ratio between $s_Q(t)$ and $s_I(t)$ since the feedback

loop continuously forces $\Delta\phi \ll 1$. This can be intuited as the orthogonal axes of the I and Q demodulating LOs rotating with those of the input signal. When $\Delta\phi \ll 1$, the demodulated I and Q components from equations (3.4) and (3.5) become

$$s_I(t) = \frac{AB}{2}\cos(\Delta\phi) \approx \frac{AB}{2} \tag{3.13}$$

$$s_Q(t) = \frac{AB}{2}\sin(\Delta\phi) \approx \frac{AB}{2}\Delta\phi \tag{3.14}$$

since $\cos(\theta) \approx 1$ and $\sin(\theta) \approx \theta$ for $\theta \ll 1$. Phase information $\Delta\phi$ is therefore contained entirely within the quadrature demodulated component $s_Q(t)$, and amplitude information is contained entirely within the in-phase demodulated component $s_I(t)$. This has the benefit of removing any need to divide $s_Q(t)$ by $s_I(t)$ since

$$\tan^{-1}\left(\frac{\sin(\Delta\phi)}{\cos(\Delta\phi)}\right) \approx \tan^{-1}\left(\Delta\phi\right) \approx \Delta\phi$$

where $\tan^{-1}(\theta) \approx \theta$ for $\theta \ll 1$. Phase-locked loops therefore have an advantage over IQ demodulation in that no division or inverse tangent calculations must be performed to retrieve a continuous phase error signal proportional to $\Delta\phi$, making them well suited to high throughput applications when implemented using field-programmable gate-arrays.

Despite their advantages over ZCD and IQ demodulation phasemeters, PLLs are considerably more complicated to design and optimise because of their dependence on feedback control. Phase-locked loops also experience an inherent bandwidth limitation due to the frequency dependent gain of the control system, limiting their ability to measure phase disturbances that occur faster than the loop bandwidth. Furthermore, the phase readout is only reliable when the PLL is behaving linearly, which is when the small angle approximation is valid (i.e., when $\Delta\phi \ll 1$). If this condition is not satisfied then the measured phase will suffer non-linearities proportional to $\sin(\Delta\phi)$.

3.2 Phase-locked loop design and implementation

Phase-locked loops are a core component of the internally sensed optical phased array. The relative phase error between the input signal and the LO model is detected using a digital multiplier, and the second harmonic generated by the multiplication process is attenuated using a low-pass filter. The detected phase error is fed back through a controller to update the phase of the LO model, which is implemented digitally as a numerically controlled oscillator (NCO). The NCO is used as a reference against which to measure the phase of the input signal (essentially defining $\phi_2 = 0$ in equation (3.9). This means that any detected phase errors are proportional only to disturbances in the phase of the input signal, ϕ_1.

The PLL design presented here was in large part inspired by work performed by Daniel Shaddock et al. [75] and Oliver Gerberding et al. [76, 77] for the Laser Interferometer Space Antenna gravitational wave observatory [78].

3.2.1 Design requirements

The performance of the PLL can be characterised by three performance metrics:

1. **Sensitivity** $\phi_{\min}$, which defines how sensitive the PLL is to phase disturbances of the input signal;

2. **Bandwidth** f_{BW}, which describes the range of frequencies over which the PLL can track phase; and
3. **Throughput** f_{update}, which describes the frequency with which the PLL provides an updated phase measurement.

The PLL presented here has been designed to achieve a measurement sensitivity of $\phi_{\mathrm{min}} < 1\ \mu\mathrm{cycle}/\sqrt{\mathrm{Hz}}$[1], a bandwidth f_{BW} of up to 100 kHz; and a throughput of $f_{\mathrm{update}} > 1$ MHz.

3.2.2 Discrete-time transfer function models

The PLL presented here was designed using a combination of transfer function (TF) and time-domain models prior to implementation in an FPGA. This approach afforded us the flexibility to quickly evaluate different PLL configurations. Details of the PLL's transfer function model are provided in Section 3.3.

Each component in the PLL is represented as either a gain or discrete-time transfer function in the z-domain[2]. Components with frequency dependent transfer functions are described using the parameter $z = \exp(j2\pi f/f_s)$ where f_s represents the sampling frequency of the discrete-time system and f represents the continuous frequency variable from 0 to f_s [79].

3.2.3 Physical architecture

The physical architecture of the digitally implemented PLL is shown in Figure 3.4, identifying the different sub-components used to perform the functions shown in Figure 3.3.

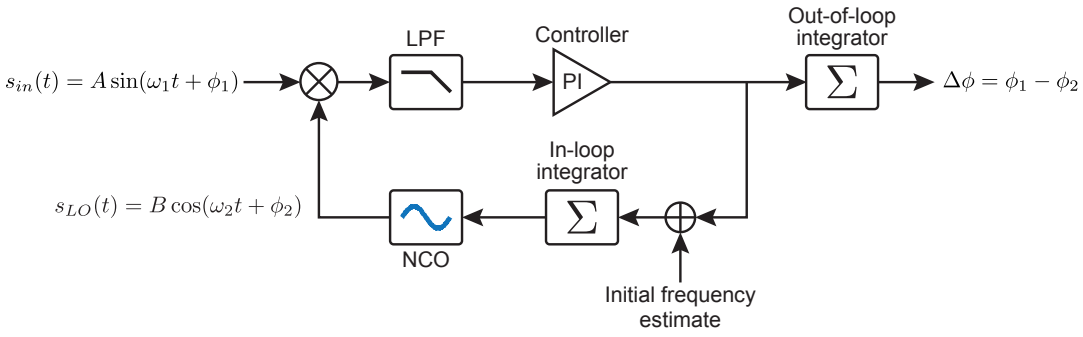

Figure 3.4: Physical architecture of a digitally implemented phase-locked loop. The input signal is mixed with a local oscillator that is generated using a numerically controlled oscillator. The mixed signal is low-pass filtered and passed through a controller before being fed back to update the local oscillator. The output of the controller is proportional to frequency error, which is integrated out-of-loop to recover phase.

The input signal $s_{in}(t) = A\sin(\omega_1 t + \phi_1)$ is mixed with a local oscillator $s_{LO}(t) = B\cos(\omega_2 t + \phi_2)$ that is generated using a numerically controlled oscillator. The mixed signal is then low-pass filtered to suppress the second harmonic and passed through a controller. The output of the controller is proportional to the instantaneous frequency error

[1]This requirement is adapted from the Laser Interferometer Space Antenna mission requirement for a phasemeter precision of approximately 1 μcycle [77].

[2]The z-domain is the discrete time analogue of the Laplace domain.

between $s_{in}(t)$ and $s_{LO}(t)$, which is used to correct the PLL's initial frequency estimate. The corrected frequency is then integrated in-loop over a fixed range (proportional to 0 to 2π) to generate a periodic phase signal used by the NCO to generate the local oscillator. The output of the controller is proportional to frequency error, which is converted into phase using an out-of-loop integrator.

3.2.4 Digital multiplier

The digital multiplier is used to estimate the phase-error between the input and LO signals. This process is described mathematically in Section 3.1.3.

The gain for the digital multiplier is:

$$M = \frac{AB}{2} \cdot 2^{-Q} \tag{3.15}$$

where A and B are the amplitudes of the input and NCO, and Q is a truncation factor used to resize the output of the multiplier. The overall gain of the PLL is sensitive to the amplitude of the input signal, which could potentially lead to unstable behaviour if the PLL's controller is not configured appropriately.

One way to overcome the mixer's amplitude dependent gain is to implement an automatic-gain algorithm that scales the controller gain proportional to the input signal's amplitude.

3.2.5 Low-pass filter

A rate-reducing low-pass filter is used to suppress the second harmonic generated by the mixer, and must be designed carefully to ensure that the attenuated non-linear second harmonic does not corrupt the PLL's measurement bandwidth. The type of filter used in this system is a constant coefficient finite impulse response (FIR) second order cascaded integrator comb (CIC) filter, selected primarily for its linear phase response and efficient implementation in FPGA [80]. A block diagram of a second order CIC filter is shown in Figure 3.5.

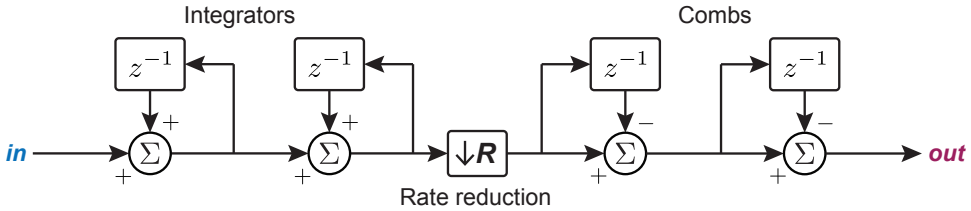

Figure 3.5: Block diagram of a second order CIC filter. The input is integrated across two consecutive discrete-time integrators. The output of the second integrator is then re-sampled at f_s/R and then passed through two consecutive differentiators (combs).

Cascaded integrator-comb filter

Cascaded integrator-comb filters consecutively integrate the input signal N times, re-sample the output of the N'th integrator, and then pass the rate-reduced integrated signal through N consecutive differentiators (which are often referred to as 'combs'[3]).

[3]Differentiators are often referred to as 'combs' because their frequency-domain transfer functions look like hair-combs.

The transfer function of N consecutive discrete-time integrators all running at the full sampling rate f_s is:

$$I(z) = \left(\frac{1}{1 - z^{-1}}\right)^N \tag{3.16}$$

The transfer function of a rate-reduced discrete-time differentiator operating at f_s/R is:

$$D(z) = \left(1 - z^{-R}\right)^N \tag{3.17}$$

The overall transfer function of an N'th order CIC filter is thus:

$$L(z) = \left(\frac{1 - z^{-R}}{1 - z^{-1}}\right)^N \tag{3.18}$$

The gain of an N'th order CIC filter at frequencies $f \ll f_s$ is equal to R^N:

$$\begin{aligned} |H_F(f)|_{f \ll f_s} &= \left(\frac{f_s}{\pi f}\sqrt{\sin^2\left(\pi R \frac{f}{f_s}\right)}\right)^N \\ &= \left(\frac{f_s}{\pi f}\sqrt{\left(\pi R \frac{f}{f_s}\right)^2}\right)^N \\ &= R^N \end{aligned} \tag{3.19}$$

3.2.6 Controller

The controller used in the PLL combines proportional (P) and integral (I) control elements. The proportional controller responds immediately and proportionally to the magnitude of any detected phase disturbances. The integral controller provides long term feedback by integrating the instantaneous phase error over time to eliminate residual steady-state error (SSE). Both P and I controllers are scaled by a tunable 'global' gain K_G, allowing a fixed ratio between K_P and K_I. The PI controller configuration is shown in Figure 3.6.

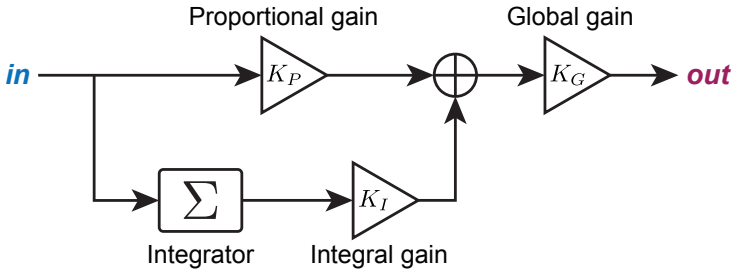

Figure 3.6: Proportional-plus-integral controller implemented in the PLL. The proportional controller provides immediate feedback, proportional to the magnitude of the input phase error Φ_{error}. The integral controller provides long term stability by integrating Φ_{error} over time to eliminate steady-state error. The combined proportional plus integral control output is scaled by a tunable global gain.

The motivation for using proportional-plus-integral (PI) control (as opposed to purely proportional control) is that it can provide a non-zero constant output even when the phase-error at the input is zero. This is important because a) it is possible for the frequency of the input signal to change over time, and b) if the initial estimate of the input frequency

is incorrect then there will be a constant frequency error that must be compensated. Integral controllers overcome SSE by applying feedback proportional to the time integrated SSE, which increases in magnitude the longer the error persists.

The discrete-time transfer function of the controller's integrator is

$$K_I(z) = \frac{K_I}{1 - z^{-R}} \tag{3.20}$$

where K_I is the integral gain. Equation (3.20) can be represented in terms of frequency f by substituting the discrete time z-parameter

$$z = \exp\left(j2\pi\frac{f}{f_s}\right) \tag{3.21}$$

where f represents the variable frequency, R is the number of points in the LPF, and f_s/R is the effective sampling rate of the integrator. This gives us:

$$\begin{aligned} K_I(f) &= \frac{K_I}{1 - \exp\left(-j2\pi\frac{Rf}{f_s}\right)} \\ &= \frac{K_I}{\left[\exp\left(j\pi\frac{Rf}{f_s}\right) - \exp\left(-j\pi\frac{Rf}{f_s}\right)\right]\exp\left(-j\pi\frac{Rf}{f_s}\right)} \\ &= \frac{K_I}{\left[2j\sin\left(\pi\frac{Rf}{f_s}\right)\right]\exp\left(j\pi\frac{Rf}{f_s}\right)} \\ &\approx \frac{K_I}{j2\pi\left(\frac{Rf}{f_s}\right)} \end{aligned} \tag{3.22}$$

where the small angle approximation has been used to simplify $\sin\left(\pi\frac{Rf}{f_s}\right) \approx \pi\frac{Rf}{f_s}$ and $\exp\left(j\pi\frac{Rf}{f_s}\right) \approx 1$ since we are interested in frequencies $Rf \ll f_s$. The transfer function of the overall controller is:

$$\begin{aligned} K(z) &= \left(K_I(z) + K_P\right)K_G \\ &= \left(\frac{K_I}{j2\pi\left(\frac{Rf}{f_s}\right)} + K_P\right)K_G \end{aligned} \tag{3.23}$$

The controller's gains K_P, K_I, and K_G are used to tune the gain of the PLL.

3.2.7 Frequency compensator

The frequency compensator updates the frequency of the LO model f_{LO} by adding the instantaneous frequency error f_{error} to the initial frequency estimate $f_{initial}$:

$$f_{LO} = f_{initial} + f_{error} \tag{3.24}$$

The frequency compensator has unity gain.

Representing frequency in the FPGA

Frequency f is represented in the FPGA as an integer (identified by capital F with units 'counts') according to

$$F = \text{floor}\left\{\frac{f}{f_s}2^D\right\} \tag{3.25}$$

where f represents the *true* frequency in Hertz, f_s is the sampling frequency of the discrete time system in Hertz, and 2^D is the size of the in-loop integrator used to convert frequency into phase (described in Section 3.2.8).

The smallest frequency $f_{\min}$ (when $F = 1$) that can be represented is:

$$f_{\min} = \frac{f_s}{2^D}$$

For example, if the sampling rate of the FPGA is $f_s = 40$ MHz and the desired frequency resolution is 100 mHz, then the exponent of the integrator must be at least:

$$\begin{aligned} D &= \text{ceil}\left\{\log_2\left(\frac{f_s}{f_{\min}}\right)\right\} \\ D &= \text{ceil}\left\{\log_2\left(\frac{40 \text{ MHz}}{100 \text{ mHz}}\right)\right\} \\ &= \text{ceil}\left\{28.58\right\} = 29 \text{ bits} \end{aligned}$$

3.2.8 In-loop integrator

The in-loop integrator accumulates F_{LO} at the full sampling rate f_s over a finite range, forcing it to 'wrap' periodically at f_{LO}, generating a signal that is proportional to the instantaneous phase of the local oscillator modulo 2π. True phase ϕ (in radians) is mapped into an integer phase Φ (in counts) via:

$$\Phi = \text{floor}\left\{\frac{\phi}{2\pi} 2^D\right\} \tag{3.26}$$

The size of the integrator 2^D is balanced against other parameters in the system to minimise quantisation noise and resource utilisation.

The in-loop integrator's transfer function is

$$\begin{aligned} P(z) &= \frac{G_P}{1 - z^{-1}} \\ &= \frac{G_P}{j2\pi\left(\frac{f}{f_s}\right)} \text{ for } f \ll f_s \end{aligned}$$

where G_P is a frequency independent gain, f is variable frequency, and f_s is the integrator's sampling rate.

The periodic phase signal Φ produced by the phase accumulator produces the address Λ used to access elements in the sine look-up table (LUT). If the number of entries in the LUT 2^L does not equal to the size of the phase accumulator 2^D, then Φ must be scaled by $G_P = 2^T = 2^{L-D}$:

$$\Lambda = \Phi G_P = \Phi 2^T = \Phi 2^{L-D}$$

When $T < 0$ the scaling operation is a division by a power of 2, which is performed in the FPGA by bit-shifting Φ to the right by T bits.

3.2.9 Numerically controlled oscillator

The numerically controlled oscillator generates the local oscillator signal using a LUT containing a single cycle of a digitised sine wave. Each address in the LUT corresponds

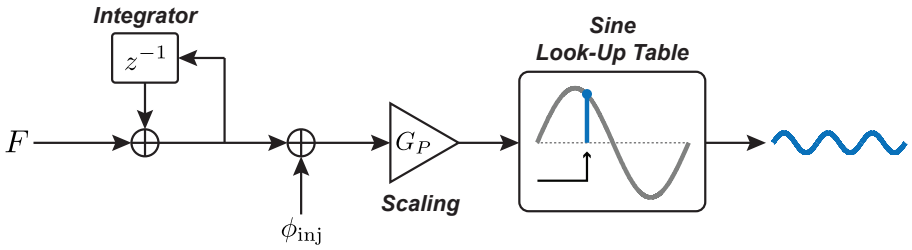

Figure 3.7: Block diagram of a numerically controlled oscillator. Frequency F is continuously integrated and scaled by the frequency independent gain parameter G_P to create an address used to access the elements of a sine look-up table. The phase of the local oscillator can be manipulated by injecting a phase signal ϕ_{inj} in to the output of the integrator.

to a particular phase of the sine wave between 0 and 2π. For a LUT with 2^L addresses, the address Λ is mapped to phase ϕ (in radians) by

$$\phi = 2\pi \frac{\Lambda}{2^L} = S\Lambda \tag{3.27}$$

where $S = \frac{2\pi}{2^L}$ is the effective gain of the NCO. Figure 3.7 shows the physical arrangement of the NCO.

3.2.10 Phase recovery

Phase error ϕ_{error} is recovered by:

1. integrating the frequency error F_{error} at a known sampling rate f_{out};
2. converting its units into radians using equation (3.25); and then
3. dividing by the rate at which F_{error} is integrated (f_{out}).

This is represented mathematically as:

$$\phi_{error} = \left(\frac{f_s}{2^D} \cdot \frac{1}{f_{out}} \right) \sum F_{error} \tag{3.28}$$

The rate at which F_{error} is integrated, f_{out}, relates to the FPGA sampling rate f_s by

$$f_{out} = \frac{f_s}{N}$$

where N is the total rate decimation prior to the out-of-loop integrator. Using this, equation (3.28) simplifies to:

$$\phi_{error} = \left(\frac{N}{2^D} \right) \sum F_{error} \tag{3.29}$$

Whilst the integration of f_{error} is performed on the FPGA, the scaling by $N/2^D$ should be performed using a CPU to calculate ϕ_{error} with floating point precision.

The transfer function of the phase recovery integrator is:

$$I(z) = \frac{N}{2^D} \cdot \frac{1}{1 - z^{-R}}$$

Dynamic range

The dynamic range (DR) of the PLL describes the ratio between the maximum and minimum measurable phase signals, often expressed in decibels:

$$\mathrm{DR}_{\mathrm{dB}} = \log_{10}\left(\frac{\phi_{max}}{\phi_{min}}\right)$$

The smallest signal the PLL can measure depends on its noise-floor (discussed in Section 3.5.3). The largest signal the can be measured by the PLL depends on the size of the out-of-loop integrator used to convert the frequency error into phase. If the size of the out-of-loop integrator is 2^P, then the maximum range of the integrated frequency error $\sum F_{error}$ is $\pm 2^{P-1}$. We can then use equation (3.29) to calculate the maximum range of the phase measurement:

$$\phi_{max} = \pm\left(\frac{N}{2^D}\right)2^{P-1}$$

The total measurement range is therefore:

$$-\left(N\frac{2^P}{2^{D+1}}\right) \leq \phi \leq \left(N\frac{2^P}{2^{D+1}}\right)$$

3.2.11 Using in-phase demodulation to measure amplitude

Assuming the PLL is locked and behaving linearly, it is possible to measure the amplitude of the input signal $s_{in}(t) = A\sin(\omega t + \phi_{in})$ by demodulating it using a second, in-phase local oscillator:

$$s_{LO,2}(t) = B\cos\left(\omega t + \phi + \frac{\pi}{2}\right) = B\sin(\omega t + \phi)$$

This requires the generation of a second LO that is locked 90 degrees out of phase with the primary LO used in the PLL, and the implementation of a second digital multiplier and LPF as shown in Figure 3.8.

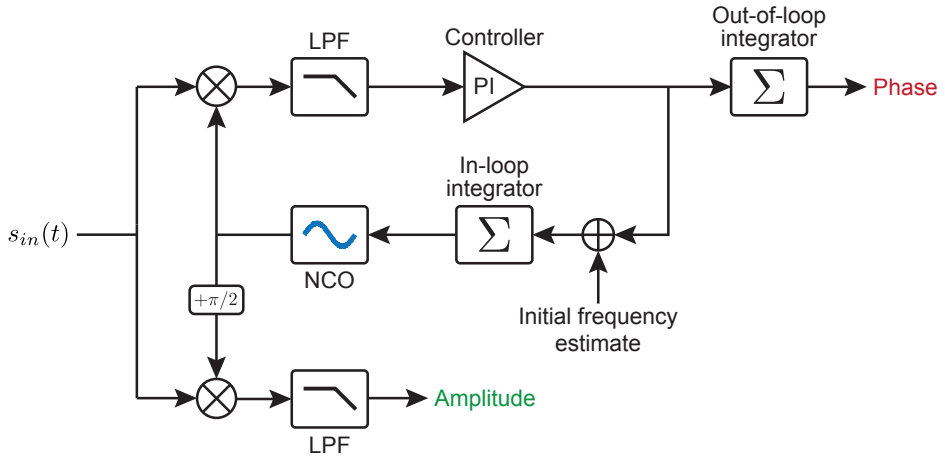

Figure 3.8: Modification to the physical architecture of the PLL enabling it to measure amplitude using a second out-of-loop demodulator and local oscillator orthogonal to the one used in the feedback loop.

Assuming $f \ll f_s$, the signal leaving the LPF is

$$s_{LPF}(t) = \frac{AB}{2} \cdot 2^{-Q} \cdot R$$

where $\frac{AB}{2}$ is the mixer gain, 2^{-Q} is the truncation at the output of the mixer, and R is the number of points in the LPF. The amplitude of the input signal is thus:

$$A = \frac{s_{LPF}(t) 2^{Q+1}}{B \cdot R}$$

3.3 Linear transfer function analysis

Phase-locked loops behave non-linearly due to their multiplicative feedback. Despite this, it is possible to analyse PLLs under the assumption it is operating linearly. By assuming the LO is phase-locked to the input signal, and that disturbances in the phase of the input signal are small and sufficiently slow[4], we are able to construct a linear transfer function model to investigate the PLL's measurement sensitivity, bandwidth, and stability.

The PLL's measurement sensitivity is limited by quantisation noise that is introduced whenever a signal is re-quantised by truncation[5]. Signal truncation is necessary to avoid excessive utilisation of FPGA resources, so it is crucial to understand exactly how it influences the PLL's performance. The purpose of this analysis is to understand how noise entering the loop at different locations affects the PLL's measurement sensitivity, and in the process develop an intuition for designing PLL's that are balanced in terms of performance and resource utilisation.

The linear transfer function model of the digitally implemented PLL is shown in Figure 3.9. For this analysis we will consider quantisation noise that is injected into the loop after the mixer, controller, and in-loop phase accumulator due to truncation, as well as amplitude quantisation noise introduced at both the ADC and NCO. The analysis will establish how to intuit the effect quantisation noise has on the phase measurement when introduced elsewhere in the loop (for example after the low-pass filter).

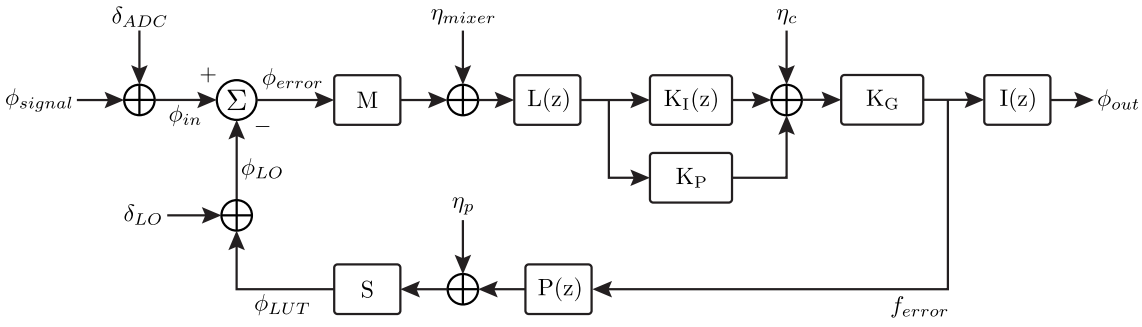

Figure 3.9: Linear transfer function model of the PLL shown in Figure 3.4. Noise is injected into the loop at various locations due to amplitude and phase quantisation noise. The linear approximation of the system is only valid when $\phi_{error} \ll 1$.

[4] Any phase disturbances that occur faster than the tracking bandwidth of the PLL cannot be tracked and will result in non-linear behaviour.

[5] Truncation is the division of a digital signal by a power of 2.

We'll begin our analysis by defining ϕ_{error}. We can see immediately from Figure 3.9 that

$$\phi_{error} = \phi_{in} - \phi_{LO} \tag{3.30}$$

$$\phi_{in} = \phi_{signal} + \delta_{ADC} \tag{3.31}$$

$$\phi_{LO} = S\big(\eta_p + P(z) f_{error}\big) + \delta_{LO} \tag{3.32}$$

and

$$f_{error} = K_G\big[\,(K_I(z) + K_P)\,L(z)\,(\eta_{mixer} + M\phi_{error}) + \eta_c\big] \tag{3.33}$$

Substituting Equation (3.33) into Equation (3.32) gives us

$$\phi_{LO} = S\eta_p + SP(z)\Big[K_G\big[\,(K_I(z) + K_P)\,L(z)\,(\eta_{mixer} + M\phi_{error}) + \eta_c\big]\Big] + \delta_{LO}$$

which, when substituted into Equation (3.30) yields

$$\phi_{error} = \phi_{in} - \Big[S\eta_p + SP(z)\Big[K_G\big[A(z)\big(\eta_{mixer} + M\phi_{error}\big) + \eta_c\big]\Big] + \delta_{LO}\Big] \tag{3.34}$$

where

$$A(z) = L(z)\big(K_I(z) + K_P\big) \tag{3.35}$$

and:

$$B(z) = SP(z) \tag{3.36}$$

Grouping the ϕ_{error} terms on the left-hand side and expanding the right-hand side gives us:

$$\phi_{error}\Big[1 + MA(z)K_G B(z)\Big] = \phi_{in} - S\eta_p - B(z)A(z)\eta_{mixer} - B(z)\eta_c - \delta_{LO} \tag{3.37}$$

The $MA(z)K_G B(z)$ component on the left hand side of Equation (3.37) is called the *open-loop gain*, $G_{OL}(z)$:

$$G_{OL}(z) = ML(z)K_G\big(K_I(z) + K_P\big)P(z)S \tag{3.38}$$

We can rearrange equation (3.37) to reveal the effect each signal entering the loop has on ϕ_{error}:

$$\phi_{error} = \frac{\phi_{in} - S\eta_p - A(z)K_G B(z)\eta_{mixer} - K_G B(z)\eta_c - \delta_{LO}}{1 + G_{OL}(z)} \tag{3.39}$$

Equation (3.39) tells us how the loop influences ϕ_{error}, but we are more interested in the effect the loop has on the output measurement ϕ_{out}. Substituting ϕ_{error} into equation (3.33) and multiplying by the out-of-loop integrator $I(z)$ gives us:

$$\begin{aligned}
\phi_{out} &= I(z) f_{error} \\
&= I(z)K_G\Big[A(z)\big(\eta_{mixer} + M\phi_{error}\big) + \eta_c\Big] \\
&= I(z)K_G\left[A(z)\left(\eta_{mixer} + M\frac{\phi_{in} - S\eta_p - C(z)\eta_{mixer} - K_G B(z)\eta_c - \delta_{LO}}{1 + G_{OL}(z)}\right) + \eta_c\right]
\end{aligned} \tag{3.40}$$

where the substitution $C(z) = A(z)K_G B(z)$ has been made for clarity. The noise contributions η_c and η_{mixer} can be grouped and simplified further since

$$\begin{aligned} A(z)\eta_{mixer} - \frac{MA(z)^2 K_G B(z)}{1+G_{OL}}\eta_{mixer} &= A(z)\left(\frac{K_G + G_{OL} - G_{OL}}{1+G_{OL}}\right)\eta_{mixer} \\ &= \frac{A(z)K_G}{1+G_{OL}}\eta_{mixer} \end{aligned}$$

and

$$\begin{aligned} \eta_c - \frac{MA(z)K_G B(z)}{1+G_{OL}(z)}\eta_c &= \frac{1+G_{OL}(z)-G_{OL}(z)}{1+G_{OL}(z)}\eta_c \\ &= \frac{K_G}{1+G_{OL}(z)}\eta_c \end{aligned}$$

Condensing equation (3.40) and rearranging yields:

$$\phi_{out} = I(z)K_G\left[A(z)M\left(\frac{\phi_{in} - S\eta_p - \delta_{LO}}{1+G_{OL}(z)}\right) + \frac{A(z)}{1+G_{OL}(z)}\eta_{mixer} + \frac{1}{1+G_{OL}(z)}\eta_c\right] \tag{3.41}$$

Equation (3.41) describes the forward loop gain of each signal entering the loop, and tells us how they influence the output measurement ϕ_{out}. The individual transfer functions from each signal entering the loop to ϕ_{out} are:

$$\phi_{out}\Big|_{\phi_{in}} = \frac{I(z)ML(z)\left(K_I(z)+K_P\right)K_G}{1+G_{OL}(z)}\phi_{in} \tag{3.42}$$

$$\phi_{out}\Big|_{\eta_{mixer}} = \frac{I(z)L(z)\left(K_I(z)+K_P\right)K_G}{1+G_{OL}(z)}\eta_{mixer} \tag{3.43}$$

$$\phi_{out}\Big|_{\eta_c} = \frac{I(z)K_G}{1+G_{OL}(z)}\eta_c \tag{3.44}$$

$$\phi_{out}\Big|_{\eta_p} = -\frac{I(z)ML(z)\left(K_I(z)+K_P\right)K_G S}{1+G_{OL}(z)}\eta_p \tag{3.45}$$

$$\phi_{out}\Big|_{\delta_{LO}} = -\frac{I(z)L(z)\left(K_I(z)+K_P\right)K_G M}{1+G_{OL}(z)}\delta_{LO} \tag{3.46}$$

In the high gain limit when $G_{OL}(z) \gg 1$ we can approximate the denominator as

$$1 + G_{OL}(z) \approx G_{OL}(z)$$

allowing us to further simplify equations (3.42) to (3.46):

$$\phi_{out}\Big|_{\phi_{in}} = \frac{I(z)}{P(z)S}\phi_{in} \tag{3.47}$$

$$\phi_{out}\Big|_{\eta_{mixer}} = \frac{I(z)}{MP(z)S}\eta_{mixer} \tag{3.48}$$

$$\phi_{out}\Big|_{\eta_c} = \frac{I(z)}{ML(z)\left(K_I(z)+K_P\right)SP(z)}\eta_c \tag{3.49}$$

$$\phi_{out}\Big|_{\eta_p} = -\frac{I(z)}{P(z)}\eta_p \tag{3.50}$$

$$\phi_{out}\Big|_{\delta_{LO}} = -\frac{I(z)}{P(z)S}\delta_{LO} \tag{3.51}$$

We can see from equations (3.47) to (3.51) that the worst location to truncate is at the output of the in-loop phase accumulator where η_p is suppressed only by $P(z)$. The best location to truncate is at the output of the controller where η_c is suppressed by the full open-loop gain. This is somewhat counter intuitive, but it does highlight an important lesson: any signal injected into the loop is suppressed by all transfer functions that precede it in the loop on the way back to measurement point. Noise introduced at the mixer η_{mixer} is preceded by M, S, and $P(z)$ on the way back through the loop to ϕ_{out}. We can therefore infer that any noise introduced at the output of the low-pass filter η_{LPF} would be suppressed by $L(z)$, M, S, and $P(z)$:

$$\phi_{out}\Big|_{\eta_{LPF}} = \frac{I(z)}{L(z)MSP(z)}\eta_{LPF} \tag{3.52}$$

Now that we have expressions that define what kind of influence quantisation noise has on the output measurement ϕ_{out}, it is necessary to quantify it.

3.4 Quantisation noise

Quantisation noise is a consequence of the discrete resolution of digital signals and is related to a signal's least-significant bit (LSB). Quantisation noise is added to the system when an analogue signal is digitised (at an analogue-to-digital converter (ADC) for example), or whenever a digital signal is truncated (or re-quantised). Quantisation noise can generally be assumed to be purely random and uniformly distributed (spectrally white) across one LSB. This is an idealisation for the purposes of analysis, so it should be noted that quantisation noise is not always necessarily white[6], an example being digitised sine waves.

Let x be a random variable representing the LSB of a signal. The standard deviation of x is

$$\sigma = \sqrt{\int_{-\infty}^{\infty} \rho(x)(x-\bar{x})^2\, dx} \tag{3.53}$$

where $\bar{x}$ represents the statistical mean of x and $\rho(x)$ is its probability density function. If the quantisation noise is uniformly distributed across $-\text{LSB}/2$ and $\text{LSB}/2$ then the PDF is described as:

$$\rho(x) = \begin{cases} \frac{1}{\text{LSB}} & \text{for } -\frac{\text{LSB}}{2} \geq x \geq \frac{\text{LSB}}{2} \\ 0 & \text{elsewhere} \end{cases} \tag{3.54}$$

then the statistical mean of $\rho(x)$ is $\bar{x} = 0$. The standard deviation is thus

$$\begin{aligned} \sigma &= \sqrt{\int_{-\text{LSB}/2}^{\text{LSB}/2} \frac{1}{\text{LSB}} x^2\, dx} \\ &= \sqrt{\frac{1}{\text{LSB}} \frac{x^3}{3}\Big|_{-\text{LSB}/2}^{\text{LSB}/2}} \\ &= \sqrt{\frac{1}{\text{LSB}} \left[\frac{\text{LSB}^3}{24} + \frac{\text{LSB}^3}{24}\right]} \\ &= \frac{\text{LSB}}{\sqrt{12}} \end{aligned}$$

[6] Quantisation can be forced white by dithering the LSB with pseudo-random noise.

and the variance, σ^2, is:

$$\sigma^2 = \frac{\text{LSB}^2}{12}$$

The variance of random noise digitised at a sampling frequency f_s is related to the single-sided power spectral density (PSD) $S_x(f)$ by:

$$\sigma^2 = \int_0^{f_s/2} S_x(f)df$$

If $S_x(f)$ is constant (spectrally white) then the relationship becomes

$$\sigma^2 = \frac{S_x(f)f_s}{2} \tag{3.55}$$

The single-sided PSD of uniformly distributed quantisation noise is thus:

$$S_x(f) = \frac{2\sigma^2}{f_s} = \frac{\text{LSB}^2}{6f_s} \tag{3.56}$$

A more useful representation is the root power spectral density (RPSD), $\tilde{S}_x(f)$, which is normalised to a 1 second averaging window and is proportional to amplitude:

$$\tilde{S}_x(f) = \frac{\text{LSB}}{\sqrt{6f_s}}$$

The amplitude normalised RPSD of quantisation noise for a signed N-bit signal is:

$$\tilde{S}_x(f) = \frac{2^{-(\text{N}-1)}}{\sqrt{6f_s}}$$

For an unsigned N-bit integer it is:

$$\tilde{S}_x(f) = \frac{2^{-\text{N}}}{\sqrt{6f_s}}$$

3.4.1 The effect of amplitude quantisation noise on phase

Figure 3.10 illustrates how amplitude quantisation noise affecting the input and local oscillator signals translates into phase noise. Rotations of $s(t)$ in a static reference frame correspond to changes in the signal's phase, ϕ_{signal}, and any noise η produces random phase fluctuations $\delta\phi$ around ϕ_{signal}. The extent to which $\delta\phi$ corrupts the true phase of $s(t)$ depends on the relative magnitudes of $s(t)$ and η and their orientation. Phase noise $\delta\phi$ is greatest when η is perpendicular to $s(t)$, and assuming the amplitude of the signal is much greater than that of the noise, we can infer that:

$$\delta\phi = \tan^{-1}\left(\frac{\eta}{|s(t)|}\right) \approx \frac{\eta}{|s(t)|} \text{ for } \eta \ll s(t)$$

Amplitude quantisation noise introduced by the ADC and NCO is transformed into phase noise at the PLL's digital multiplier. The inputs to the digital multiplier are

$$s_{in}(t) = A\sin(\omega t + \phi_{in}) + \eta_{in} \tag{3.57}$$

$$s_{lo}(t) = B\cos(\omega t + \phi_{lo}) + \eta_{lo} \tag{3.58}$$

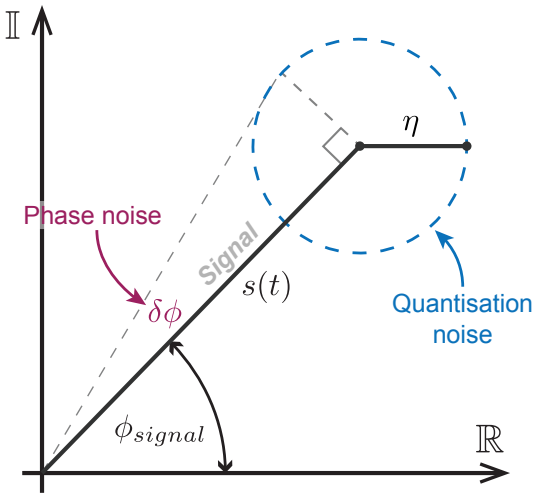

Figure 3.10: Phasor diagram illustrating the effect that additive quantisation noise η has on the phase of an oscillating signal $s(t)$. Phase noise $\delta\phi$ is greatest when η is orthogonal to $s(t)$, and lowest when it is parallel. In this figure the magnitude of η relative to $s(t)$ has been greatly exaggerated.

with amplitudes A and B; angular frequency ω; phase ϕ_{in} and ϕ_{lo}; and additive quantisation noise η_{in} and η_{lo}. The output of the digital multiplier $s_M(t)$ is:

$$\begin{aligned} s_M(t) &= s_{in}(t)s_{lo}(t) \\ &= \big(A\sin(\omega t+\phi_{in})+\eta_{in}\big)\big(B\cos(\omega t+\phi_{lo})+\eta_{lo}\big) \\ &= AB\sin(\omega t+\phi_{in})\cos(\omega t+\phi_{lo}) + B\eta_{in}\cos(\omega t+\phi_{lo}) + A\eta_{lo}\sin(\omega t+\phi_{in}) + \eta_{in}\eta_{lo} \\ &= AB\left(\frac{\sin(\phi_{in}-\phi_{lo})+s_{2\omega}(t)}{2} + \frac{\eta_{in}}{A}\cos(\omega t+\phi_{lo}) + \frac{\eta_{lo}}{B}\sin(\omega t+\phi_{in}) + \frac{\eta_{in}\eta_{lo}}{AB}\right) \end{aligned} \tag{3.59}$$

The second harmonic $s_{2\omega}(t)$ can be ignored as it will be strongly attenuated by the PLL's low-pass filter. We can also safely assume that η_{in} and η_{lo} are much less than A and B, which means the final term $\frac{\eta_{in}\eta_{lo}}{AB}$ will have negligible impact on $s_M(t)$ and can therefore be discarded. Under the assumption that $\Delta\phi = \phi_{in} - \phi_{lo} \ll 1$ we can apply the small angle approximation for sine to further simplify equation (3.59), giving us:

$$s_M(t) = \frac{AB}{2}\left(\Delta\phi + \frac{2\eta_{in}}{A}\cos(\omega t+\phi_{lo}) + \frac{2\eta_{lo}}{B}\sin(\omega t+\phi_{in})\right) \tag{3.60}$$

The remaining sine and cosine components in equation (3.60) can approximated by their RMS amplitudes, $\frac{1}{\sqrt{2}}$, since the noise is quantified as a standard deviation:

$$s_M(t) = \frac{AB}{2}\left(\Delta\phi + \frac{2\eta_{in}}{\sqrt{2}A} + \frac{2\eta_{lo}}{\sqrt{2}B}\right) \tag{3.61}$$

The amplitude quantisation of the input and LO signals $s_{in}(t)$ and $s_{lo}(t)$ can thus be represented in the linear TF model as:

$$\begin{aligned} \delta_{in} &= \frac{2\eta_{in}}{\sqrt{2}A} \\ \delta_{lo} &= -\frac{2\eta_{lo}}{\sqrt{2}B} \end{aligned} \tag{3.62}$$

3.5 Modelling and simulation

The linear transfer function model of the PLL was simulated in MATLAB using the z-domain transfer functions described in Section 3.2. The model represents the PLL architecture used in all of the OPA experiments. The only variables in the model are the amplitude of the input signal, A, integral gain, K_I, and global gain, K_G. The transfer functions are summarised below:

$$
\begin{aligned}
M &= \frac{AB}{2} \cdot 2^{-Q} \\
L(z) &= \left(\frac{1 - z^{-R}}{1 - z^{-1}} \right)^2 \\
P(z) &= 2^{L-D} \cdot \frac{1}{1 - z^{-1}} \\
S &= \frac{2\pi}{2^L} \\
I(z) &= \frac{R}{2^D} \cdot \frac{1}{1 - z^{-R}} \\
K_I(z) &= K_I \cdot \frac{1}{1 - z^{-R}}
\end{aligned}
$$

In this phasemeter the output of the digital multiplier is truncated by $M = 15$ bits; the number of points in the LPF is $R = 40$; the NCO has $2^L = 2^{12}$ elements, and the size of the in-loop phase accumulator is $D = 32$ bits. The amplitude of the input signal was set nominally at a value at half the full range of the 16-bit ADC: $A = 2^{14}$. The amplitude of the NCO is $B = 2^{15}$.

3.5.1 Open-loop gain

An example trace of the PLL's open-loop gain $G_{OL}(z)$ is shown in Figure 3.11 with a unity gain frequency of 63.3 kHz and a phase margin of 63.7 degrees. The transfer function in the region of the unity gain frequency is proportional to the inverse Fourier frequency, which means PLL bandwidth in this region will scale proportionally with K_G and M.

If the amplitude of the input signal A is reduced by some arbitrary factor, then $G_{OL}(z)$ can be compensated by increasing K_G by the same factor. A global gain of 33 results in a unity gain frequency of 63.3 kHz when the input signal is at half the full range of the ADC. This gain can be used as a reference $K_G^\dagger$ from which to automatically adjust the gain of the loop to maintain a desired bandwidth. A unity gain frequency of 100 kHz would be achieved by multiplying the reference gain $K_G^\dagger = 33$ by $100/63.3 = 1.58$, resulting in an updated global gain of $K_G = 52$. This is only applicable when the transfer function crosses unity gain with a $1/f$ profile.

According to equation (3.62), however, the phase noise caused by amplitude quantisation at the ADC depends on the amplitude of the input signal A. This makes it more appropriate to instead maximise A to increase the UGF by increasing $M = \frac{AB}{2}$.

3.5.2 Selecting appropriate controller gains

The process of selecting appropriate values for K_P, K_I, and K_G begins with the consideration of the PLL's open-loop gain $G_{OL}(z)$ when $K_P = 1$, $K_I = 0$, and $A = 2^{C-2}$ (where C is the resolution of the ADC). When $K_I = 0$ the open-loop gain is dominated by a single

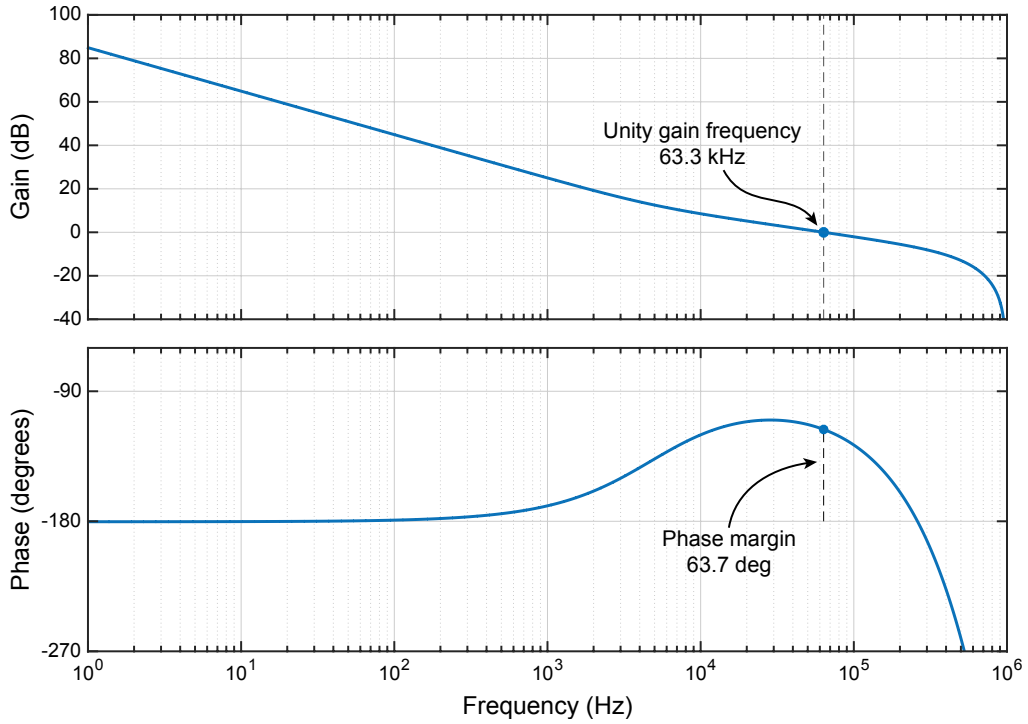

Figure 3.11: Open loop gain of the PLL.

integrator $P(z)$ with a $1/f$ profile within the passband of the LPF. The PLL's unity gain frequency (bandwidth) can then be controlled by adjusting K_G.

The integral and proportional gains are then selected based on the frequency at which their individual transfer functions intersect in order to maximise gain at low frequencies (purpose of the integrator), without causing the PLL to become unstable at high frequencies. This comes down to ensuring that $G_{OL}(z)$ has a phase margin of at least 60 degrees at the unity gain frequency [81].

K_P and $K_I(z)$ intersect when:

$$K_P = \frac{K_I}{j2\pi\left(\frac{Rf}{f_s}\right)} \tag{3.63}$$

What matters here is the ratio between K_I and K_P, so to simplify things we set $K_P = 1$ and focus only on K_G:

$$1 = \frac{K_I}{j2\pi\left(\frac{Rf}{f_s}\right)}$$

The frequency at which the integrator intersects unity gain is therefore

$$\therefore f_{\text{int}} = \frac{K_I f_s}{j2\pi R} \tag{3.64}$$

which is important because of the additional 90 degree phase delay introduced to the loop whenever $K_I(z) > K_P$. The integrator's 90 degree phase delay combined with that from the phase accumulator results in a total phase delay approaching 180 degrees at frequencies $f < f_{\text{int}}$. This is shown in Figure 3.12.

As a general rule, the frequency f_{int} should be set approximately one or two orders of magnitude lower the PLL bandwidth to ensure a phase margin of at least 60 degrees. A

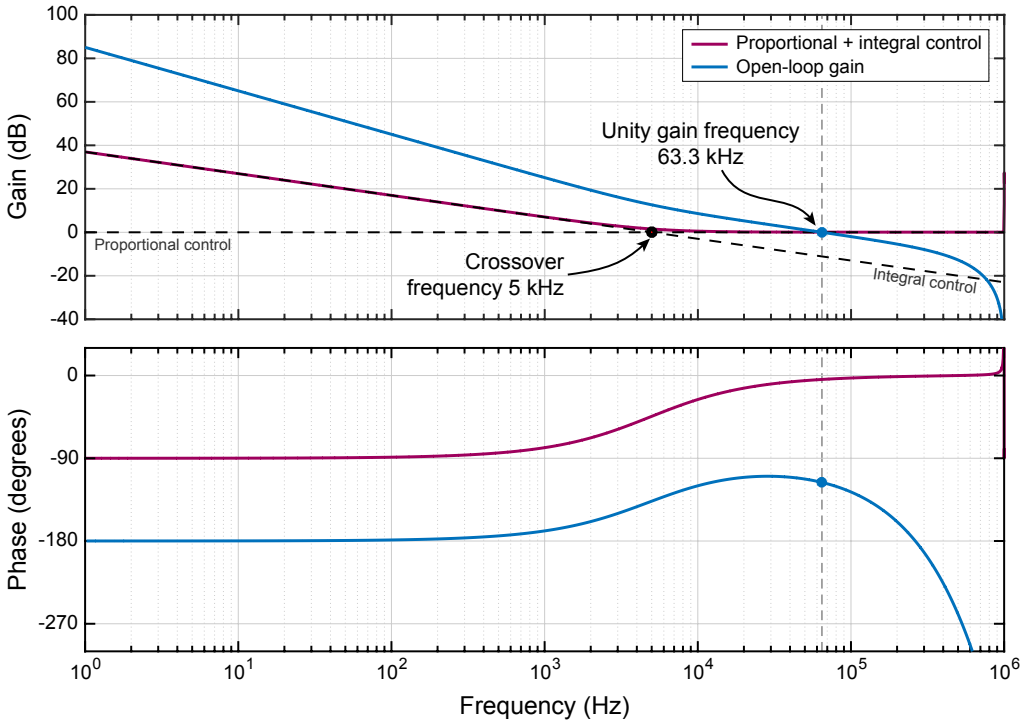

Figure 3.12: Relationship between $K_I(z) + K_P$ (identified individually in red, and combined in blue) and the open-loop gain (orange). The frequency at which $K_I(z)$ and K_P intersect is 50 times lower than the unity gain frequency to ensure a phase margin greater than 60 degrees.

safe option for a PLL bandwidth of 66.3 kHz is $f_{\text{int}} \approx 5$ kHz. Assuming a sampling rate $f_s = 40$ MHz, decimation factor $R = 40$, and ignoring the imaginary phase delay $j = e^{j\frac{\pi}{2}}$, an appropriate integral gain is

$$\begin{aligned} K_I &= \left| \frac{j2\pi R f_{\text{int}}}{f_s} \right| \\ &= \left| \frac{j2\pi \cdot 40 \cdot 10^4}{40 \times 10^6} \right| \\ &= 0.0314 \approx 2^{-5} = 0.0312 \end{aligned} \tag{3.65}$$

The integral gain is thus a division by a power of two, and is handled easily in the FPGA by bit-shifting the output of the integrator to the right by 5 bits. This re-quantisation of the integrator's output explains why quantisation noise is introduced at the output of the controller.

3.5.3 Quantisation noise-floor

Equations (3.42) to (3.46) describe how quantisation noise affects the PLL's phase measurement. The input and NCO signals are amplitude quantised, and the signal within the PLL is truncated at the output of the mixer, controller, and phase accumulator. The forward loop contributions of each source of quantisation noise are shown in Figure 3.13 along with the total combined noise.

The simulated quantisation noise levels in Figure 3.13 belongs to the specific configuration described earlier in this section. The total estimated quantisation noise-floor of 30.2

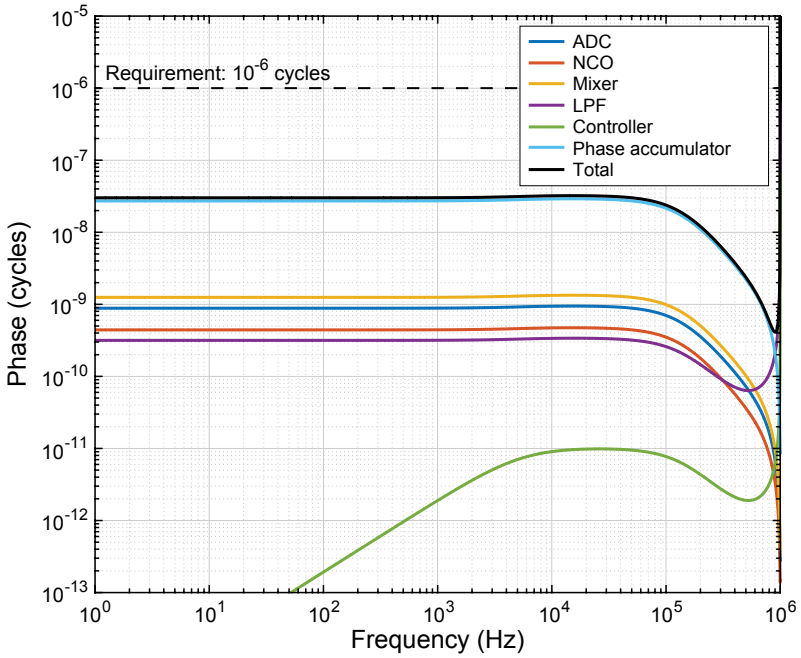

Figure 3.13: The simulated quantisation noise-floor of the PLL is approximately 30.2 nano-cycles/$\sqrt{\text{Hz}}$, limited primarily by re-quantisation after the phase accumulator. The noise-floor can change depending on the specific gains settings, so the result presented here belongs to the specific transfer functions and gains presented in this section.

nano-cycles is limited by quantisation noise introduced at the output of the phase accumulator, which is truncated by $S - D = -20$ bits. This is the most severe re-quantisation in the PLL, and is only suppressed by the gain of the phase accumulator $P(z)$:

$$\begin{aligned} \phi_{out}\Big|_{\eta_p} &= -\frac{I(z)}{P(z)}\eta_p \\ &= -\frac{R\left(1-z^{-1}\right)}{2^S\left(1-z^{-R}\right)}\eta_p \end{aligned} \tag{3.66}$$

Equation (3.66) states that the η_p can be suppressed by increasing the number of elements in the NCO (2^L), and increasing overall FPGA resource utilisation; or by reducing the number of points in the LPF, R. It is unnecessary to change either in this case since the noise-floor of 30.2 nano-cycles exceeds the design requirement of 1 μcycle by almost two orders of magnitude.

It is important to reiterate that the noise-floor of the PLL is sensitive to changes in gain—especially that of the mixer. A reduction in the amplitude of the input signal reduces the mixer gain $M = \frac{AB}{2}$ by the same amount. This will reduce the loop's suppression of any truncation noise introduced after the mixer (equations (3.48) and (3.49)), and amplify the effect of amplitude quantisation noise introduced by the ADC (equation (3.57)).

Quantisation noise introduced after the controller (green trace) decreases towards lower frequencies because it is suppressed by both integrators in the loop.

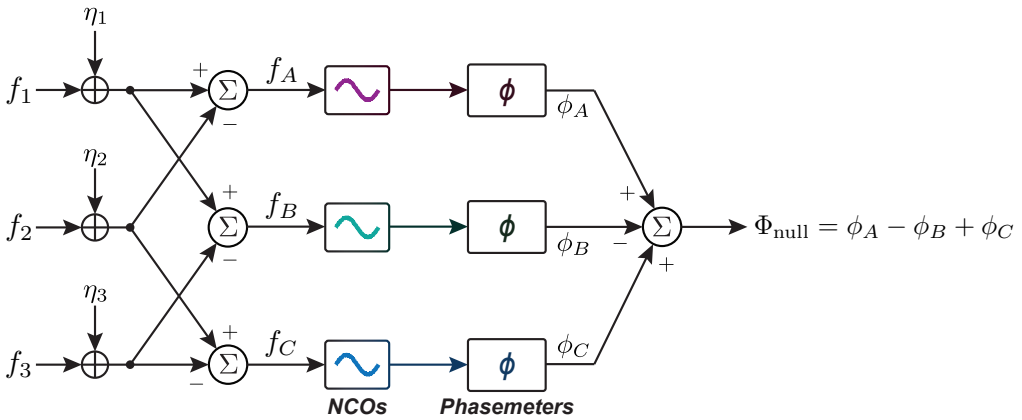

Figure 3.14: Block diagram of the three noise test implemented on an FPGA. Three noisy seed frequencies $\tilde{f}_1$, $\tilde{f}_2$, and $\tilde{f}_3$ are combined to produce three noise coupled signals at frequencies f_A, f_B, and f_C using NCOs. The phase of these three signals are measured using independent but identical and clock-synchronised PLLs, and then combined in such a way that $\phi_A - \phi_B + \phi_C = 0$.

3.6 Three noise test

The three-noise test (shown schematically in Figure 3.14) is used to measure the PLLs sensitivity and to reveal non-linearities (e.g., aliased second harmonics) in the measurement band. The test is performed by generating three signals in such a way that the linear combination of their phases is zero:

$$\phi_A - \phi_B + \phi_C = 0 \tag{3.67}$$

Any non-linearities or noise uncommon to the three signals (e.g., quantisation noise introduced by the PLL) will not cancel to zero and will remain in the combined measurement.

The three-noise test presented here is performed digitally using signals generated within the FPGA. Whilst the test can be performed using externally generated analogue signals (e.g., using the heterodyne beat notes from a photo-detector), it would require three independent laser sources and a single FPGA that can support three ADCs. The two FPGA platforms used in this system can support only two ADC inputs each, ruling out the possibility of performing the three-noise test under realistic conditions.

The generation of the three test signals begins with three noisy frequencies $\tilde{f}_1$, $\tilde{f}_2$, and $\tilde{f}_3$

$$\begin{aligned}\tilde{f}_1 &= f_1 + \eta_1 \\ \tilde{f}_2 &= f_2 + \eta_2 \\ \tilde{f}_3 &= f_3 + \eta_3\end{aligned}$$

where f_1, f_2, and f_3 represent constant frequencies and η_1, η_2, and η_3 represent random noise. Three noise coupled frequencies f_A, f_B, and f_C are then defined by combining $\tilde{f}_1$, $\tilde{f}_2$, and $\tilde{f}_3$

$$\begin{aligned}f_A &= \tilde{f}_1 - \tilde{f}_2 \\ f_B &= \tilde{f}_1 - \tilde{f}_3 \\ f_C &= \tilde{f}_2 - \tilde{f}_3\end{aligned}$$

such that:

$$\begin{aligned} f_{\text{sum}} &= f_A - f_B + f_C \\ &= (\tilde{f}_1 - \tilde{f}_2) - (\tilde{f}_1 - \tilde{f}_3) + (\tilde{f}_2 - \tilde{f}_3) \\ &= 0 \end{aligned}$$

Numerically controlled oscillators are used to generate signals at frequencies f_A, f_B, and f_C with amplitude B:

$$\begin{aligned} s_A(t) &= B_A \sin(2\pi f_A t) + \\ &= B_A \sin\left(2\pi \left(f_1 - f_2 + \eta_1 - \eta_2\right) t\right) \end{aligned} \tag{3.68}$$

Three independent but identical and clock-synchronised PLLs are used to measure the phase error of each signal (which is the time integral of the frequency error, φ)

$$\begin{aligned} \sum^{\infty} f_A &= \sum^{\infty} (\eta_1 - \eta_2) \\ &= \varphi_1 - \varphi_2 \end{aligned}$$

plus any uncorrelated noise and/or non-linearities (e.g., quantisation noise, additive white Gaussian noise) introduced at the PLL, ϑ_A, and NCO, ψ_A. The phase measurement is therefore:

$$\phi_A = \varphi_1 - \varphi_2 + \vartheta_A + \psi_A$$

Combining the three phase measurements described in equation (3.67) eliminates any common phase signals, leaving only the linear combination of uncorrelated noise introduced by the NCO and PLLs:

$$\begin{aligned} \Phi_{\text{null}} &= \phi_A - \phi_B + \phi_C \\ &= (\varphi_1 - \varphi_2 + \vartheta_A + \psi_A) - (\varphi_1 - \varphi_3 + \vartheta_B + \psi_B) + (\varphi_2 - \varphi_3 + \vartheta_C + \psi_C) \\ &= (\vartheta_A + \psi_A) - (\vartheta_B + \psi_B) + (\vartheta_C + \psi_C) \end{aligned}$$

If the three PLLs are identical (i.e., have same transfer functions), then the RPSDs of their quantisation noise-floors, $\tilde{\Theta}$, will also also be identical and equal to the sum of all the forward loop quantisation noise contributions described in Section 3.5.3.

From the analysis in Section 3.5.3 we know that the dominant source of quantisation noise in the PLL is introduced at the output of the phase accumulator just before the PLL's NCO, η_p. The three input signals used in the three noise test are generated using the same phase accumulators and NCOs as those in the PLL, which means they will experience quantisation noise η_{sg} with the same magnitude, η_p. The effect this has on the PLL's phase measurement is the same as that of the PLL's phase accumulator since the magnitude of their forward loop transfer functions are identical:

$$\begin{aligned} \phi_{out}\Big|_{\eta_{sg}} &= \frac{I(z)ML(z)\left(K_I(z) + K_P\right)K_G S}{1 + G_{OL}(z)}\eta_{sg} \\ \phi_{out}\Big|_{\eta_p} &= \frac{I(z)ML(z)\left(K_I(z) + K_P\right)K_G S}{1 + G_{OL}(z)}\eta_p \end{aligned}$$

We can see from Figure 3.13 that the PLL's overall quantisation noise-floor is approximately equal to the forward loop quantisation noise of the phase accumulator,

$$\tilde{\Theta} \approx \phi_{out}\Big|_{\eta_p}$$

which is equal to the RPSD of the quantisation noise of the three input signals. This means that the RPSD of the null measurement will be:

$$\tilde{\Phi}_{\text{null}} \approx \tilde{\Theta}\sqrt{6}$$

where the RPSD of the linear combination of N incoherent and identical noise sources scales with $\sqrt{N}$ (since the combined noise power scales with N). The quantisation noise-floor of each PLL is then approximately:

$$\tilde{\Theta} = \frac{\Phi_{\text{null}}}{\sqrt{6}} \tag{3.69}$$

3.7 Simulating noise on the FPGA

The aim of the three-noise test is to investigate how the PLL behaves under realistic conditions. Physical noise sources likely to affect the input to the PLL have Gaussian probability distribution (e.g., shot noise and laser frequency noise), which can be simulated in the FPGA using pseudo-random noise. The specific type of PRN codes used in this system are called maximal-length sequences, which can be generated efficiently using linear feedback shift registers (as discussed in Section 2.3.1).

3.7.1 Uniform noise distribution

An N-bit LFSR will generate every value between 1 and $2^N - 1$ in a random order, producing uniformly distributed random noise with length $2^N - 1$. At a sampling rate f_s the code will repeat every $T_{\text{code}} = (2^N - 1)/fs$ seconds, producing harmonics at integer multiples of the code repetition frequency $f_{\text{code}} = 1/T_{\text{code}}$.

The LFSRs used to generate the noise for the three-noise test are 64-bits long, which at a sampling rate of 40 MHz will repeat every 14.5 thousand years.

A single 32-bit LFSR (truncated to 30-bits) was used to generate 500 thousand samples of uniformly distributed noise on an FPGA. The histogram and expected theoretical probability density function (p.d.f.) of the noise is shown in Figure 3.15(a). The noise has been interpreted as a signed 30-bit integer, which has a range of -2^{29} to $2^{29} - 1$.

3.7.2 Triangular noise distribution

A triangular noise distribution can be generated by adding or subtracting the outputs of two LFSRs with different feedback configurations [82]. Where the probability distribution of the single LFSR is analogous to rolling a single dice, the distribution of the sum (or difference) of two LFSRs is the same as rolling two die.

A 500 thousand sample histogram of the noise generated by subtracting the outputs of two separate 30-bit LFSRs is shown in Figure 3.15(b). The predicted triangular p.d.f. is also shown.

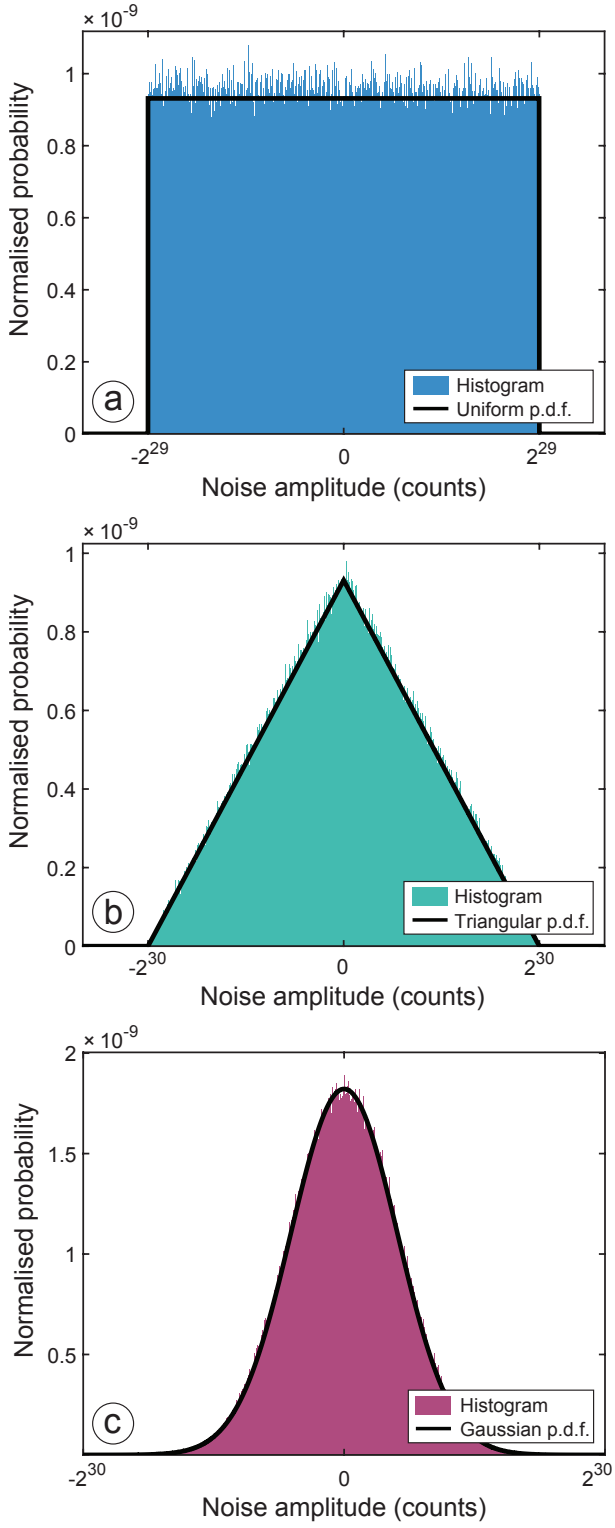

Figure 3.15: (a) Histogram (blue) and PDF (black) of uniformly distributed random noise generated by a 30-bit LFSR implemented on an FPGA. (b) Histogram (green) and PDF (black) of triangularly distributed random noise generated by linearly combining (adding or subtracting) the noise generated by two different LFSRs implemented on an FPGA. (c) Histogram (magenta) and PDF (black) of Gaussian-like distributed random noise generated by linearly combining the noise generated by eight separate LFSRs implemented on an FPGA.

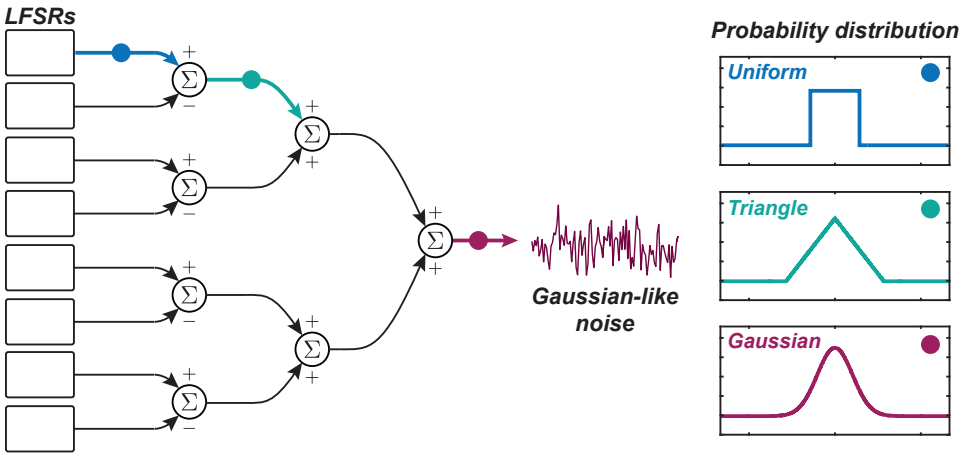

Figure 3.16: Block diagram showing how noise with uniform (blue), triangle (green), and Gaussian (magenta) probability distributions is generated in the FPGA. Each LFSR has a different feedback configuration and initial seed to ensure their combined outputs are random.

3.7.3 Gaussian-like noise distribution

The central limit theorem states that the distribution of the sum (or average) of a large number of independent, identically distributed random variables will be approximately Gaussian, regardless of the underlying distribution of the random variables [83]. This means that Gaussian noise can be created by summing the outputs of 'a large number' of LFSRs. It was found through trial and error that summing eight separately configured LFSRs generates Gaussian-like noise to a good approximation.

Approximately white Gaussian noise was generated by summing the outputs of eight 64-bit LFSRs, each truncated to 28-bits to produce a single 31-bit noise signal[7]. A 500 thousand sample histogram and theoretical p.d.f. of the generated Gaussian noise in Figure 3.15(c).

3.7.4 Block diagram of noise generator

A schematic of how the noise generator can be implemented on an FPGA is shown in Figure 3.16. Individually configured LFSRs produce independent but identically distributed (i.i.d.) pseudo-random noise with a uniform probability distribution (red). Subtracting the output of one LFSR from another produces noise with a triangle probability distribution (green). The combination of many LFSRs produces noise with a Gaussian probability distribution.

3.7.5 Spectral characteristics of pseudo-random noise

The power spectrum of pseudo-random noise has a sinc^2 envelope with nulls at integer multiples of the PRN chip frequency, f_{chip}, and is only approximately white at frequencies much lower than f_{chip}. An RPSD of 31-bit Gaussian-like noise generated at $f_s = 40$ MHz and translated into frequency noise using equation (3.25) is shown in Figure 3.17, revealing a noise shelf with an amplitude of $\sim$790 Hz/$\sqrt{\text{Hz}}$ at frequencies less than 1 MHz. The

[7]The noise was truncated to 31-bits to avoid potential overflow when combined with a second noise source of equal size, which in this case will produce 32-bit noise. Truncating the 64-bit noise does not affect the PRN repetition rate—it merely scales the noise's amplitude.

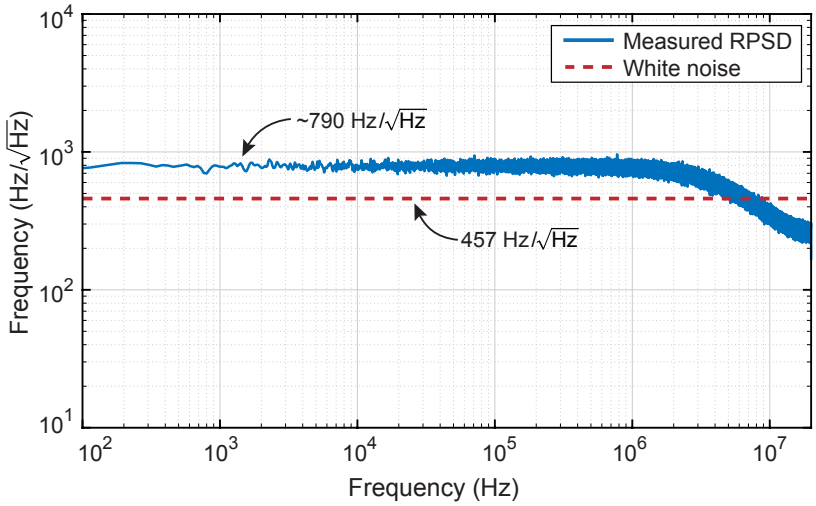

Figure 3.17: RPSD of Gaussian noise generated on the FPGA. The spectrum is not white across the full noise bandwidth, which means the theoretical RPSD cannot be calculated under the assumption that it is spectrally white. A more reliable estimate of the PRN's RPSD is $\sqrt{3}$ times the estimate provided under the assumption that the spectrum is white.

noise amplitude declines as the frequency approaches Nyquist and the spectral null of the Sinc^2 envelope.

It was stated in equation (3.55) that the variance σ^2 of spectrally white noise with a bandwidth f_s is related to its single sided PSD $S_x(f)$ by

$$\sigma^2 = \int_0^{fs/2} S_x(f)df$$

which, under the assumption that $S_x(f)$ is constant, provides us with an equation to estimate the noise PSD:

$$S_x(f) = \frac{2\sigma^2}{f_s}$$

The RPSD $\tilde{S}_x(f)$ is then related to the standard deviation σ by:

$$\tilde{S}_x(f) = \sigma\sqrt{\frac{2}{f_s}} \tag{3.70}$$

The standard deviation of the noise in Figure 3.17 was calculated to be $\sigma = 2.042$ MHz from 100 Hz to 20 MHz. Applying equation (3.70) implies a constant RPSD of 457 Hz/$\sqrt{\text{Hz}}$, $\sqrt{3}$ lower than the amplitude of the noise shelf in Figure 3.17. This $\sqrt{3}$ discrepancy is not unique to this particular noise generator, either; the same effect was observed using MATLAB to simulate the combined outputs of eight 15-bit LFSRs. We do not fully understand the origin of the $\sqrt{3}$ discrepancy, however our observations indicate that the peak RPSD of the digitally generated Gaussian noise is consistently $\sqrt{3}$ greater than the estimate provided by equation (3.70), which can be modified to account for the discrepancy:

$$\tilde{S}_x(f)^\dagger = \sigma\sqrt{\frac{6}{f_s}} \tag{3.71}$$

The RPSD of the unscaled frequency noise leaving the Gaussian noise generator is:

$$\begin{aligned}\tilde{f}_{\text{raw}}(f) &= 2.402 \text{ MHz } \times \sqrt{\frac{6}{40 \text{ MHz}}} \\ &= 790.9 \text{ Hz}/\sqrt{\text{Hz}}\end{aligned}$$

Noise shaping

Different types of noise can be simulated by manipulating the output of the white Gaussian noise generator. White noise is produced by scaling the generator's output by K_P, and $1/f$ noise is produced by integrating it.

Scaling the output of the Gaussian noise generator by $K_P = 2^{-8}$ produces white noise with an RPSD:

$$\begin{aligned}\tilde{f}_{\text{white}}(f) &= 2^{-8} \times \tilde{f}_{\text{raw}}(f) \\ &= 3.08 \text{ Hz}/\sqrt{\text{Hz}}\end{aligned}$$

Pink ($1/f$) noise can be produced by integrating white noise at a particular rate, f_s, and then scaling the integrated noise by K_I to adjust its amplitude. The RPSD of the $1/f$ noise can be estimated by multiplying $\tilde{f}_{\text{raw}}(f)$ by the transfer function of discrete time integrator:

$$\begin{aligned}\tilde{f}_{\text{pink}}(f) &= \text{abs}\left[\left(\frac{1}{1-z^{-1}}\right) \cdot K_I \cdot \tilde{f}_{\text{raw}}(f)\right] \\ &= \left(\frac{K_I f_s}{2\pi f}\right) \tilde{f}_{\text{raw}}(f)\end{aligned}$$

where

$$\frac{1}{1-z^{-1}} \approx \frac{f_s}{j2\pi f} \text{ for } f \ll f_s$$

The gain K_I needed to produce $1/f$ frequency noise with an amplitude $f_{\text{pink}} = 10 \text{ kHz}/\sqrt{\text{Hz}}$ at 1 Hz is:

$$\begin{aligned}K_I &= \frac{2\pi}{f_s}\left(\frac{10 \text{ kHz}/\sqrt{\text{Hz}}}{790.9 \text{ Hz}/\sqrt{\text{Hz}}}\right) \\ &= 1.99 \times 10^{-6} \approx 2^{-19}\end{aligned}$$

The closest power of two to the ideal gain is selected to simplify the $1/f$ noise generator's implementation on the FPGA.

A combination of white and pink noise with gains $K_P = 2^{-8}$ and 2^{-19} was used to simulate the laser frequency noise of a typical Nd:YAG non-planar ring oscillator (NPRO) laser, which is proportional to the inverse Fourier frequency and has an amplitude of $\delta f \approx 10$ $\text{kHz}/\sqrt{\text{Hz}}$ at 1 Hz. The noise was measured on the FPGA at the full sampling rate f_s. The measured and predicted RPSDs are shown in Figure 3.18 and are consistent.

3.8 Performance measurements

Two three noise tests were performed to simulate 1) the phase of a path-length matched Mach-Zehnder interferometer; and 2) the relative phase noise between two free running Nd:YAG NPRO lasers.

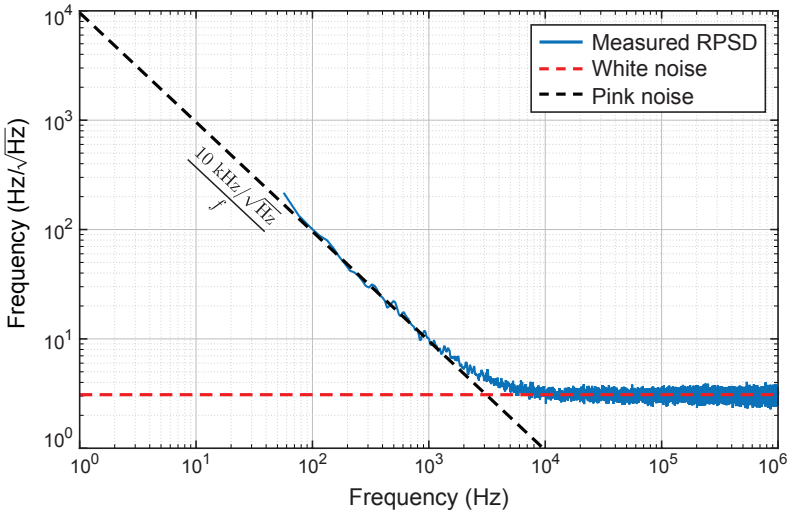

Figure 3.18: Measured RPSD of Gaussian-like noise that has been shaped to simulate $1/f$ laser frequency noise with a magnitude of 10 kHz/$\sqrt{\text{Hz}}$ at 1 Hz. The white noise was produced by scaling the raw output of the Gaussian noise generator by $K_P = 2^{-8}$, and the pink noise was created by integrating the raw output of the generator and then scaling the result by $K_I = 2^{-19}$.

3.8.1 Test one: white frequency noise

White frequency noise with an RPSD of $\delta\tilde{f} = 3.08$ Hz/$\sqrt{\text{Hz}}$ was added to seed frequencies 13, 7, and 4 MHz to produce noisy beat notes at 6, 9, and 3 MHz. The effective noise amplitude of each signal is $\delta\tilde{f}\sqrt{2}$ since f_A, f_B, and f_C are formed by linearly combining two uncorrelated sources of noise.

A time-series measurement of a 10 hour three noise test with white frequency noise is shown in Figure 3.19(a). The $1/f$ noise behaviour of the measured phase is caused by the integration of white frequency noise. The inset shows a five second segment of the null measurement $\phi_A - \phi_B + \phi_C$ with an approximately 3×10^8 zoom in amplitude.

Measurements were recorded at different data capture rates to produce the composite RPSD shown in Figure 3.19(b) from 1 mHz up to 10 kHz. The RPSD of the null measurement (black) is spectrally flat with a magnitude of 74.5 nano-cycles/$\sqrt{\text{Hz}}$. This indicates a PLL quantisation noise-floor of 30.4 nano-cycles/$\sqrt{\text{Hz}}$, consistent with the theoretical noise-floor of 30.2 cycles/$\sqrt{\text{Hz}}$ predicted in Section 3.5.3.

The phase noise measured by the PLL due to frequency noise can be calculated using equation (3.29):

$$\begin{aligned}\tilde{\phi}_{\text{noise}} &= \left(\frac{1}{1-z^{-1}}\right)\left(\frac{1}{f_s}\right)\tilde{f}_{\text{noise}} \\ &= \frac{f_s}{j2\pi f}\left(\frac{1}{f_s}\right)\tilde{f}_{\text{noise}} \\ &= \frac{1}{j2\pi f}\tilde{f}_{\text{noise}}\end{aligned} \tag{3.72}$$

White frequency noise with an amplitude 4.36 Hz/$\sqrt{\text{Hz}}$ has a corresponding phase noise of 0.69 cycles/$\sqrt{\text{Hz}}$ at 1 Hz, consistent with the results shown in Figure 3.19.

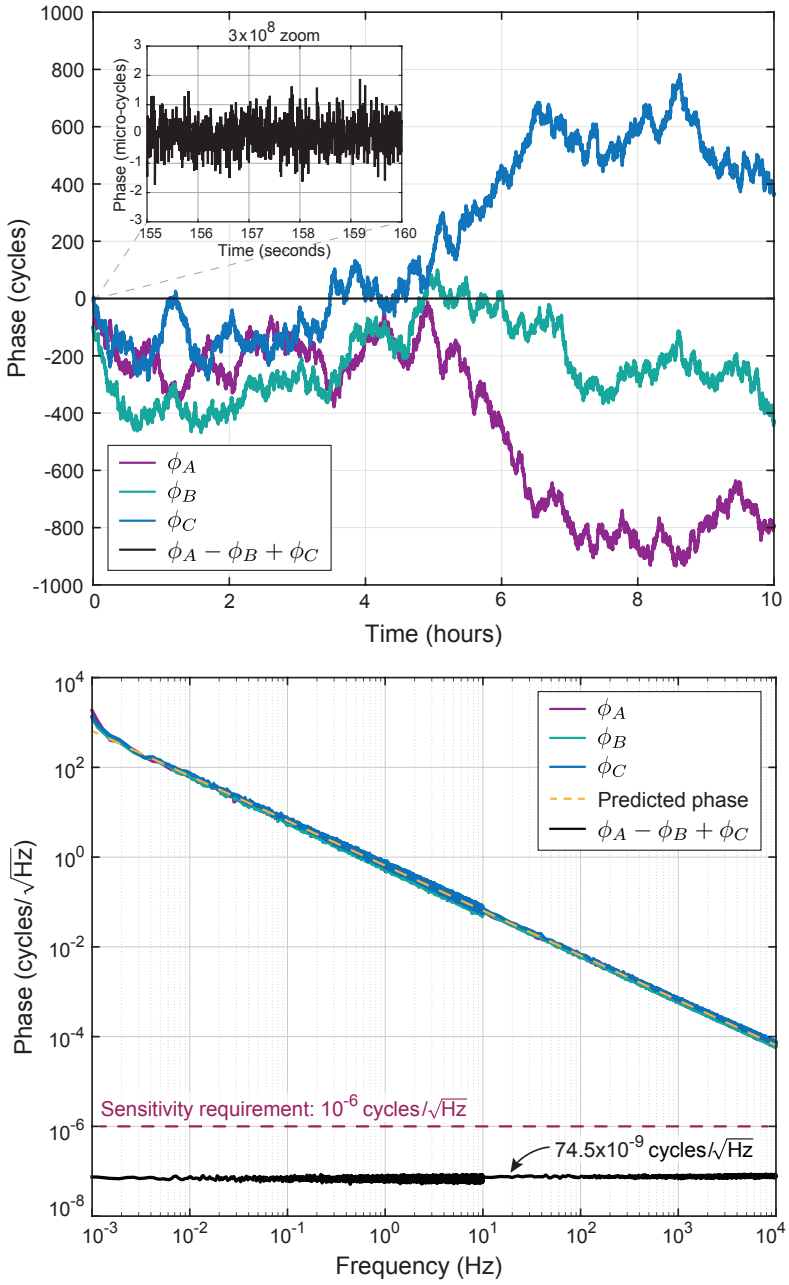

Figure 3.19: (a) Time series measurement of phase for a 10 hour three noise test with white frequency noise. The $1/f$ phase noise of ϕ_A, ϕ_B, and ϕ_C is caused by the integration of white frequency noise. The inset shows a 5 second segment of the null measurement ($\phi_A - \phi_B + \phi_C$) with $\sim 3 \times 10^8$ zoom in amplitude. (b) RPSD of the individual phase measurements and corresponding null measurement from the three-noise test data shown in (a). The RPSD of the null measurement (black) is spectrally flat with an RPSD of $\approx$ 74.5 nano-cycles/$\sqrt{\text{Hz}}$. Applying equation (3.69) resolves the quantisation noise-floor of each PLL to be 30.4 nano-cycles/$\sqrt{\text{Hz}}$.

3.8.2 Test two: pink frequency noise

Pink frequency noise with an RPSD $\delta\tilde{f} \approx 10$ kHz/$\sqrt{\text{Hz}}$ at 1 Hz was added to seed frequencies 18, 12, and 4 MHz to produce noisy beat notes at 6, 14, and 8 MHz.

The time-series measurement of the 8000 second long three noise test with pink frequency noise is shown in Figure 3.20(a). The corresponding RPSD is shown in Figure 3.20(b) revealing a null measurement RPSD at 76.5 nano-cycles/$\sqrt{\text{Hz}}$, corresponding to a PLL quantisation noise-floor of 31.2 nano-cycles/$\sqrt{\text{Hz}}$. The dynamic range of this measurement is $\sim 10^{15}$ at 200 mHz.

3.9 Summary

Phasemeters are the heart of the internally sensed OPA. They are used to precisely measure the phase of detected heterodyne beat-notes in order to control the relative output phase of the OPA, enabling it to perform coherent beam combination. In the context of the OPA's control system, the phasemeter is the sensor. The phasemeter's design and implementation is therefore crucial since its sensitivity, bandwidth, and through-put will all influence the OPA's performance.

Three types of phasemeter were introduced: 1) zero-crossing detectors, which time how long it takes for consecutive zero-crossings of a periodic signal; 2) in-phase and quadrature demodulation, which calculates the phase using trigonometry; and 3) phase-locked loops, which use continuous feedback control to *lock* the phase of a local oscillator to that of the input signal. Zero-crossing detectors are by far the simplest to implement, however they are severely disadvantaged by the fact that noise at integer harmonics of the signal's frequency alias into the measurement band. In-phase and quadrature demodulation is extremely effective, however it cannot be performed entirely on an FPGA—the calculation of phase must be performed on a CPU—which means it is not suitable for applications that require low latency feedback of phase information, for example the OPA's high bandwidth control system. Phase-locked loops on the other hand are: a) high throughput and well suited to applications requiring low-latency feedback of phase information; b) extremely precise, capable of sensing phase disturbances as small as 1 μcycle; and c) implementable on an FPGA. It is for these reasons that PLLs are used in the internally sensed OPA.

The PLL was designed to have a measurement sensitivity ϕ_{min} of less than 1 μcycle/$\sqrt{\text{Hz}}$; a bandwidth f_{BW} of up to 100 kHz; and an update rate of 1 MHz. The design was optimised using transfer function models in MATLAB and then tested in simulation and hardware. The PLL's measurement sensitivity was measured using a three noise test to resolve residual quantisation noise and any non-linearities present in the measurement band. The sensitivity was measured to be $\phi_{min} = 30$ nano-cycles/$\sqrt{\text{Hz}}$, over an order of magnitude more sensitive than the design requirement of 1 μcycle/$\sqrt{\text{Hz}}$. The PLL can operate at a bandwidth of 100 kHz, however exceeding this bandwidth will lead to unstable behaviour as the loop's phase margin drops below 60 degrees. The number of points in the PLL's LPF was set to $R = 40$, resulting in an update rate of 1 MHz at an FPGA sampling frequency of $f_s = 40$ MHz. The PLL's dynamic range was measured to be 10^{15} at 200 mHz.

The PLL presented in this Chapter was used during all three stages of the OPA's development.

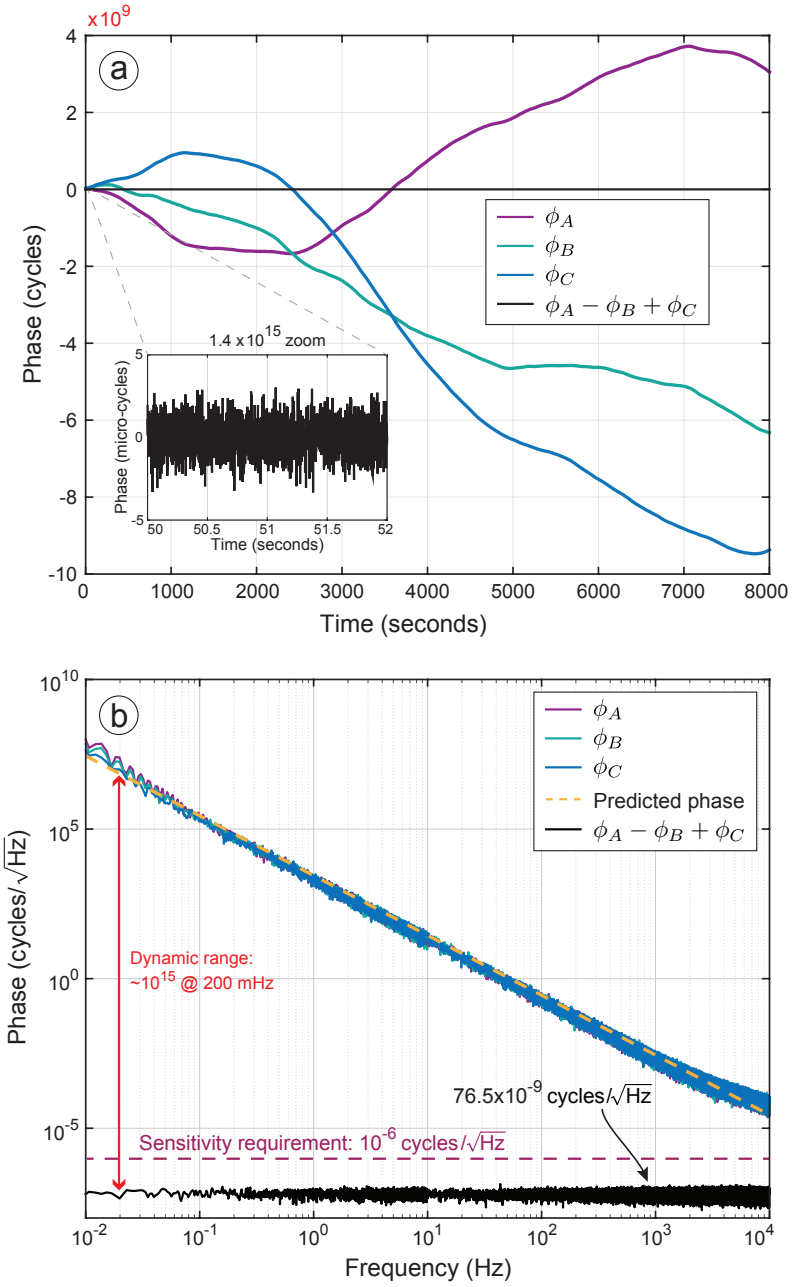

Figure 3.20: Time-series measurement of phase during an 8000 second three noise test simulating free running laser frequency noise of two Nd:YAG NPRO lasers. The $1/f^2$ phase noise of ϕ_A, ϕ_B, and ϕ_C is caused by the integration of pink frequency noise. The inset shows a 2 second segment of the null measurement ($\phi_A - \phi_B + \phi_C$) with $\sim 1.4 \times 10^{15}$ zoom in amplitude. (b) RPSD of the individual phase measurements, predicted phase measurement, and null measurement from the three-noise test with $1/f$ frequency noise. The RPSD of the null measurement (black) is spectrally flat with an RPSD of ≈ 76.5 nano-cycles/$\sqrt{\text{Hz}}$. Applying equation (3.69) resolves the quantisation noise-floor of each PLL to be 31.2 nano-cycles/$\sqrt{\text{Hz}}$. The measurement reveals a dynamic range of at least 10^{15} at 200 mHz.

Proof-of-concept OPA

4.1 Introduction

The first experimental demonstration of an internally sensed optical phased array was performed by Bowman et al. [58] using digitally enhanced heterodyne interferometry to measure and control the phase of three emitters in a 1D array. Whilst this demonstration concluded that internal sensing was possible, the OPA's functionality and performance were not fully characterised. This made it difficult to determine whether or not internal sensing would be capable of satisfying the requirements for coherent beam combining established in Section 1.2.2. To develop some confidence in the internal sensing technique, a proof-of-concept OPA was established to fully characterise its performance, explore its limitations, and address some of the weaknesses associated with the original design presented in [58].

This chapter begins with an analysis of the original internally sensed OPA in Section 4.2, establishing context for the design of the proof-of-concept OPA, and highlighting various ways in which to improve the original design. The design of the proof-of-concept OPA is described in Section 4.3, followed by details of its implementation both optically and in digital signal processing in Sections 4.4 and 4.5 respectively. Some of the challenges encountered during the proof-of-concept's implementation are discussed in Sections 4.6, 4.7 and 4.8. An analysis of the experimental results is presented Section 4.9.

4.1.1 Aims and requirements

The aim of the proof-of-concept OPA experiment described in this chapter was to demonstrate internal sensing's viability by coherently combining three emitters with an RMS output phase stability of $\Phi_{\mathrm{RMS}} \leq \lambda/24$, without the use of any external sampling optics to stabilise the relative output phase.

4.2 Original internally sensed OPA

A simplified diagram of the original internally sensed OPA presented in [58] is shown in Figure 4.1. Light from a master laser is phase modulated with pseudo-random noise and then split into three arms with different optical path lengths to produce three unique code delays at the return path photodetector. Fibre collimators are used to transmit the light from each aperture into free-space. An optically flat reference surface (e.g., a low reflectivity mirror) is used to reflect a small fraction of the light back into the collimators. The back-reflected light propagates in the reverse direction through the optical system

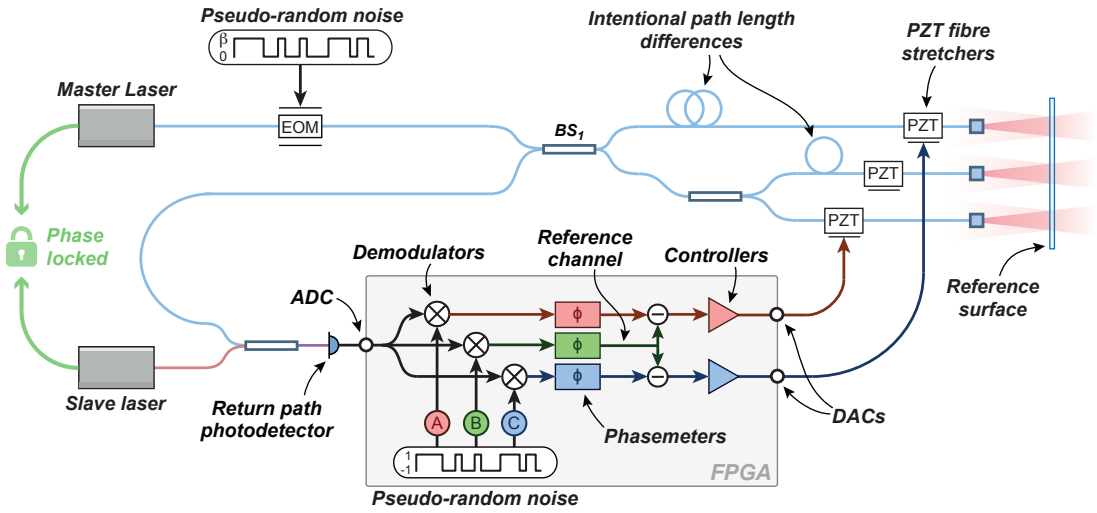

Figure 4.1: Simplified diagram of the original internally sensed OPA.

where it is interfered with a second offset phase-locked slave laser—which operates at a slightly different frequency to the master laser—to produce a heterodyne beat-note at the photodetector. The signal generated by the photodetector is digitally sampled using an analogue-to-digital converter interfaced with a field-programmable gate-array. The reflected electric fields originating in each of the three arms are isolated in the FPGA by demodulating the digitised signal with locally generated PRN codes with delays corresponding to their different optical path lengths: the longest optical path length in the system will experience a greater round-trip code delay before arriving at the photodetector.

The phase of the demodulated signals (ϕ_A, ϕ_B, ϕ_C) is measured using independent phasemeters implemented on an FPGA. The relative phase error between the channels is then calculated with respect to a common reference phase (in this example ϕ_B):

$$\Delta\phi_{B-A} = \phi_B - \phi_A$$

$$\Delta\phi_{B-C} = \phi_B - \phi_C$$

This phase error calculation eliminates any phase common to all three emitters (i.e., all optical path lengths in the system except the three arms between BS_1 and the reference surface). The calculated phase errors are passed through controllers to generate feedback signals needed to stabilise the relative optical path lengths of the OPA using piezoelectric transducer (PZT) fibre stretchers.

In practice, the scalability of the original internally sensed OPA is limited due to its need for meter-scale optical path length differences between arms to allow isolation of decoded signals. The minimum acceptable path-length difference between arms is

$$L_{\text{min}} = \frac{c}{n f_{\text{chip}}}$$

where c is the speed of light in vacuum, n is the effective refractive index of the optical fibre, and f_{chip} is the PRN chip frequency. The PRN chip frequency used in the original OPA was 40 MHz, corresponding to a minimum required path length difference of 5 meters. The complication here is that large differences in optical path length introduce a susceptibility to laser frequency noise, which can lead to an effect called 'beam-squint' that can cause the position of the interfered beam to change [84].

The proof-of-concept OPA presented in this chapter improves upon the original design by including electro-optic modulators in each arm, allowing each channel to be encoded

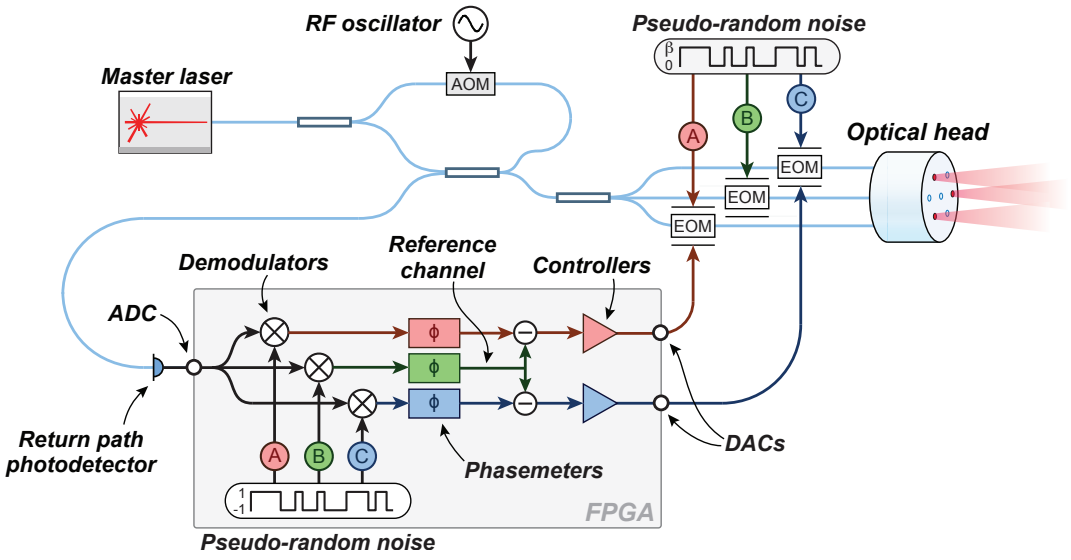

Figure 4.2: Idealised proof-of-concept internally sensed OPA.

individually using delayed versions of the same PRN code, relaxing the need for unique and increasingly large optical path length differences, and thus reducing its susceptibility to laser frequency noise and beam-squint. Including an EOM in each arm can also be used for phase-stabilising control and PRN modulation, foregoing the need to use PZT fibre stretchers which require high voltage amplifiers, and are typically bandwidth limited to ∼30 kHz due to mechanical resonances [85]. The EOMs used in the proof-of-concept OPA have a bandwidth of 150 MHz.

An alternative to using an offset phase-locked slave laser for heterodyne detection is to shift the frequency of the reference beam using an acousto-optic modulator.

4.3 Design

The proof-of-concept internally sensed OPA shown in Figure 4.2 comprises two sub-systems: 1) the optical system; and 2) digital signal processing implemented on an FPGA (including all analogue-to-digital interfaces between the optical and DSP sub-systems). Light from a free-running continuous wave laser is separated into two arms. The first (upper arm in the figure) is frequency shifted by f_h using a free-space AOM to generate the local oscillator for heterodyne detection. The second (lower arm in the figure) is split into three channels, each connected to a fibre waveguide EOM for PRN phase modulation and feedback control, the outputs of which are connected to an integrated optical head assembly. A small fraction of the light that leaves the OPA (∼4%) is reflected back into the fibre due to Fresnel reflection, propagating in the reverse direction through the fibre arrangement until it interferes with the frequency shifted local oscillator at a high-bandwidth return path photodetector. All optical fibre and fibre-coupled components are polarisation maintaining (PM). The optical system is described in more detail in Section 4.4.

The signal produced by the return path photodetector is digitised using a a high bandwidth ADC interfaced with a field-programmable gate-array. The digitised signal is digitally demodulated in parallel using delay matched versions of the PRN code used to encode the phase of each emitter. Multiple digitally implemented phasemeters are used to measure the phase of each of the decoded signals to generate feedback for relative path length stabilisation. The feedback signals exiting the controllers are converted into analogue

voltages using digital-to-analogue converters (DACs). The PRN codes are converted into analogue voltages using the FPGA's digital outputs. The PRN codes and feedback control signals are combined electronically using bias-tees. The digital signal processing system is described in more detail in Section 4.5.

4.3.1 Optical phase at the forward and return path photodetectors

The simplified three channel optical system shown in Figure 4.3 illustrates the various phase contributions from each of the optical path lengths travelled by the light on its way from the laser to the photodetector.

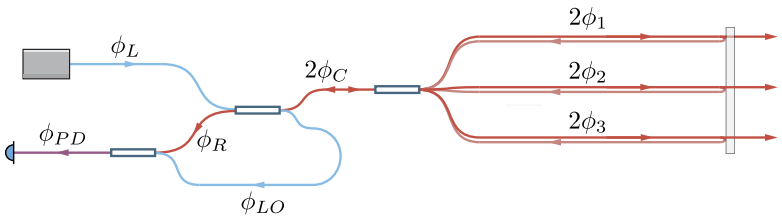

Figure 4.3: Individual phase contributions from every unique length of optical fibre in the system.

The measured phase of the three emitters at the photodetector is:

$$\begin{aligned}\Phi_1 &= (\phi_L + 2\phi_C + 2\phi_1 + \phi_R + \phi_{PD}) - (\phi_L + \phi_{LO} + \phi_{PD}) \\ &= 2\phi_C + 2\phi_1 + \phi_R - \phi_{LO} \qquad (4.1)\\ \Phi_2 &= (\phi_L + 2\phi_C + 2\phi_2 + \phi_R + \phi_{PD}) - (\phi_L + \phi_{LO} + \phi_{PD}) \\ &= 2\phi_C + 2\phi_2 + \phi_R - \phi_{LO} \qquad (4.2)\\ \Phi_3 &= (\phi_L + 2\phi_C + 2\phi_3 + \phi_R + \phi_{PD}) - (\phi_L + \phi_{LO} + \phi_{PD}) \\ &= 2\phi_C + 2\phi_3 + \phi_R - \phi_{LO} \qquad (4.3)\end{aligned}$$

The relative phase error of the three channels is calculated with respect to a common reference (in this example channel 1) to reject common phase signals (including noise), and generate error signals for feedback control:

$$\Phi_{error1,2} = \Phi_2 - \Phi_1 = 2\phi_2 - 2\phi_1 \qquad (4.4)$$

$$\Phi_{error1,3} = \Phi_3 - \Phi_1 = 2\phi_3 - 2\phi_1 \qquad (4.5)$$

Equations (4.4) and (4.5) show that the calculated phase error signals are sensitive only to the uncommon optical path lengths of the three emitters. Stabilising $\Phi_{error1,2}$ and $\Phi_{error1,3}$ will stabilise the relative OPA's output phase.

4.3.2 Phase ambiguity

A consequence of the internally sensed architecture is an ambiguity in phase due to the light's double pass through the same length of optical fibre. Consider the scenario where the output phase of two emitters are $\phi_1 = \pi$ and $\phi_2 = 0$, which have a *true* relative phase error of π at the OPA's output. Due to the round-trip propagation of the back-reflected signals, the measured phases at the return path photodetector will be $2\phi_1$ and $2\phi_2$, and

so the paths appear to be in-phase at the detector, when in fact they are destructively interfering at the optical head.

One way to solve this π-ambiguity in phase is to use a second, coarse phase measurement of each emitter through a different return path to calibrate the array upon start-up. This can be done by bonding a low reflectivity reference surface to the output of the optical head to reflect light into neighbouring 'dark fibres' in the array. This dark fibre calibration solves the π-ambiguity by correcting the original phase measurement by π if an incongruence exists between it and the calibration measurement. Since the measurement is sensitive to changes in phase (and not the absolute phase), the π-ambiguity calibration only needs to be performed once during initialisation.

The dark fibre calibration was not the focus of this research and is not considered in this thesis.

4.4 Optical implementation

The proof-of-concept OPA was implemented experimentally using the optical system shown in Figure 4.4.

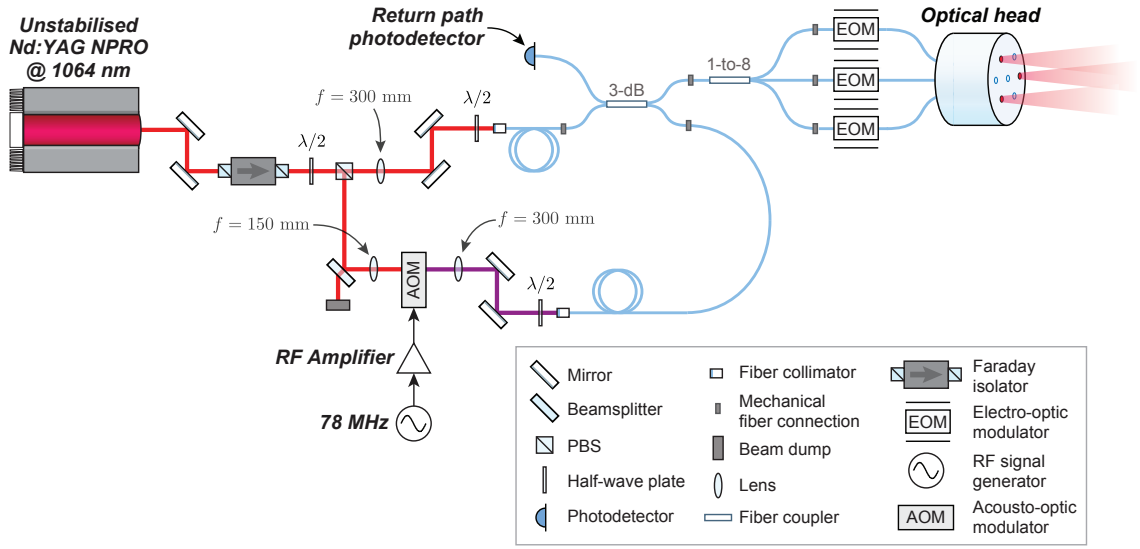

Figure 4.4: Detailed experimental setup of the proof-of-concept OPA.

Light from an Nd:YAG NPRO laser with planar polarisation was directed through a Faraday isolator using two steering mirrors to reduce back-reflections re-entering the laser cavity. A half-wave plate was used to adjust the relative amount of optical power leaving the two output ports of a polarising beam-splitter (PBS). The *signal arm* (upper arm in the diagram) was coupled into polarisation maintaining fibre using a 300 mm lens and two steering mirrors to mode-match the light entering a fibre collimator. Half-wave plates were used to align the polarisation of the light to the slow-axis of all fibres. The signal arm was optically connected to a PM 3-dB fibre coupler using a mechanical fibre-coupled / angled physical contact (FC/APC) connector. One of the outputs was connected to a 1-to-8 PM fibre waveguide splitter that uses fibre-coupled / flat physical contact (FC/FPC) connectors; three of the eight channels were connected to EOMs using FC/FPC to FC/APC conversion fibre patch cables. The remaining five outputs of the 1-to-8 splitter remained disconnected, which in hindsight was likely to have caused parasitic interference at the return path photodetector since the disconnected FC/FPC connectors were likely

to reflect ∼4% of the light back into the fibre due to Fresnel reflection. A simple way to suppress these back-reflections would have been to use commercially available FC/FPC beam-dumps, or to immerse the fibre ends in refractive index matched fluid (e.g., glycerin). The optical head was fusion spliced to the EOMs to minimise scattering and errors in the relative alignment of the fibre polarisation axes.

The reference arm (lower arm in the diagram) is directed into an AOM using a steerable 3-dB beam-splitter (to reduce power) and a focussing lens. The frequency shifted beam is re-collimated using another lens, and directed into PM fibre the same way as the signal arm. The reference arm is connected to the second output of the 3-dB fibre coupler where it interferes at the return path photodector with the back-reflected light from the optical head.

4.4.1 Laser

An Innolight Prometheus NE20 Nd:YAG NPRO CW laser [86] with a wavelength of 1064 nm was used for its narrow linewidth (<1 kHz) and correspondingly high coherence length. Narrow linewidth lasers reduce susceptibility to laser frequency noise in the presence of macroscopic optical path length differences.

The unstabilised Prometheus has a laser frequency noise that is proportional to the inverse of the Fourier frequency and has an amplitude of $\delta f \approx 10\text{kHz}/\sqrt{\text{Hz}}$ at 1 Hz [86]. While NPRO lasers with low frequency noise characteristics are desirable, they are not strictly necessary. The system presented here was not limited by laser frequency noise of the unstabilised NPRO.

Frequency noise, δf, delayed by a differential optical path length ΔL translates into phase noise, $\delta\phi$, according to

$$\delta\phi = \frac{2\pi n \Delta L}{c}\delta f \tag{4.6}$$

at frequencies much less than $c/(n\Delta L)$, where c is the speed of light in vacuum and n is the refractive index of optical fibre. The relationship between ΔL and δf means that (assuming a fixed maximum output phase noise requirement) a system using a laser with higher laser frequency noise will require smaller optical path length differences. Selecting a suitable laser for this architecture depends on its specific frequency noise characteristics, any optical path length differences between channels, and the required phase stability at the output of the OPA.

4.4.2 Acousto-optic modulator

An Isomet 1205C-843 AOM [87] (shown in Figure 4.5) was used to shift the frequency of the reference beam by 78 MHz. The 78 MHz RF tone was generated using an Agilent 33250A function generator [88]. An amplifier was used to amplify the RF tone up to the AOM's recommended drive power of 1.2 Watts. The relevant specifications for the AOM are provided in Table 4.1.

4.4.3 Electro-optic modulators

Photline NIR-MPX-LN-0.1 fibre coupled EOMs were used because of their DC to 150 MHz actuation bandwidths and good polarisation properties. The relevant technical specifications of the EOMs are presented in Table 4.2.

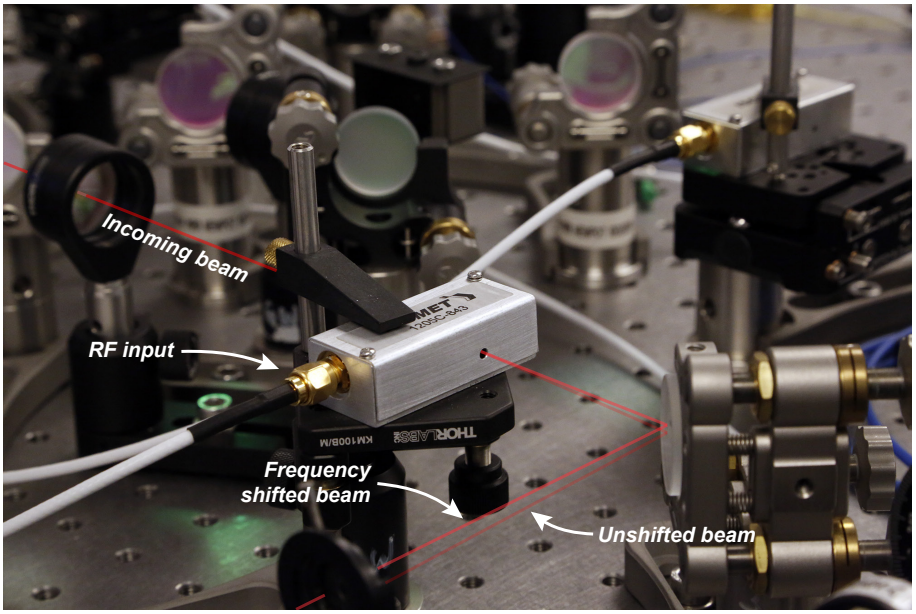

Figure 4.5: Photograph of an Isomet 1205C-843 acousto-optic modulator in the optical system.

Table 4.1: Technical specifications for the Isomet 1205C-843 AOM.

Parameter	**Specification**	**Unit**
Operating wavelength	1064	nm
Centre frequency, f_c	80	MHz
Tuned RF bandwidth, Δf	$\pm$ 15	MHz
Diffraction efficiency	$\geq$ 80.0	%
Max. RF input power	$<$ 1.3	W

Table 4.2: Technical specifications for the Photline NIR-MPX-LN-0.1 electro-optic modulator.

Parameter	**Specification**	**Unit**
Electro-optic bandwidth	150	MHz
V_π	2.5	V
RF input impedance	10,000	Ω
Max. optical input power	100	mW
Max. RF input voltage	$\pm$ 20	V

4.4.4 Photodetector

The interference between the back-reflection and local oscillator was measured using a 125 MHz bandwidth NewFocus 1811 photodetector [89].

4.4.5 Integrated optical head assembly

The custom designed integrated optical head assembly shown in Figure 4.6 was fabricated by Fibreguide Industries to test the OPA in a realistic configuration, and to avoid the need for any external sampling optics at the output of the array. The optical head assembly was designed to support up to 61-emitters to evaluate how tolerant the manufacturing processes were in regard to fibre positioning and alignment. It also provides an opportunity to scale the OPA to more emitters in the future.

The output fibres are polarisation maintaining and aligned to within 3° in polarisation across the array surface to maximise interference power in the far field. The relative pointing alignment has a maximum tolerance of $< 0.9°$. The emitters are distributed in a hexagonal topology with a pitch of $250 \pm 3\mu$m to resemble an approximately circular aperture. The surface of the array in the vicinity of the selected output fibres is flat to within 5 nm (as measured by the manufacturer) resulting in a constant output phase offset of approximately $\lambda/200$ between channels. Large inconsistencies in the array's surface flatness can result in permanent random phase offsets that may distort the interfered beam. In this experiment the inconsistencies are negligible, but if they weren't then they could be corrected by applying a constant corrective phase shift to each emitter.

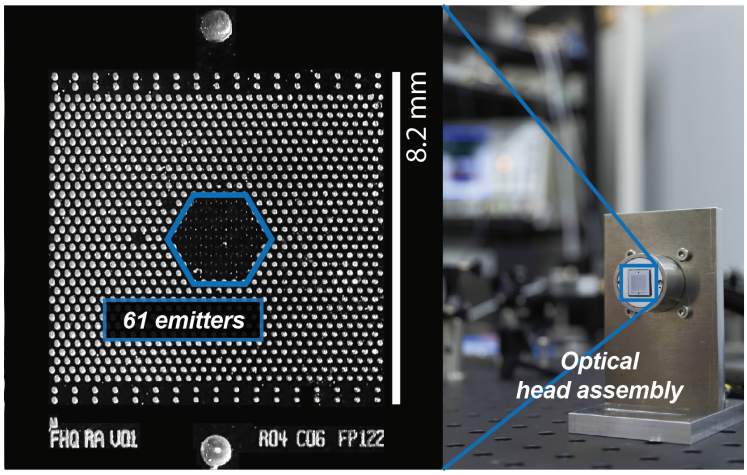

Figure 4.6: Photograph of the integrated optical head assembly.

4.5 Digital signal processing

The measurement and feedback control system implemented on the FPGA is shown in Figure 4.7. The voltage signal produced by the photodetector is digitised at 40 MHz using a 16-bit ADC. The digitised signal is demodulated in parallel with a delay matched version of the PRN code used to modulate the phase of the light in the optical system. The respective phase of the demodulated signals are measured using multiple phasemeters in parallel. Phase error signals are produce by subtracting the phase of the reference channel

(identified as channel R in Figure 4.7) from channels 2 and 3. Each phase error signal is passed through a controller to generate feedback for relative path-length stabilisation. The output of each controller is passed through a wrapping algorithm to ensure that the feedback signal does not exceed the maximum RF input voltage of the EOMs (see Table 4.2). The wrapped feedback signals are converted into analogue voltages using 1 MS/s DACs (DAC_1 and DAC_2). PRN codes are converted into analogue voltages using the FPGA's digital outputs (DO_1, DO_2, and DO_3). The analogue feedback signals for channels 2 and 3 and their respective PRN codes are combined electronically using bias-tees.

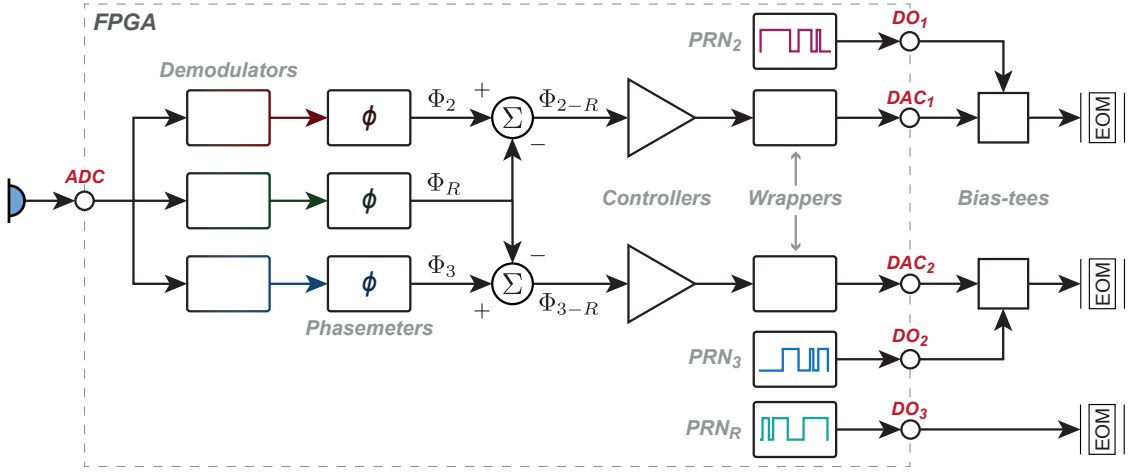

Figure 4.7: Digital signal processing architecture for the proof-of-concept OPA. The analogue feedback and PRN codes are combined electronically using bias-tees.

4.5.1 FPGA development platform

The FPGA development platform used in system was a National Instruments PXIe-1082e chassis [90] containing two PXI-7852R Virtex-5 LX50 FPGAs [91], and a PXIe-8133 real-time (RT) computer [92]. Peripheral ADCs (described in Section 4.5.2) were connected directly to the PXI-7852R. The FPGA's DACs and digital outputs were accessed through a custom-built breakout box with SMA connectors.

4.5.2 Analogue-to-digital converter

The ADCs used in this system are Maxim Integrated MAX1165EVKIT evaluation boards [93] that have been modified to interface with the National Instruments PXI-7852 R-series FPGA modules. The relevant technical specifications of the ADCs are presented in Table 4.3.

The ADCs have two inputs: a clock, and an analogue signal. The 40 MHz clock was generated using an Agilent 33250A function generator, which was phase locked to the FPGA's clock using a 10 MHz reference. The FPGA and ADC clocks must be synchronised to prevent glitching (random impulses in amplitude).

Table 4.3: Technical specifications for the Maxim Integrated MAX1168EV ADC.

Parameter	Specification	Unit
Resolution	16	bits
Input voltage range	±1.28	V
Sampling rate	$20 \leq f_{\text{clk}} \leq 100$	MS/s
Analogue bandwidth (-3 dB)	120	MHz

The ADC's conversion from Volts to counts is

$$\begin{aligned} A(z) &= \frac{2^{R-1}}{V_{\max}} \\ &= \frac{2^{15}}{1.28} \left[\frac{\text{counts}}{\text{Volt}}\right] \end{aligned} \tag{4.7}$$

where R represents the ADC's resolution (in bits) and $V_{\max}$ represents its maximum input voltage.

4.5.3 Digital-to-analogue converter

The analogue feedback control signals are generated using the NI PXI-7852R FPGA development platform's in-built DACs. The relevant technical specifications for the DACs are presented in Table 4.4.

Table 4.4: Technical specifications for the NI PXI-7852R DACs.

Parameter	Specification	Unit
Resolution	16	bits
Voltage range	±10	V
Impedance	0.5	Ω
Slew rate	10	V/μs
Update rate	1	MS/s
Channels	8	–

The DAC's conversion from counts to Volts is

$$\begin{aligned} D(z) &= \frac{V_{\max}}{2^{R-1}} \\ &= \frac{10}{2^{15}} \left[\frac{\text{Volts}}{\text{count}}\right] \end{aligned} \tag{4.8}$$

where R represents the DAC's resolution (in bits) and $V_{\max}$ represents its maximum output voltage

4.5.4 Digital outputs

The PRN code are generated using linear feedback shift registers (as described in Section 2.3), and converted into analogue voltages using the PXI-7852R's digital I/O ports.

The relevant technical specifications for the digital outputs are presented in Table 4.5.

Table 4.5: Technical specifications for the NI PXI-7852R digital outputs.

Parameter	**Specification**	**Unit**
Low voltage	$0 \leq V_{\text{low}} \leq 0.4$	V
High voltage	$2.4 \leq V_{\text{high}} \leq 3.3$	V
Output current	4	mA
Minimum pulse width	12.5	ns
Impedance[1]	$1 < Z < 100$	Ω
Output channels	8	–

The digital output ports of the FPGA are fed through a comparator; the supply voltage of the comparator can be adjusted to set the level of digital signal between 2.4 and 5 Volts.

4.5.5 Demodulators

Demodulators were implemented on the FPGA using digital multipliers as described in Section 2.5.

4.5.6 Phasemeters

The phasemeters used in this system are described in Chapter 3.

4.5.7 Controllers

Proportional-plus-integral controllers (shown in Figure 4.8) were used to generate the feedback signals needed to stabilise the relative path-lengths of the three channels. The proportional controller provides immediate feedback proportional to the amplitude of the error signal. The integral controller provides slow feedback proportional to the time integral of the error signal, providing additional gain at low frequencies. The proportional and global gains K_P and K_G are implemented using digital multipliers. The integral gain K_I is implemented using a right-pointing bit-shift, limiting division to inverse powers of two.

The inputs to each controller are the phase errors defined in equations (4.4) and (4.5).

4.5.8 Phase wrapping algorithm (wrappers)

The EOMs used to stabilise the optical path lengths have a limited actuation range of ~ 8 cycles, which can be extended by exploiting the 2π ambiguity of phase. A phase wrapping algorithm implemented at the output of each controller detects whenever the feedback

[1]The output impedance for the PXI-7852R is non-linear due to the circuitry used to provide 5V tolerance. The output impedance is low (<50 Ω) for output voltages near 0 V. As the output voltage increases to around 3.3 V, the impedance rises to approximately 100 Ω. As the voltage rises beyond 3.3 V, the impedance rises very quickly to MΩ [94].

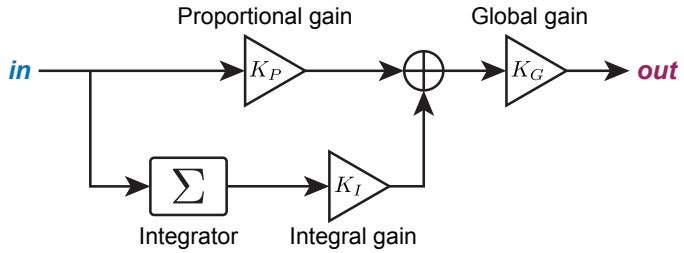

Figure 4.8: Proportional-plus-integral controller used in the proof-of-concept OPA.

signal exceeds two predefined thresholds and immediately wraps it by an integer number of cycles. The thresholds are padded to prevent rapid and repeated wrapping caused by noisy signals. In this system the phase is wrapped by one cycle, with padded thresholds located at ±0.6 cycles. The wrapping concept is illustrated in Figure 4.9.

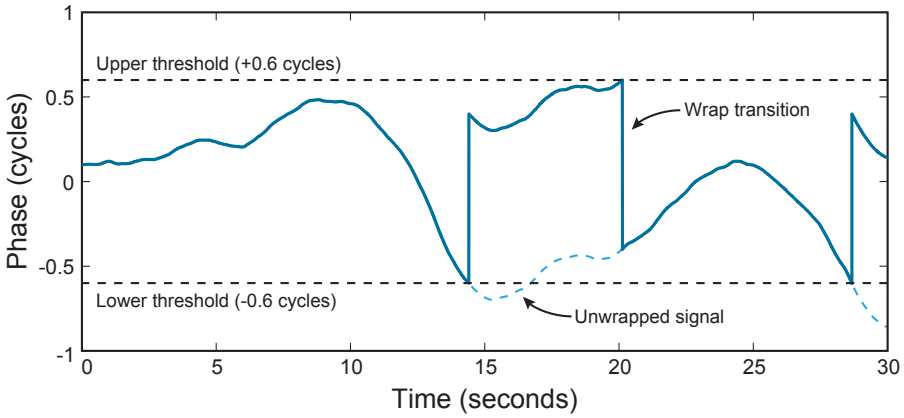

Figure 4.9: Phase wrapping algorithm used to exploit the 2π ambiguity of phase. At ~14 seconds the feedback signal encounters the lower threshold at -0.6 cycles. This prompts the wrapping algorithm to add 1 cycle to the feedback signal, which immediately jumps to +0.4 cycles.

The phase wrapping algorithm is implemented in DSP between the controller and DAC. When the feedback signal exceeds one of the two user-defined boundaries it will immediately wrap it by one cycle. But in order to implement this in hardware, it is necessary to understand the relationship between the applied EOM voltage and its resulting phase shift.

The EOM's response is typically linear (i.e., $\phi_{\text{EOM}} \propto V_{\text{applied}}$), and generally characterised by the specific voltage V_π required to shift the phase of the light by π radians. The phase shift ϕ_{EOM} imparted by an applied voltage V_ϕ is therefore

$$\phi_{\text{shift}} = \left(\frac{\pi}{V_\pi}\right) V_\phi$$

where the coefficient π/V_π relies exclusively on the EOM's V_π specification. The voltage V_ϕ required to shift the phase by ϕ_{shift} is thus:

$$V_\phi = \left(\frac{V_\pi}{\pi}\right) \phi_{\text{shift}} \tag{4.9}$$

Consider an R-bit DAC with a maximum output voltage range of $\pm V_{\max}$ Volts. The feedback signal in the FPGA S_{FB} is converted from bits to volts by:

$$V_{FB} = \left(\frac{V_{\max}}{2^{R-1}}\right) S_{FB}$$

The bit amplitude corresponding to a specific feedback voltage V_{FB} is therefore

$$S_{FB} = \text{floor}\left\{\left(\frac{2^{R-1}}{V_{\max}}\right) V_{FB}\right\} \tag{4.10}$$

where the 'floor' operation indicates the conversion from floating point to an integer. Equations (4.10) and (4.9) can be combined to define how large the feedback signal $S_{2\pi}$ must be in order to produce a phase shift of 2π at the EOM:

$$\begin{aligned} S_{2\pi} &= \text{floor}\left\{\left(\frac{2^{R-1}}{V_{\max}}\right) V_{2\pi}\right\} \\ &= \text{floor}\left\{\left(\frac{2^{R-1}}{V_{\max}}\right)\left(\frac{V_{\pi}}{\pi}\right) 2\pi\right\} \\ &= \text{floor}\left\{\left(\frac{V_{\pi}}{V_{\max}}\right) 2^{R}\right\} \end{aligned} \tag{4.11}$$

If $S_{2\pi}$ is defined incorrectly and instead produces a phase shift of $\phi = 2\pi \pm \phi_{error}$, then the phasemeters will measure a step change in phase equal in magnitude to ϕ_{error} that the control system will attempt to correct. The most likely reason for this to occur is an incorrect value for V_{π}, which may be slightly different for each EOM. The specific V_{π} for each EOM should therefore be calibrated to prevent the controller from injecting unwanted noise into the feedback loop.

If the slew rate of the wrapped feedback signal does not exceed the phasemeter's bandwidth then it will be measured, forcing the controller to impose it on the generated feedback signal. In this system the slew rate of the wrap should only be limited by the analogue bandwidth of the DAC, which is 500 kHz (half the sampling rate of the 1 MS/s DAC). This means that any wrap transitions should occur well outside the bandwidth of the phasemeters, which is typically between 10 and 50 kHz.

4.6 Decoding signals with low PRN modulation depths

The electric field at the detector consists of the PRN modulated back-reflected signals from the optical head, interfered with the local oscillator field. Ignoring losses, the measured interference signal $s(t)$ is:

$$s(t) = \sum_{1}^{n} A_n \sin(\omega_h t + \Phi_n + \beta c(t - \tau_n)) \tag{4.12}$$

where the index n represents the emitter number, A_n is the amplitude of the interfered field at the detector, ω_h is the angular heterodyne frequency, Φ_n the total phase of the n'th emitter, and β is the modulation depth of the binary PRN code $c(t - \tau) \in [0, 1]$ delayed by τ_n.

To recover the signal belonging to a specific emitter in the OPA the detected signal $s(t)$ is digitised using a 40 MHz analogue-to-digital converter and then decoded using a delay

matched version of the PRN code $p(t-\tau_n)$ belonging to that emitter. The PRN code used for demodulation is generated using the equation

$$p(t-\tau) = 1 - 2c(t-\tau) \tag{4.13}$$

to map $c(t) \in [0,1]$ to $p(t) \in [1,-1]$. Demodulation is performed in the FPGA by multiplying $s(t)$ by $p(t-\tau_k)$

$$d_k(t) = p(t-\tau_k) \sum_{1}^{n} A_n \sin(\omega_h t + \Phi_n + \beta c(t-\tau_n)) \tag{4.14}$$

where $d_k(t)$ represents the demodulated signal for the k'th emitter in the OPA. The PRN modulated phase component $\beta c(t-\tau_n)$ can be separated from $(\omega_h t + \Phi_n)$ using the trigonometric identity $\sin(a \pm b) = \sin(a)\cos(b) \pm \cos(a)\sin(b)$ where $a = \omega_h t + \Phi_n$ and $b = \beta c(t-\tau_n)$. Equation (4.14) then becomes

$$d_k(t) = p(t-\tau_k) \sum_{1}^{n} A_n \left[\sin(\omega_h t + \Phi_n)\cos(\beta c(t-\tau_n)) + \cos(\omega_h t + \Phi_n)\sin(\beta c(t-\tau_n))\right] \tag{4.15}$$

When the modulation depth is very small (i.e., $\beta \ll 1$) the small angle approximation for sine and cosine can be applied, simplifying equation (4.15) to

$$d_k(t) = p(t-\tau_k) \sum_{1}^{n} A_n \left[\sin(\omega_h t + \Phi_n) + \cos(\omega_h t + \Phi_n)\beta c(t-\tau_n)\right] \tag{4.16}$$

where $\cos(\beta c(t-\tau_k)) \approx 1$ and $\sin(\beta c(t-\tau_n)) \approx \beta c(t-\tau_n)$. Suppose we want to recover the signal for emitter 1 (i.e., $k=1$). Equation (4.16) becomes:

$$\begin{aligned} d_1(t) = {} & p(t-\tau_1) A_1 \left[\sin(\omega_h t + \Phi_1) + \cos(\omega_h t + \Phi_1)\beta c(t-\tau_1)\right] \\ & + p(t-\tau_1) \sum_{2}^{n} A_n \left[\sin(\omega_h t + \Phi_n) + \cos(\omega_h t + \Phi_n)\beta c(t-\tau_n)\right] \end{aligned} \tag{4.17}$$

The demodulated signal $d_1(t)$ has two components: i) the correctly demodulated signal $r_1(t)$, and ii) the sum of the remaining incorrectly demodulated signals, η_{noise}. The incorrectly demodulated signals are remodulated at full modulation depth $\beta = \pi$ since

$$p(t-\tau)\sin(\omega t) = \sin(\omega t + \pi c(t-\tau)) \tag{4.18}$$

and appear as broadband noise in the measurement of $d_1(t)$. Substituting $c(t-\tau) = \frac{1}{2}(1 - p(t-\tau))$ into the correctly demodulated component of $d_1(t)$ yields

$$r_1(t) = A_1 p(t-\tau_1) \left(\sin(\omega_h t + \Phi_1) + \frac{\beta}{2}\cos(\omega_h t + \Phi_1)\big(1 - p(t-\tau)\big)\right) \tag{4.19}$$

which, recalling that $p(t-\tau_k)\cdot p(t-\tau_k) = 1$, simplifies to:

$$r_1(t) = A_1 \left(p(t-\tau_1)\Big[\sin(\omega_h t + \Phi_1) + \frac{\beta}{2}\cos(\omega_h t + \Phi_1)\Big] - \frac{\beta}{2}\cos(\omega_h t + \Phi_1)\right) \tag{4.20}$$

The component of $r_1(t)$ that we are interested in measuring (highlighted green in equation (4.20)) is the only term *not* modulated by PRN. At this point the full demodulated

signal is:

$$\begin{aligned} d_1(t) &= r_1(t) + \eta_{noise} \\ &= A_1 \left[p(t-\tau_1) \sin(\omega_h t + \Phi_1) + \frac{\beta}{2} p(t-\tau_1) \cos(\omega_h t + \Phi_1) - \frac{\beta}{2} \cos(\omega_h t + \Phi_1) \right] \\ &\quad + p(t-\tau_1) \sum_2^n A_n \left[\sin(\omega_h t + \Phi_n) + \frac{\beta}{2} \cos(\omega_h t + \Phi_n)(1 - p(t-\tau_n)) \right] \end{aligned} \tag{4.21}$$

The demodulated signal's signal-to-noise ratio (SNR) can be improved by averaging over the full length of the PRN code $p(t-\tau_1)$ to exploit the auto-correlation properties of m-sequences as demonstrated in Section 2.5. Integrating $d_1(t)$ over L points improves the amplitude of the correctly demodulated signal by a factor of L relative to all other signals remodulated by $p(t-\tau_1)$.

Combining the remaining noise signals into a blanket noise term η_b allows us to condense equation (4.21) to

$$\sum^L d_1(t) \approx -\left(L A_1 \frac{\beta}{2} \cos(\omega_h t + \Phi_1) + \eta_b \right) \tag{4.22}$$

where

$$\begin{aligned} \eta_b &= A_1 \left(p(t-\tau_1) \sin(\omega_h t + \Phi_1) + \frac{\beta}{2} p(t-\tau_1) \cos(\omega_h t + \Phi_1) \right) \\ &\quad + p(t-\tau_1) \sum_2^n A_n \left(\sin(\omega_h t + \Phi_n) + \frac{\beta}{2} \cos(\omega_h t + \Phi_n)(1 - p(t-\tau_n)) \right) \end{aligned}$$

This analysis demonstrates that whilst it is possible to successfully isolate signals at low modulation depths, doing so will lead to a substantially reduced SNR of the demodulated signal. The consequence of this reduction in SNR is the necessary reduction in phasemeter bandwidth to avoid non-linear behaviour (e.g., cycle-slipping[2] [95, 96]), and to limit the introduction of residual PRN noise into the feedback loop. Furthermore, reducing the phasemeter's bandwidth also reduces the bandwidth of the controller, limiting it's ability to stabilise the relative output phase of the OPA. The balance between modulation depth, phasemeter bandwidth, controller bandwidth, and coherence should therefore be selected carefully.

4.7 Double-passing the EOM

One of the complications with the proof-of-concept OPA is the fact that the back-reflected light passes through the EOM twice. This means the back-reflected signal is modulated twice with slightly different delayed versions of the PRN code as illustrated in Figure 4.10.

Consider the electric field of the light entering the EOM:

$$E_0(t) = E_0 e^{i(\omega t + \phi)}$$

[2]Cycle slips are non-deterministic jumps in phase caused by the PLL's non-linear response to random noise. The likelihood of a cycle slip event occurring is low for high SNR input signals, however the probability increases dramatically with increasing noise power. This makes low SNR input signals more susceptible to cycle slipping, introducing step errors into the phase measurement.

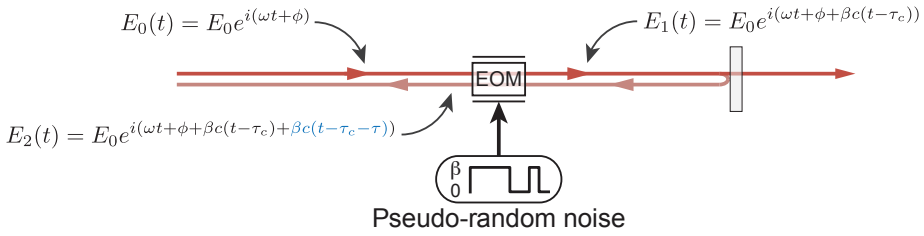

Figure 4.10: Double modulation of the back-reflected signal.

The light exiting the EOM on its first pass is modulated by the PRN code $c(t-\tau_c) \in [0,1]$ at the specific code delay τ_c with modulation depth β:

$$E_1(t) = E_0 e^{i(\omega t+\phi+\beta p(t-\tau_c))}$$

The modulated forward propagating beam travels towards the output of the fibre wave guide where a small fraction of the optical power is reflected back towards the EOM. Some time later the backward propagating reflected signal will re-enter the EOM and be modulated a second time by the same code delayed by the total round-trip delay τ:

$$E_2(t) = E_0 e^{i(\omega t+\phi+\beta c(t-\tau_c)+\beta c(t-\tau_c-\tau))} \tag{4.23}$$

The round-trip delay τ is related to the round-trip optical path length L_{opt} by

$$\tau = \frac{nL_{\text{opt}}}{c} \tag{4.24}$$

where n is the refractive index of the fibre and c is the speed of light in vacuum. Whether or not the double-modulation is a problem depends on i) how big τ is compared to the chip period of a single chip of the modulating PRN code T_{chip}, and ii) the sampling rate of the ADC f_{adc} which (in this system) is equal to the inverse period of a single PRN chip:

$$f_{\text{adc}} = \frac{1}{T_{\text{chip}}}$$

If we assume the code delay to be an integer k number of PRN chips

$$\tau_c = kT_{\text{chip}}$$

and the round-trip delay to be a fraction of the PRN chip

$$\tau = \delta T_{\text{chip}}$$

then we can rewrite equation (4.23) as:

$$E_2(t) = E_0 e^{i(\omega t+\phi+\beta c(t-kT_{\text{chip}})+\beta c(t-(k+\delta)T_{\text{chip}}))}$$

The ideal scenario is when the fractional delay $\delta \ll 1$ (and therefore $\tau \ll T_{\text{chip}}$) because the back-reflected signal will be modulated by the same delayed version of the PRN code, effectively doubling its modulation depth. If we assume for convenience that $T_{\text{chip}} = 1$, and that $\delta \ll 1$, then:

$$\begin{aligned} E_2(t) &= E_0 e^{i(\omega t+\phi+\beta c(t-k)+\beta c(t-k-\delta))} \\ &\approx E_0 e^{i(\omega t+\phi+2\beta c(t-\tau_c))} \end{aligned}$$

This highlights an important design consideration, which is that if the round-trip delay is sufficiently smaller than the period of a single PRN chip, then the modulation depth of the signal at the return path photodetector is effectively doubled without further degrading the coherence of the combined field. This is advantageous because it was established in Section 4.6 that increasing β improves the SNR of the demodulated signal, reducing uncertainty in the phase measurement.

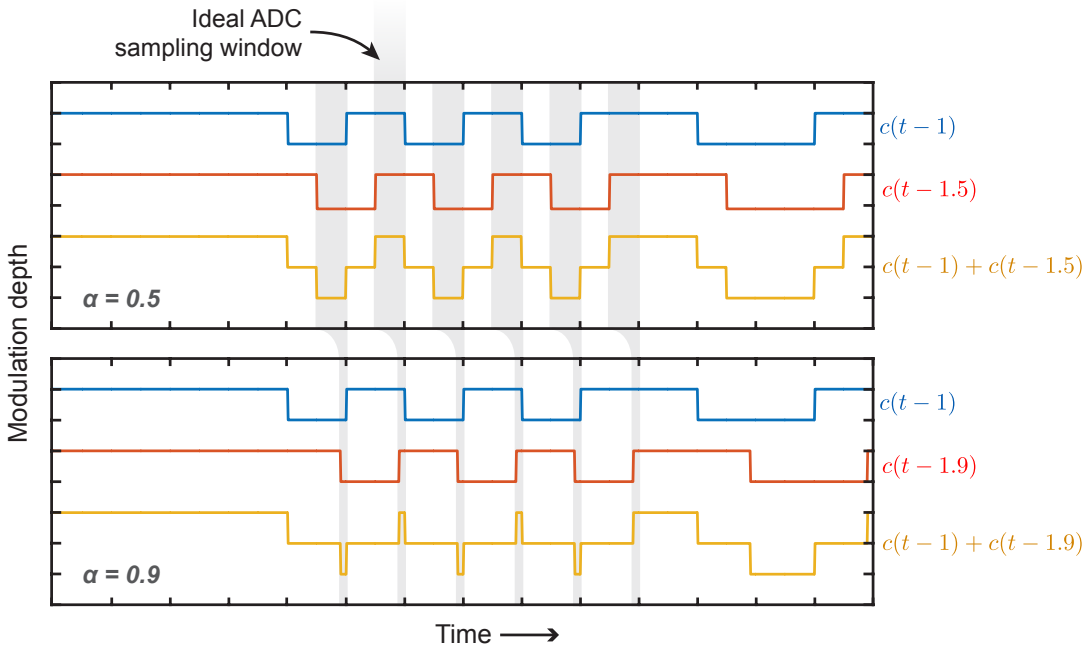

Figure 4.11: Effect of double modulation on the modulation depth of the back-reflected light when $\delta = 0.5$ (top) and $\delta = 0.9$ (bottom). Each marker indicates the period of single chip in the PRN code. The ideal ADC sampling window indicates the region where the two modulating codes overlap and double the modulation depth of the back-reflected light.

A more realistic scenario is that the fractional round-trip delay δ is not small due to the potentially large round-trip optical path between the EOM and optical head. This means that the back-reflected light will end up being re-modulated by at a different delay relative to the forward propagating light. The consequence of this is that only a fraction $(1 - \delta)$ of the back-reflected light will experience a doubling in modulation depth. Figure 4.11 illustrates this effect when $\delta = 0.5$ and 0.9. When $\delta > 1$, the two codes are offset by more than one chip and there is no possibility of doubling the modulation depth.

In order to take advantage of increased modulation depth, we need to ensure that the ADC samples the back-reflected light in the correct window. Figure 4.11 highlights the ideal ADC sampling window based on the fractional round-trip delay δ, which gets narrower as δ approaches 1. If the signal is sampled within this window then the PRN code will appear to be $2\beta c(t - \tau_c)$. If the signal is sampled outside this window (but still within the same chip) then it will appear to be encoded at two consecutive delays $c(t - \tau_c)$ and $c(t - \tau_c - 1)$. This means that both delays may be used to demodulate the signal, but because they are only separated by one chip there is a risk that the demodulated signals will experience mutual cross-talk, degrading the SNR even further.

This analysis assumes that the analogue bandwidth of the ADC is much greater than its sampling frequency. If the analogue bandwidth is too low then high frequency features (such as the narrow pulses in the yellow curve in the bottom panel of Figure 4.11) will be

low-pass filtered, reducing any benefit of sampling in the ideal window.

The ADC can be configured to sample the detected signal within the correct window (as shown in Figure 4.12) by adjusting the phase of the clock that drives it while observing the amplitude of the demodulated signal using the phasemeter. Scanning the phase of the ADC clock by a full cycle will shift the sampling point through an entire PRN chip assuming $f_{\text{adc}} = 1/T_{\text{chip}}$. If the phase of the ADC's clock cannot be adjusted directly then the same effect can be achieved by adding intentional electronic path length delays in the system either between the photodetector and ADC, or between the digital PRN code outputs and the EOMs, allowing the phase of each channel to be adjusted separately.

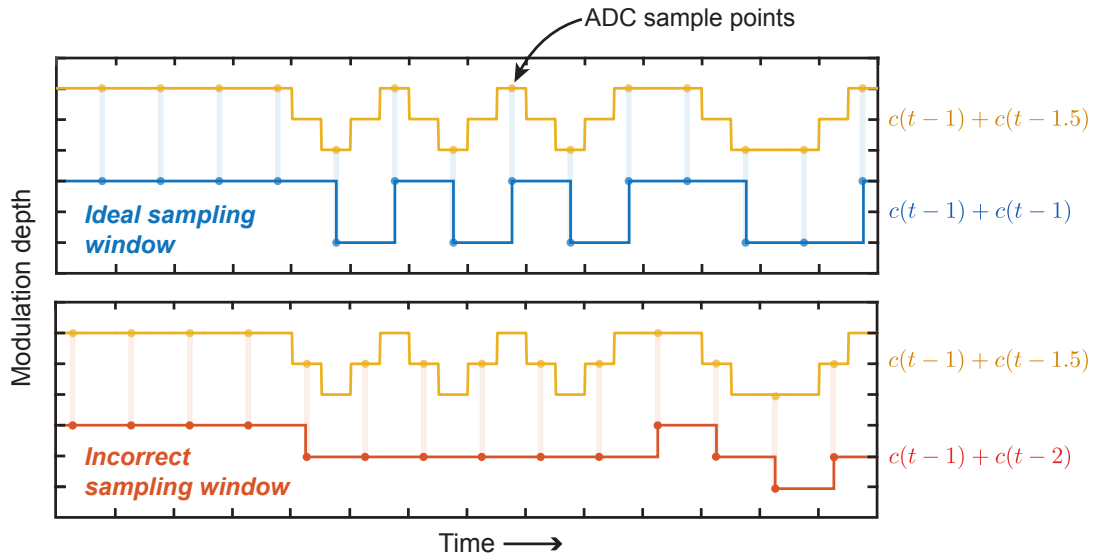

Figure 4.12: Illustration showing the effect of sampling the double-modulated back-reflection within the ideal and incorrect sampling windows.

4.7.1 Experimental round-trip optical path length

The round-trip optical path length between the optical head and EOMs was measured to be approximately 2.5 m. Assuming a typical refractive index of glass $n = 1.5$, speed of light in vacuum $c = 3 \times 10^8$ ms^{-1}, and PRN chip period $T_{\text{chip}} = 25$ ns, the fractional round-trip time delay δ was calculated to be:

$$\begin{aligned}\delta &= \frac{\tau}{T_{\text{chip}}} \\ &= \frac{1.5 \times 2.5}{25 \times 10^{-9} \times 3 \times 10^8} \cdot \left[\frac{\text{m}}{\text{s} \cdot \text{ms}^{-1}}\right] \\ &= 0.5\,[\text{chips}]\end{aligned} \tag{4.25}$$

The phase of the ADC clock was adjusted by adding additional signal cables between the photodetector and ADC to maximise the SNR of the demodulated signal. An alternative solution would have been to adjust the phase of the 40 MHz oscillator used to drive the ADC clock input.

4.8 Feedback actuation and PRN modulation using the same EOM

The low-frequency (DC to 100 kHz) feedback and RF (40 MHz) PRN sequence are combined electronically using a bias-tee. This allows PRN phase modulation and feedback actuation to be performed using the same EOM, reducing the number of components required in the optical system.

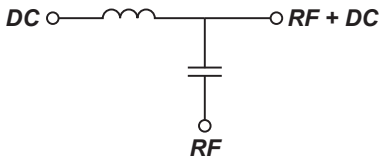

Figure 4.13: Circuit diagram of a bias-tee. The capacitor in the RF port blocks the DC signal, and the inductor in the DC port blocks the RF signal.

The circuit diagram for a bias-tee is shown in Figure 4.13. The bias-tee's RF port has a capacitor to block DC, and its DC port contains an inductor to block RF. The bias-tee used in this system has a rated DC cut-off frequency of 100 kHz and an RF bandwidth of 6 GHz. This means that any RF signals lower than 100 kHz should not enter the DC port, and any DC greater than 100 kHz cannot enter the RF port.

When the bias-tee was first tested it was noticed that the low-frequency feedback signal was being filtered far more aggressively than anticipated; the PRN code was not filtered. This limited the slew rate of feedback wrap transitions to approximately 20 kHz, far lower than the 100 kHz that was originally expected. This reduction in slew rate meant that wraps in the feedback signal were occurring slow enough for the phasemeter to measure. This resulted in the controller attempting to correct correct each wrap, quickly resulting in instability.

This reduction in slew-rate due to the bias-tee's low-pass filtering effect was never resolved since the bandwidth of the phasemeters ended up being reduced to approximately 1 kHz anyway to filter residual phase noise introduced by imperfect demodulation (as discussed in Section 4.6). This reduction in phasemeter bandwidth prevented them from being able to measure the slew-rate limited feedback wraps. If this were not the case, the most likely solution would have been to use a bias-tee with a higher cut-off frequency.

4.8.1 Total combined signal amplitude

The total combined amplitude of the PRN and feedback signals cannot exceed the maximum recommended operating voltage of the EOMs (± 20 V). The maximum output voltage range of the digital output and comparator circuit is ± 2.5 V. The maximum output voltage of the DACs is ± 10 V. The total combined amplitude of the combined signal could therefore potentially reach ± 12.5 V, well within the safe operating range of the EOMs.

4.9 Results and analysis

4.9.1 Intensity scaling and fringe visibility

The output intensity of the OPA was measured using a beam profiler for one, two, and three emitter configurations. Figure 4.14 presents the measured intensities normalised to the intensity of the central emitter (solid lines), accompanied by simulated expectations (dashed lines). The simulation assumes measured values of the fibre beam divergence angle (6.6°), projection range (60 ± 1 mm), and the relative power of each emitter (which were different by up to 5%). The simulation also assumed the fibre separation given by the manufacturer (250 μm), that the three emitters are co-planar, and that the modulation depth $\beta = 0$. The measurement agrees well with simulation.

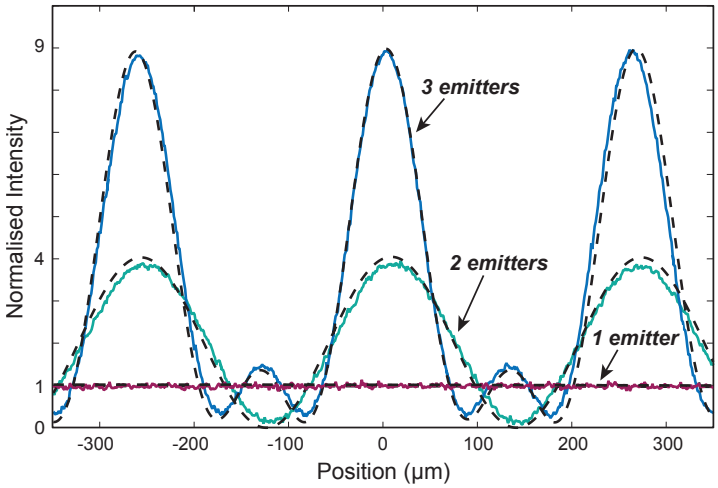

Figure 4.14: Far field interference profiles for one, two, and three emitter configurations. (Solid) measured; (dashed) simulation.

Table 4.6: Measured peak intensity of the central interference fringe for one, two, and three emitters. The measured results are normalised to the intensity of the central emitter.

No. of emitters	I_{norm} (measured)	I_{norm} (quadratic)
1	1.00	1
2	3.87	4
3	9.03	9

The peak intensities of the central interference lobes in Figure 4.14 are presented in Table 4.6, revealing quadratic scaling of intensity. The single emitter measurement appears flat because the beam profiler was looking at a small segment of the Gaussian beam. Fringe visibility was calculated to be 95.2% and 95.9% for the two and three emitter configurations respectively. The results agree well with simulation, and meet the requirements established in Section 4.1.1.

4.9.2 RMS output phase stability

A common method for measuring Φ_{RMS} involves interfering the output of the OPA with a passively stable frequency shifted reference beam to generate a heterodyne beat-note at a free-space photodetector [57]. This method assumes that fluctuations in the phase of the combined beam are far greater in magnitude than those experienced by the passively stable reference, and that both arms in the interferometer have equal optical path lengths to limit susceptibility to laser frequency noise. While this assumption *may* be valid over sufficiently short time-frames, it is almost certainly invalid at longer time scales where fibre phase noise caused by environmental disturbances dominates. The authors of [57] acknowledge this and consequently limit their measurement period to 1 ms, preventing them from evaluating their system's RMS output phase stability at sub-kHz frequencies. We decided against this particular method because we were not confident that the resulting measurement would accurately reflect the true RMS output phase error of the OPA, and were interested in measuring Φ_{RMS} over longer time scales (> 1 ms).

In this experiment, Φ_{RMS} was inferred by measuring the RMS angular disturbance of the output beam θ_{RMS} caused by phase noise between emitters at the OPA's output. The relationship between θ_{RMS} and Φ_{RMS} is [97]

$$\Phi_{\text{RMS}} = \frac{2\pi D}{\lambda} \cos(\theta)\ \theta_{\text{RMS}} \tag{4.26}$$

where D represents the emitter separation (250 μm), λ is the wavelength (1064 nm), θ is the mean steering error, $\theta_{\text{RMS}} = \tan^{-1}(x_{\text{RMS}}/z)$ where x_{RMS} is the RMS displacement error measurement at the detector, and z is the distance from the optical head to the beam profiler.

To measure θ_{RMS} a single fringe of the far field interference pattern was apertured using two razor blades mounted on adjustable platform stages located at the aperture of the beam profiler, positioned 65 ± 1 mm from the optical head. By measuring the location of the peak for 10 minutes at a resolution of 1.2 μm and sampling rate of 5 Hz, θ_{RMS} was calculated to be $\pm$ 35 μrad. Applying equation (4.26) with the previously stated parameters results in a Φ_{RMS} of 0.052 radians, or $\lambda/120$. Because this measurement was recorded at a sampling rate of 5 Hz, it was unable to capture higher frequency disturbances that were most likely to limit Φ_{RMS}.

It was observed that the output phase stability deteriorated significantly as the phasemeter's bandwidth was increased. This is likely because the phasemeter was able to track higher bandwidth noise introduced by the imperfect demodulation of the PRN encoded signal.

It was also found that the proof-of-concept OPA was very sensitive to shock; accidentally knocking the optical bench would often cause the control system to fail. The effect was less pronounced at higher PRN modulation depths as the corresponding increase in SNR of the demodulated signal permitted higher control bandwidths.

4.9.3 PRN modulation depth

The analysis in Section 4.6 states that the signal-to-noise ratio (SNR) of the demodulated signal (from which the phase is measured) scales with modulation depth at the expense of power delivered to the central interference lobe. The higher the SNR, the higher the potential phasemeter bandwidth, and therefore greater the potential forward-loop controller bandwidth. The measurements presented in Section 4.9.2 were performed at a

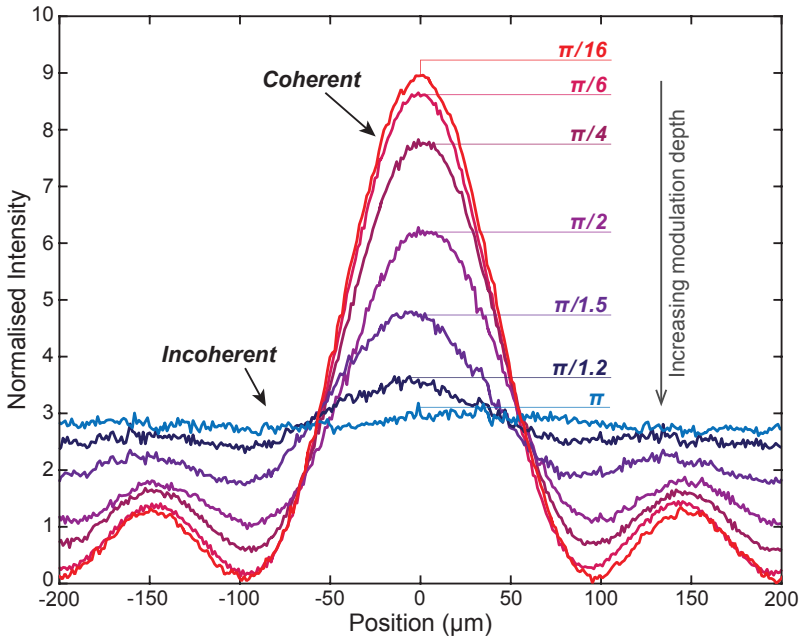

Figure 4.15: Measured degradation of fringe visibility and intensity scaling of a three element OPA as the PRN modulation depth increases from $\pi/16$ to π.

single-pass modulation depth of $\pi/16$ (double-pass modulation depth of $\pi/8$) as it had no noticeable impact on fringe visibility, and supported an RMS output phase stability of $\Phi_{RMS} = \lambda/120$. The controller did however have difficulty maintaining lock at lower modulation depths due to cycle-slipping triggered by excessive noise. If it weren't for the internal sensing architecture's inherent π phase ambiguity (Section 4.3.2) it is unlikely that cycle-slips would be catastrophic to the OPA's performance since the 2π phase slips would still result in a relative output phase error of zero. Instead, cycle-slips can cause emitters to re-stabilise out of phase by π, consequently distort the interfered beam.

The relationship between modulation depth and coherence was investigated qualitatively using the beam profiler to record the 1D far field intensity of the interfered light at different modulation depths *without* the benefit of double modulation. The reason for foregoing double modulation was to allow the PRN modulation depth to be increased to π without compromising DEHI's ability to isolate channels at the return path detector. This experiment was performed with round-trip optical path lengths between the EOM and optical head of approximately 15 m, three times the PRN chip length of 5 m.

The PRN modulation depth was changed by adjusting the supply voltage of the comparators on each of the FPGA's digital outputs, and an oscilloscope was used to measure the voltage supplied to each of the EOMs to infer the modulation depth in radians using the EOM's V_π rating.

The measured far field intensity profiles at modulations depths $\beta = \pi/16$, $\pi/6$, $\pi/4$, $\pi/2$, $\pi/1.5$, $\pi/1.2$, and π are presented in Figure 4.15. The results clearly show the degradation of coherence with increasing PRN modulation depth. The fringe visibility of the combined beam at full modulation depth $\beta = \pi$ is effectively zero due to the total randomisation of the carrier's phase which appears as broadband noise in the frequency domain. The peak intensity of the combined field at full modulation depth is approximately three times the intensity of a single emitter, consistent with incoherent combination.

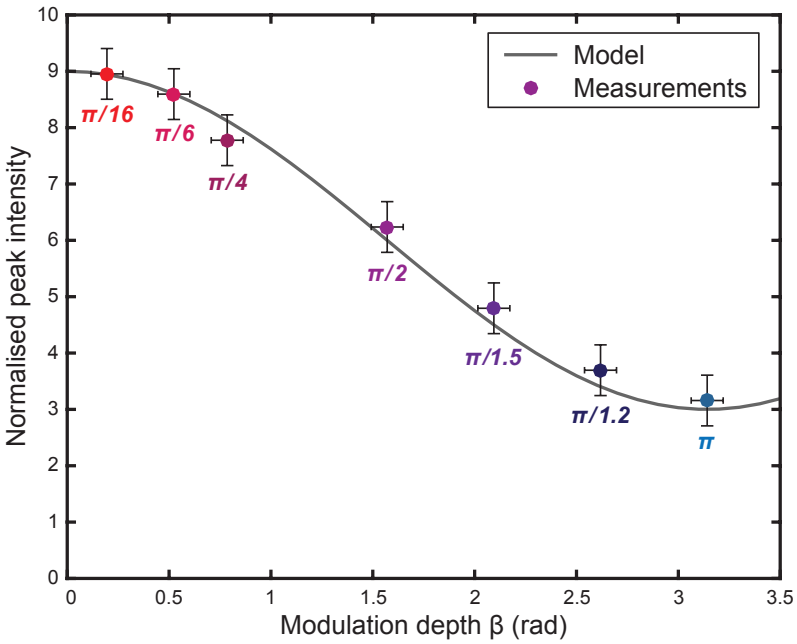

Figure 4.16: Measured degradation of fringe visibility and intensity scaling as the PRN modulation depth increases from $\pi/6$ to π.

The peak intensities of the interference profiles in Figure 4.15 are plotted against their corresponding modulation depths in Figure 4.16. The uncertainty in amplitude is attributed primarily to fluctuations in interference power caused by polarization drifts of each emitter. The error in modulation depth is attributed to the uncertainty of the true V_π ratings of the three EOMs.

The amplitude of N coherently combined signals with equal amplitude scales with N; whereas the amplitude of N incoherently combined signals (e.g., uncorrelated noise) scales with $\sqrt{N}$. The peak intensity therefore scales with N^2 for coherent combination, and N for incoherent combination. However, as shown in Figure 4.16, the peak intensity gradually decreases from N^2 at $\beta = 0$ to N at $\beta = \pi$. This transition from coherent to incoherent combination depends on the effective amplitude of the signal at the carrier frequency after being phase modulated with PRN, which is approximately proportional to $\cos\left(\frac{\beta}{2}\right)$ as shown in the phasor diagram in Figure 4.17. The power of the signal at the original carrier frequency is then proportional to:

$$\cos^2\left(\frac{\beta}{2}\right) = \frac{1+\cos(\beta)}{2} \tag{4.27}$$

The peak intensity as a function of modulation depth can then be modelled by:

$$I_p(\beta) = \left(\frac{N^2-N}{2}\right)\cos(\beta) + \frac{N^2+N}{2} \tag{4.28}$$

Equation (4.28) is shown in Figure 4.16 with the measured results, clearly showing the degradation of coherence with increasing modulation depth. The coherence will improve as the modulation depth increases towards 2π, as this would be indistinguishable from a modulation depth of zero.

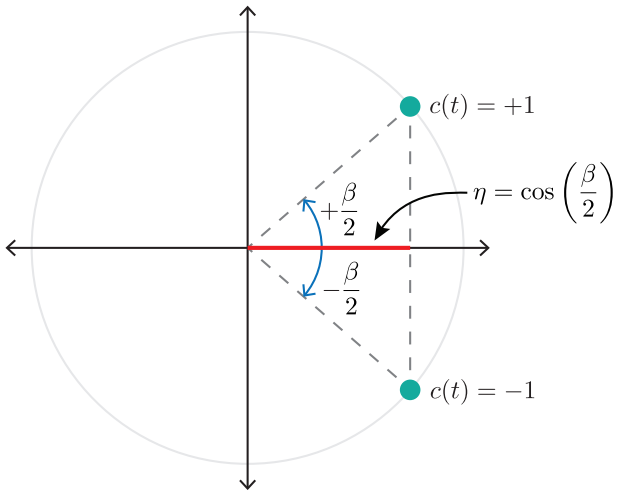

Figure 4.17: Phasor diagram illustrating the relationship between modulation depth β and approximate amplitude of , which diminishes to zero as β approaches π.

4.10 Summary

The proof-of-concept OPA presented in this chapter has validated the concept of internal sensing by coherently combining three emitters in the far field without the use of any external sampling optics to stabilise the array. The output phase of each emitter was inferred by detecting the small fraction of light reflected back into the fibre at the OPA's glass-air interface at a single photodetector. Individual reflections were isolated using digitally enhanced heterodyne interferometry, and their phase was measured using digitally implemented phasemeters. The relative phase error between emitters was then used to actively stabilise the relative output phase of the OPA using feedback.

The fringe visibility of the far field interference pattern produced by the coherent combination of three spatially separate emitters was measured to be 95%, and the peak intensity of the central interference peak was observed to scale quadratically with the number of emitters. The OPA's RMS output phase stability was measured to be $\Phi_{\mathrm{RMS}} = \lambda/120$ over a 5 Hz bandwidth, satisfying the design requirements defined in Section 4.1.1.

The total output power that can be supported by the proof-of-concept OPA is limited by the damage threshold of sensitive optical elements, which is 100 mW for electro-optic modulators. Whilst the proof-of-concept OPA was never intended to support high optical power, it must be acknowledged that the overall utility of an internally sensed OPA is weak if they are unable to support appreciably high optical powers on the order of 10s of Watts. Another weakness with the proof-of-concept OPA is the need to modulate the phase of the carrier with low modulation depths pseudo-random noise codes to maintain coherent combination in the far field. This had the effect of limiting the signal-to-noise ratio of the demodulated signal, forcing a reduction in phasemeter controller bandwidth to prevent instability.

Overcoming these limitations—in particular the handling of high optical power—is the focus of the offset phase-locked and amplifier-compatible OPA designs presented in Chapters 5 and 6.

Offset phase-locked OPA

5.1 Introduction

The offset phase-locked OPA presented in this chapter extends the power handling capabilities of the proof-of-concept OPA. Asymmetric fibre couplers are included in each emitter to isolate sensitive optical elements from high optical powers, which in this implementation is supplied by multiple offset phase-locked lasers. The use of asymmetric fibre couplers also enables digitally enhanced interferometry to be performed at full π modulation depth, improving the SNR of decoded signals without deteriorating output coherence.

The OPA design presented in this chapter relies heavily on a technique called offset phase-locking (e.g., [47, 98, 99]), which is used to actively stabilise the relative phase of one or more slave lasers to a master laser at an offset frequency. One advantage of offset phase-locking is the ability to set the specific frequency offset of each slave laser relative to the master. This provides a way to precisely characterise the OPA's RMS output phase stability by measuring the phase of the heterodyne beat-note produced by interfering two slave lasers with slightly different frequencies at a photodetector located in the far field.

This chapter begins with the design of the offset phase-locked OPA in Section 5.2, followed by details of its implementation in optics and digital signal processing in Sections 5.3 and 5.4 respectively. Section 5.5 presents a discussion of a transfer function model used to analyse the control system and select appropriate gains parameters. An analysis of experimental results is presented in Section 5.6.

5.1.1 Aims and requirements

The aim of the offset phase-locked OPA was to coherently combine three offset phase-locked slave lasers (each delivering no less than than 100 mW of optical power to the far field) with an RMS output phase stability of $\Phi_{\mathrm{RMS}} \leq \lambda/24$ using internal sensing.

5.2 Design

The architecture of an idealized offset phase-locked OPA is shown in Figure 5.1. Light from a free running continuous wave laser is separated into two paths. The first (upper arm in the figure) is split into three channels, each connected to a $>$ 100 MHz fibre waveguide EOM for PRN phase modulation. The second path forms a reference that is interfered with the back-reflected light originating from the frequency-shifted slave lasers at the return path (RP) photodetector.

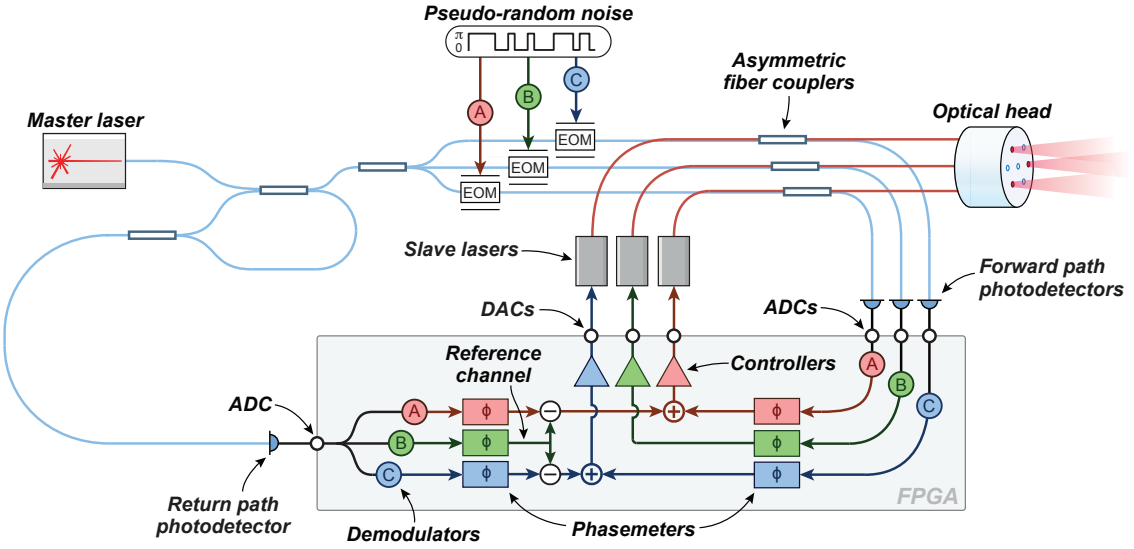

Figure 5.1: Idealised offset phase-locked optical phased array.

Asymmetric fibre couplers are used to combine light from the three slave lasers with the PRN modulated light leaving the EOMs. The asymmetric couplers are configured such that 1% of the light from each slave laser interferes with 99% of the PRN modulated light at a high bandwidth forward path (FP) photodetector. The remaining 99% of the high-power light from the slave lasers travels directly to the optical head where it propagates into free space with along with 1% of the low-power PRN modulated light. When light exits the OPA, a small fraction (∼4%) is reflected back into the fibre due to Fresnel reflection, travelling in the reverse direction back towards the asymmetric fibre couplers. 99% of the back-reflected light travels towards the slave laser where it is suppressed by an optical isolator. The remaining 1% of the back-reflected light passes through the EOMs where it is phase modulated for the first time with a PRN code before being interfered with the unshifted reference at a high-bandwidth return path photodetector.

The forward and return path photodetectors are necessary to sense the phase contributions of all uncommon optical path lengths in the optical system. A measurement of the phase in each optical path length allows us to stabilise the relative output phase of the OPA. The signal at each of the forward path detectors stems from the interference of the high-power and local oscillator paths, and has the form

$$s_{Fn}(t) = A_n \sin\left(\omega_h t + \Phi_{Fn} + \beta c(t - \tau_n)\right) \tag{5.1}$$

where the index n represents the emitter number, A_n is the amplitude of the interference at the detector, ω_h is the angular heterodyne frequency, Φ_{Fn} is the beat-note phase (which represents the phase difference between the two arms of the interferometer at the asymmetric fibre coupler), and β is the modulation depth of the binary PRN code $c(t-\tau_n) \in [0, 1]$ delayed by τ_n. When the modulation depth $\beta = \pi$ we can use the identity $\sin(\theta \pm \pi) = -\sin(\theta)$ to simplify equation (5.1) using the equation $p(t-\tau) = 1 - 2c(t-\tau)$ to map $c(t) \in [0, 1]$ to $p(t) \in [1, -1]$:

$$s_{Fn}(t) = p(t - \tau_n) A_n \sin\left(\omega_h t + \Phi_{Fn}\right)$$

The signal at the return path detector is the sum of the back-reflected signals $r_{Rk}(t)$

interfering with the local oscillator

$$\begin{aligned} s_R(t) &= \sum_k r_{Rk}(t) \\ &= \sum_k p(t-\tau_k) B_k \sin\left(\omega_h t + \Phi_{Rk}\right) \end{aligned}$$

where the index k indicates the emitter number, B_k is the amplitude of the heterodyne interference of the k'th emitter at the detector, and Φ_{Rk} represents the beat-note phase of the k'th emitter.

The signals at each detector are digitized using high bandwidth ADCs, and then digitally demodulated using correspondingly delay matched versions of the code $p(t-\tau)$ for each emitter (as described in Section 2.5). The phase of the demodulated signals is measured using digitally implemented phasemeters, the outputs of which are used to derive feedback control for relative path length stabilisation. A description of how an arbitrarily large number of PRN modulated signals incident on a single photodetector affects the signal-to-noise ratio of the phase measurement is presented in [100] and [67], which indicates that this architecture may potentially be scaled to 10s of emitters using a single return path photodetector.

The control system used to stabilise the relative output phase of the OPA has two stages: i) forward path stabilisation; and ii) return path stabilisation. Feedback for both stages is actuated using each slave laser's piezo-electric transducer frequency control to stabilise the relative laser frequency. The use of a single actuator allows us to use a controller that accepts a linear combination of forward and return path phase error signals to control the phase of each emitter. Digital feedback is converted into analogue voltages using high bandwidth DACs.

5.2.1 Optical phase at the forward and return path photodetectors

The simplified two emitter optical system shown in Figures 5.2 and 5.3 illustrates the various phase contributions from each of the optical path lengths travelled by the light on its way to each detector. The beat-note phase measured at the two forward path detectors is:

$$\begin{aligned} \Phi_{F0} &= \phi_A - (\phi_L + \phi_0 + \phi_a) \\ \Phi_{F1} &= \phi_B - (\phi_L + \phi_0 + \phi_b) \end{aligned}$$

The phase of each channel measured at the return path detector is:

$$\Phi_{R0} = (\phi_A + 2\phi_c + \phi_a + \phi_0 + \phi_R) - (\phi_L + \phi_{LO}) \tag{5.2}$$

$$\Phi_{R1} = (\phi_B + 2\phi_d + \phi_b + \phi_0 + \phi_R) - (\phi_L + \phi_{LO}) \tag{5.3}$$

If we apply feedback to each of the slave lasers to lock the relative phase at the asymmetric fibre couplers, then in the high-gain limit we get

$$\hat{\Phi}_{F0} = \phi_A - (\phi_L + \phi_0 + \phi_a) = 0$$

$$\therefore \phi_A = (\phi_L + \phi_0 + \phi_a) \tag{5.4}$$

$$\hat{\Phi}_{F1} = \phi_B - (\phi_L + \phi_0 + \phi_b) = 0$$

$$\therefore \phi_B = (\phi_L + \phi_0 + \phi_b) \tag{5.5}$$

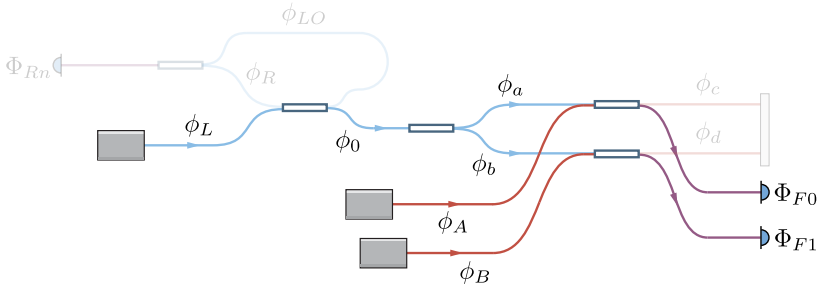

Figure 5.2: Forward path phase contributions from each unique length of fibre in the optical system. For clarity only two fibres are shown.

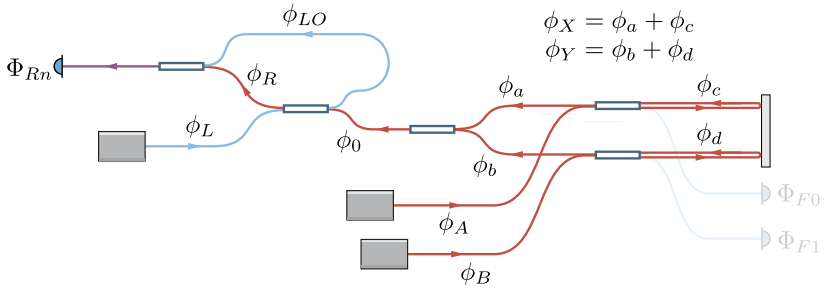

Figure 5.3: Return path phase contributions from each unique length of fibre in the optical system. For clarity only two fibres are shown.

but this only stabilises the relative path lengths up to the asymmetric coupler, and does not take into account the round-trip phase contributions $2\phi_c$ and $2\phi_d$.

Substituting equations (5.4) and (5.5) into equations (5.2) and (5.3) respectively provides the return path phase when the forward paths are locked:

$$\hat{\Phi}_{R0} = 2\phi_a + 2\phi_c + 2\phi_0 + \phi_R - \phi_{LO} \tag{5.6}$$

$$\hat{\Phi}_{R1} = 2\phi_b + 2\phi_d + 2\phi_0 + \phi_R - \phi_{LO} \tag{5.7}$$

The relative phase difference of the return path interferometers is then

$$\begin{aligned}\hat{\Phi}_{error} &= (2\phi_b + 2\phi_d) - (2\phi_a + 2\phi_c) \\ &= 2\phi_Y - 2\phi_X\end{aligned}$$

where

$$\begin{aligned}\phi_X &= \phi_a + \phi_c \\ \phi_Y &= \phi_b + \phi_d\end{aligned}$$

providing us with the information needed to stabilise the remaining uncommon optical path lengths. Combining $\hat{\Phi}_{error}$ with the forward path error signals $\hat{\Phi}_{F0}$ and $\hat{\Phi}_{F1}$ enables us to stabilise the relative path lengths of both emitters, satisfying the requirements for coherent beam combining:

$$\begin{aligned}\hat{\Phi}_{error,YX} = 2\phi_Y - 2\phi_X = 0 \\ \therefore \phi_Y = \phi_X\end{aligned}$$

The same result applies for additional channels, for example a third channel ϕ_Z:

$$\hat{\Phi}_{error,ZX} = 2\phi_Z - 2\phi_X = 0$$
$$\therefore \phi_Z = \phi_Y = \phi_X$$

5.2.2 Offset phase locking

Offset phase-locking is a technique used to stabilise the relative phase of two or more lasers at an offset frequency [98]. It is similar in principle to the operation of a phase-locked loop in that their primary functions are to: 1) detect the phase error between two oscillators; and 2) update the phase of one of the oscillators. Offset phase-locked interferometers are sometimes also referred to as optical phase-locked loops [47].

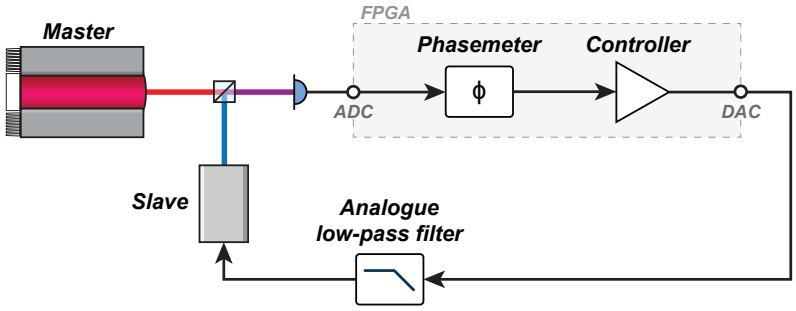

Figure 5.4: Illustration of the optical and FPGA implemented digital signal processing systems used to offset phase-lock two lasers.

A simplified diagram of a digitally implemented offset phase-locked interferometer used in this system is shown in Figure 5.4. Light from a master laser is interfered with a slave laser at a high-bandwidth photodetector to generate a heterodyne beat-note at their difference frequency. The beat-note is digitised using an ADC interfaced with an FPGA. The beat-note phase is measured using a phasemeter, producing an error signal proportional to the phase difference between the master and slave lasers at the beamsplitter. This phase error signal is passed through a controller to generate feedback, which is converted into an analogue voltage using a high-bandwidth DAC. An analogue low-pass filter is used to filter high-frequency noise produced by the DAC [101, 102]. The low-pass filtered control signal actuates the slave laser's frequency to stabilise its phase relative to the master.

5.3 Optical implementation

The offset phase-locked OPA was demonstrated experimentally and characterised using the setup shown in Figure 5.5.

Four free-space Nd:YAG NPRO lasers with a wavelength of 1064 nm were used in this experiment: one as a master, and three as slaves. Light from each laser was coupled from free-space into polarisation maintaining PM980 fibre using fibre collimators. Half-wave plates were used to align the light's polarisation to the slow-axis of the fibre. A half-wave plate and polarising beam-splitter were used to control the amount of optical power entering each fibre. The fibre-coupled power for each of the slaves was set at 150 mW to demonstrate the architecture's ability to operate at optical powers exceeding the damage threshold of the EOMs.

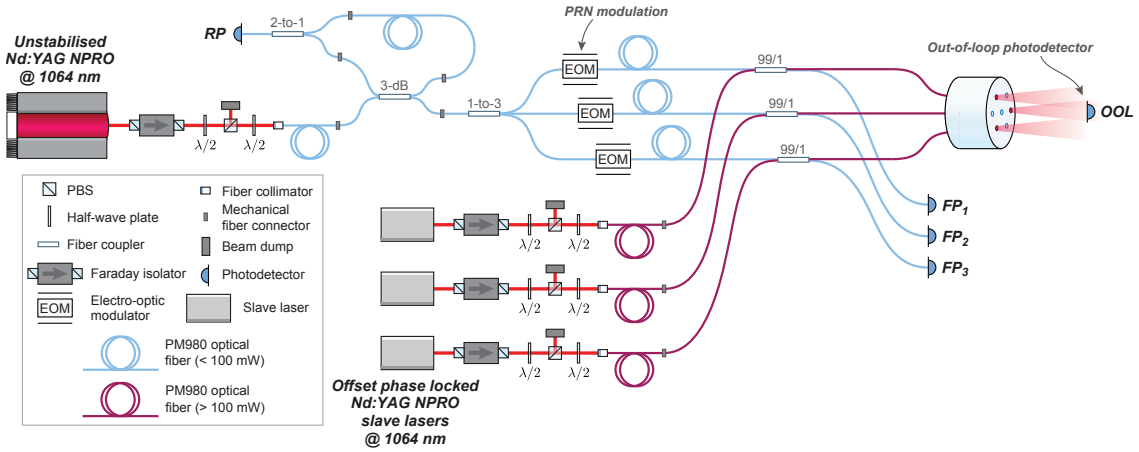

Figure 5.5: Experimental configuration of the optical system used to characterize the offset phase-locked OPA.

Each slave laser can be phase-locked to the master at different offset frequencies, allowing us to: 1) generate heterodyne beat-notes at an out-of-loop photodetector for direct measurement of the OPA's output phase stability; and 2) stabilise the OPA without digitally enhanced heterodyne interferometry, providing a useful baseline from which to gauge the effects it has on Φ_{RMS}.

The heterodyne beat-notes produced by the interference of the three emitters were detected using a commercial free-space 125 MHz bandwidth photodetector (labelled OOL in the figure), placed approximately 10 cm from the optical head. Equivalent fibre-coupled photodetectors were used to detect the interference at the return path and forward path (FP_1, FP_2, and FP_3) detection points. A rail mounted beam-scan (not shown in Figure 5.5) was used to measure the 1D intensity profile of the coherently combined beam at various distances in the far field.

Without DEHI, the light from each channel is identified at FP_1, FP_2, FP_3, and RP via RF demodulation at the corresponding heterodyne beat-notes produced by the three interfering slave lasers. No cross-talk occurs between the various beat-notes produced by the RP detector as the phasemeters apply a band-pass filter with a bandwidth of $\leq$100 kHz centred at the specific RF frequency.

5.3.1 Asymmetric fibre couplers

Opto-Link OLCPL-P-22-106-01-25-FA polarization maintaining fibre couplers with 99/1 (20 dB) splitting ratio and ability to support up to 300 mW CW optical power were used to isolate the high- and low-power sections of the optical system [103]. The relevant technical specifications for the Opto-Link fibre couplers are provided in Table 5.1.

The asymmetric fibre couplers were fusion spliced directly to the PRN phase modulating EOMs and optical head. The particular splitting ratio of the asymmetric couplers is selected to avoid saturating the forward and return path photodetectors; higher laser powers require greater asymmetry.

Table 5.1: Technical specifications for Opto-Link OLCPL-P-22-106-01-25-FA 99/1 fibre couplers.

Parameter	**Specification**	**Unit**
Max. input power	300	mW
Min. splitting ratio	17	dB
Max. splitting ratio	23	dB
Typical splitting ratio	20	dB
Polarization extinction ratio	≥ 20	dB
fibre type	PM980	–

5.3.2 Lasers

The master laser used in this experiment was an Innolight Mephisto 2000-NE. The slaves were a JDSU/Lightwave 126N-1064-700 [104], Innolight Mephisto 200-OEM-NE-ETR, and Innolight Prometheus 20NE [86]. All four lasers operate at a wavelength of 1064 nm with a specified linewidth of < 1 kHz over 100 ms. Their free-running laser frequency noise is proportional to the inverse of the Fourier frequency with an amplitude of $\delta f \approx 10\text{kHz}/\sqrt{\text{Hz}}$ at 1 Hz.

The frequency of each laser can be tuned over a range of approximately 30 GHz by changing the temperature of the laser crystal. This can be done using the front panel of the laser controller, or by applying a voltage to the laser's thermal input. The thermal input tuning coefficient for each laser can be between -3 and -10 GHz/K, where the change in temperature due to an applied voltage is 1 K/V. The thermal response bandwidth is less than 1 Hz.

Fast tuning of the laser frequency can be achieved by applying a voltage signal to the laser's PZT input. Each laser's PZT tuning coefficient can be anywhere between 1 and 3 MHz/V with a response bandwidth of approximately 100 kHz [86].

5.4 Digital Signal Processing

The FPGA implemented measurement and feedback control system used to stabilise the relative output phase of the offset phase-locked OPA is shown in Figure 5.6.

The beat-notes produced by each photodetector are digitised at 40 MHz using 16-bit ADCs. The digitised beat-notes are decoded with delay matched PRN codes. Phasemeters are used to measure the phase of the decoded signals. The phase of the second slave laser SL_2 is the reference from which the return path phase errors (highlighted blue and green in Figure 5.6) are calculated.

The measured phase leaving the two forward path phasemeters (ϕ_{FP1} and ϕ_{FP2}) are combined with the phase error signals produced by the return path phasemeters (highlighted green and magenta in Figure 5.6), producing error signals needed to stabilise the relative output phase of the array. Each combined error signal is passed through a 'fast' controller to generate high-bandwidth (up to 500 kHz, limited by the analogue bandwidth of the DAC) feedback. The fast feedback signals are converted into analogue voltages using 1 MS/s DACs, which are low-pass filtered to suppress high-frequency noise produced by the

DACs[1], which includes harmonic distortion, glitch noise, and quantisation noise [101, 102]. The outputs of the LPFs are connected directly to each laser's PZT input.

Slow controllers are used to provide long term stability by preventing the fast PZT control signals from exceeding the ±10 V range of the DACs in response to large frequency errors. Large, slow frequency disturbances are therefore corrected by the slow controller, whereas lower magnitude, high-frequency disturbances are corrected by the fast controller. The slow controller works by actuating the laser's temperature by applying a voltage to its thermal input. The feedback response bandwidth of the thermal input is typically $<$ 0.5 Hz. The slow controller feedback signal is combined with a manually tunable laser temperature control that is used to bring the laser's interference beat-notes within range of the photodetectors. The thermal control signals are converted into analogue voltages using 1 MS/s DACs.

5.4.1 Dual FPGAs

Two NI PXI-7852R FPGA platforms were required to implement the full measurement and feedback control system for the offset phase-locked OPA, primarily because each FPGA was only able support up to two ADCs. The allocation of DSP resources is shown in Figure 5.6, identifying the two PXI-7852R FPGAs ($FPGA_1$ and $FPGA_2$); the real-time computer (RT); and the FPGA's external interfaces (ADCs, DACs, and DOs). Both FPGAs were mounted in the same chassis, which automatically phase locks their clocks to a common 10 MHz back-plane. The PXI chassis includes eight trigger lines that can be configured to synchronise events (e.g., resets) across different platforms. This was particularly important for synchronising the PRN generators on the two FPGAs.

$FPGA_1$ was used to measure the phase of the first and second forward path channels (FP_1 and FP_2); and derive feedback control signals to stabilise the phase of the first and second slave lasers (SL_1 and SL_2). $FPGA_1$ has two ADC inputs (ADC_1 and ADC_2), and four DAC outputs (DAC_1, DAC_2, DAC_3, and DAC_4).

$FPGA_2$ was used to measure the phase of the second forward path channel (FP_2) and the three return path (RP) signals; generate feedback to stabilise the phase of the second slave laser (SL_2), and generate the PRN codes used to modulate the phase of each emitter. $FPGA_2$ has two ADC inputs (ADC_3 and ADC_4); two DAC outputs (DAC_3 and DAC_4); and three digital outputs (DO_1, DO_2, and DO_3).

Data was transferred between $FPGA_1$ and $FPGA_2$ using the NI PXI-8133e real-time computer. Data from one FPGA was written directly to the RT's shared memory and simultaneously accessed by the other FPGA. This data-sharing routine was executed in real-time at 1 MHz, resulting in a consistent latency of approximately 10 μs, measured by halving the time taken for a pulse to travel from one FPGA to the other and back again. This 10 μs latency did not appear to affect the controller's ability to stabilise the OPA since the transfer rate was much higher than the control system's bandwidth; a 10 μs delay at the FPGA's sampling rate of 40 MHz introduces an 18° phase delay at 10 kHz.

Some NI FPGA platforms facilitate FPGA-to-FPGA data transfer using a standard peripheral component interface express (PCIe) bus, reducing latency and guaranteeing loss-less data transfer. The PXIe-1082 chassis used in this system does not support this capability, however more modern platforms do and are able to support up to 16 FPGAs in parallel (e.g., PXIe-1085 [105]).

[1]Low-pass filters were not originally included in the design, but the slave laser's sensitivity to noise at high frequencies made them necessary.

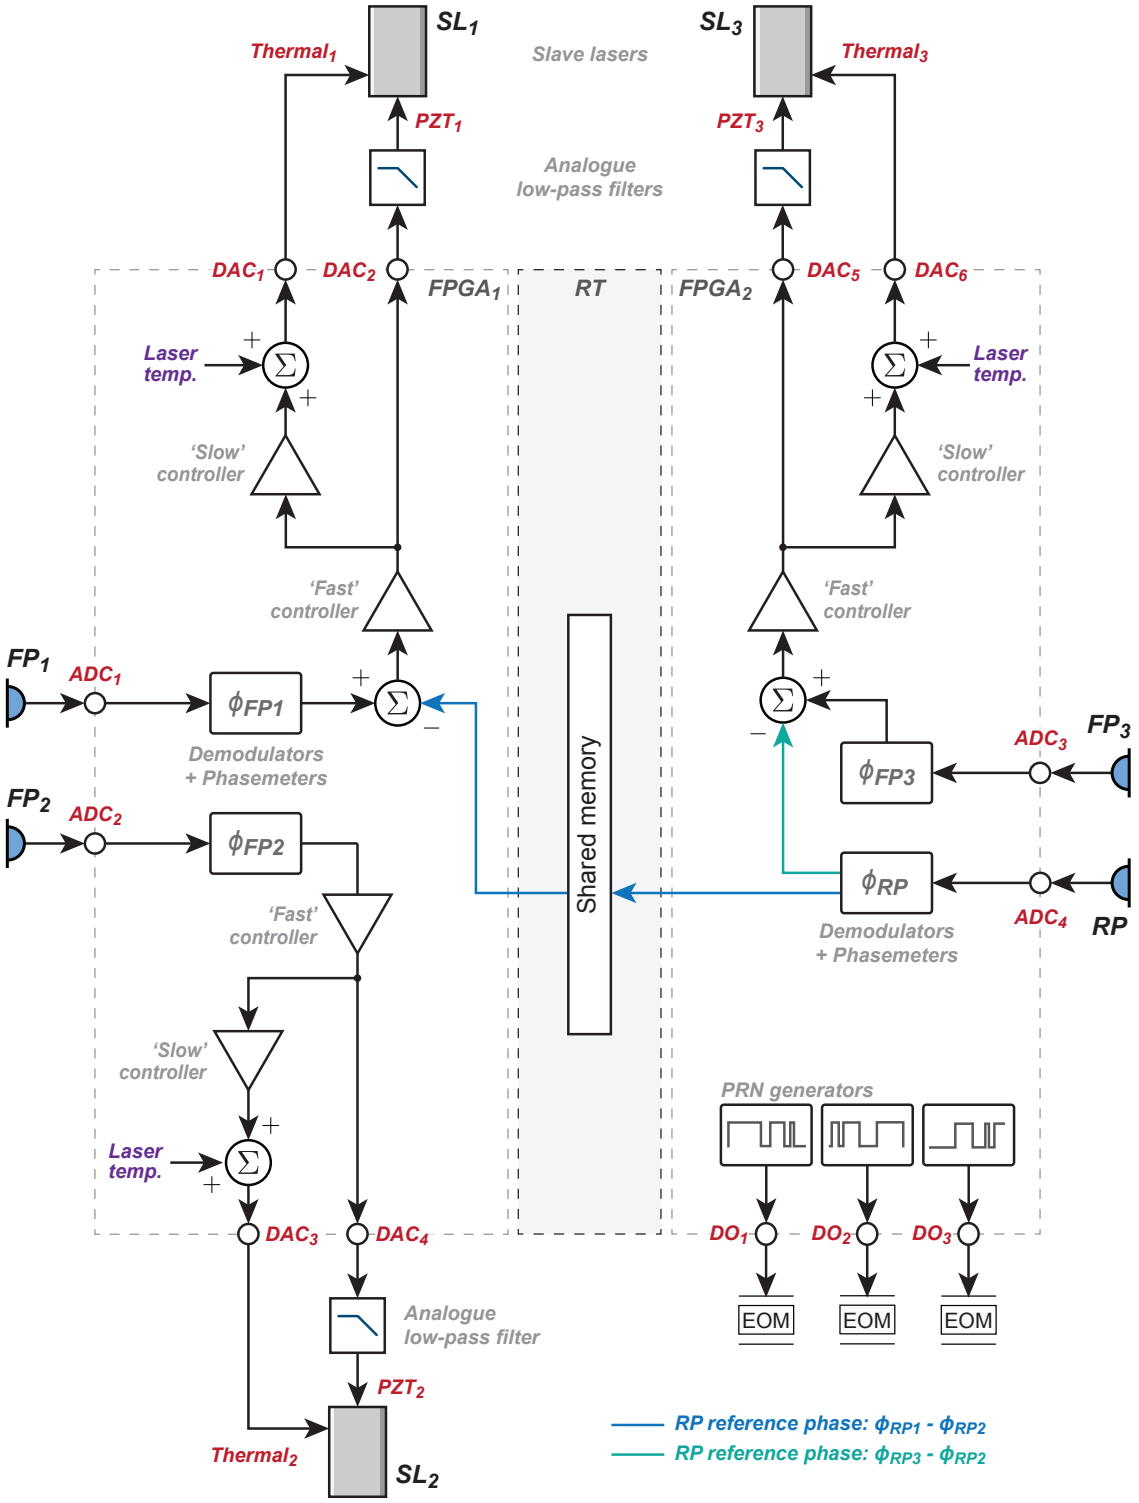

Figure 5.6: Digital signal processing architecture for the offset phase-locked OPA showing the shared allocation of DSP resources between two NI PXI-7852R FPGAs. The FPGA's clocks are phase-locked to a common 10 MHz reference. Data is transferred from one FPGA to the other using the real-time (RT) computer's shared memory.

5.4.2 Phasemeter implementation for offset phase-locking

In offset phase locking, the phase error $\Delta\phi$ is measured with respect to the desired offset frequency, f_{set}. If the absolute frequency of the heterodyne beat-note is f_h, then the absolute phase error with respect to f_{set} is:

$$\Delta\phi = \int_{\tau=0}^{\infty} \Delta f \; d\tau = \int_{\tau=0}^{\infty} \left(f_{set} - f_h\right) \; d\tau \tag{5.8}$$

The phase error $\Delta\phi$ is measured using the PLL shown in Figure 5.7, where the absolute heterodyne frequency f_h is subtracted from the desired set frequency f_{set} prior to integration.

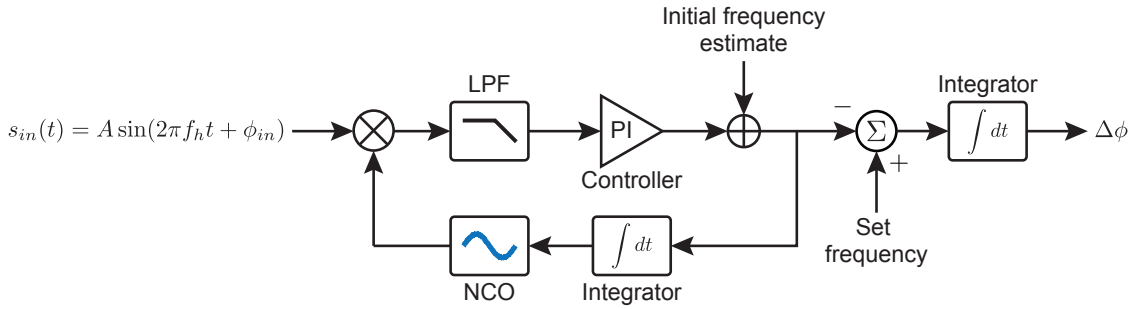

Figure 5.7: Modified phase-locked loop used to calculate phase error with respect a desired offset frequency, f_{set}.

Phasemeter auto-acquisition

One of the challenges with offset phase-locking is estimating the initial beat-note frequency produced by two interfering free running lasers, whose frequencies fluctuate proportional to their laser frequency noise. If the frequency difference between the input signal and PLL's initial frequency estimate exceeds the loop's filter bandwidth, then the resulting error signal will be suppressed, preventing the PLL from acquiring lock. It is therefore important to estimate the frequency of the input signal to within the PLL's forward-loop controller bandwidth, which in this case is approximately 100 kHz. To do this, we created an auto-acquisition routine for the phasemeter using a 4096 point fast Fourier transform (FFT) and a simple peak search algorithm to identify the strongest frequency component in the demodulated signal. The phasemeter's initial frequency could then be set to this value automatically for acquisition. The 4096 point FFT has a frequency resolution of ~ 10 kHz, producing frequency estimates within the lock-in range of the phasemeter.

Aliasing of the heterodyne beat-note frequency

An issue originating from the discrete-time nature of this offset phase locker is aliasing, which causes signals oscillating faster than the Nyquist frequency f_N to appear at frequencies between DC and f_N. This leads to an ambiguity in frequency that prevents us from determining the true frequency of the digitised beat-note, which is a problem since coherent beam combining requires each emitter to operate at the same frequency. This was overcome using a high-bandwidth spectrum analyser to check the true beat-note frequencies at the forward path photodetectors.

5.4.3 Fast controller

The offset phase-locked OPA's fast controller (shown in Figure 5.8) consists of a pre-gain, a single proportional-plus-differential stage, two proportional-plus-integral stages, and a global gain, all connected in series.

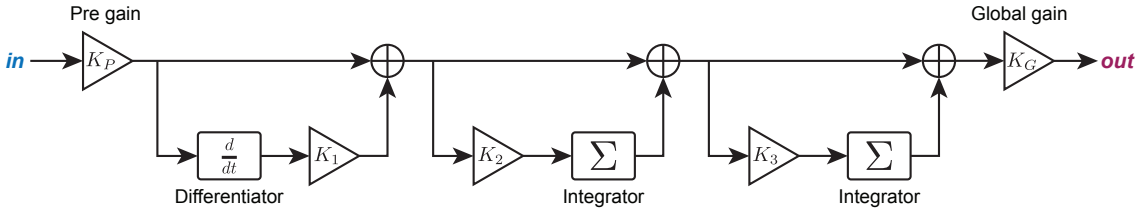

Figure 5.8: Control architecture used in the offset phase-locked OPA.

The pre-gain K_P reduces the amplitude of the input signal during initial lock acquisition to prevent saturating the integrators during wind up[2]. All three proportional-plus-integral and proportional-plus-differential control stages have unity proportional gain, allowing the system's open loop gain to be scaled using K_P and the global gain K_G, without having to adjust the differential and integral gains (K_1, K_2, and K_3).

The two proportional-plus-integral stages provide increased gain at low frequencies to suppress free-running laser frequency noise. The integral gains K_1 and K_2 are implemented as bit-shift dividers, and are situated before the integrators to prevent saturation. The transfer functions of the two proportional-plus-integral controllers are

$$K_2(z) = 1 + \frac{K_2}{1 - z^{-P}} \tag{5.9}$$

and:

$$K_3(z) = 1 + \frac{K_3}{1 - z^{-P}} \tag{5.10}$$

The z^{-P} term in equations (5.9) and (5.10) indicates that the phase error is integrated at the decimated sampling rate f_s/P, where f_s represents the sampling rate of the FPGA, and P is number of points in the phasemeter's low-pass filter (see Section 3.2.5).

The proportional-plus-differential stage is used to compensate for the reduced gain and additional phase delay caused by the analogue LPF. The transfer function of the proportional-plus-differential controller is:

$$K_1(z) = 1 + K_1\left(1 - z^{-P}\right) \tag{5.11}$$

The overall transfer function for the fast controller is:

$$\begin{aligned} C(z) &= K_P K_1(z) K_2(z) K_3(z) K_G \\ &= K_P K_G \left(1 + K_1\left(1 - z^{-P}\right)\right)\left(1 + \frac{K_2}{1 - z^{-P}}\right)\left(1 + \frac{K_3}{1 - z^{-P}}\right) \end{aligned} \tag{5.12}$$

[2]Wind up refers to the situation when an integral controller accumulates rapidly in response to a large and abrupt change in set-point, for example when the controller is first activated.

5.4.4 Analogue low-pass filter

The analogue low-pass filter is used to suppress high-frequency DAC noise. It's transfer function can be approximated in the discrete-time domain as

$$F(z) = \frac{1 - e^{-\frac{T_s}{RC}}}{1 - e^{-\frac{T_s}{RC}} z^{-1}} \tag{5.13}$$

where T_s is the sampling period of the digital system, and R and C are the analogue filter's resistance and capacitance measured in Ohms and Farads respectively. The low-pass filter used in this implementation is a first-order parallel RC circuit with resistance $R = 2.4$ kΩ and capacitance $C = 22$ nF. The LPF's cut-off frequency is $f_c = 3$ kHz (i.e., the LPF has a gain of 0.5 (or -3 dB) at 3 kHz).

Compensating for the low-pass filter

Placing a first-order LPF at the output of the DAC introduces a pole at its cut-off frequency, $f_c = 3$ kHz, which had to be compensated to prevent instability at the desired 10 kHz unity gain frequency. This was achieved by introducing a zero at f_c using a proportional-plus-differential controller with an appropriate differential gain, K_1. Given we know that the LPF's gain at 3 kHz is 0.5, we can estimate the necessary differential gain K_1 required to compensate the LPF by evaluating the expression

$$\begin{aligned} |L(f)K_1(f)| &= 1 \\ \left|1 + K_1\left(1 - e^{j2\pi\frac{Pf}{f_s}}\right)\right| &= \left|\frac{1}{L(f)}\right| \\ \therefore K_1 &= \left|\left(\frac{1}{1 - e^{j2\pi\frac{Pf}{f_s}}}\right)\left(\frac{1}{L(f)} - 1\right)\right| \end{aligned} \tag{5.14}$$

at the known cut-off frequency $f = 3$ kHz when $|L(f)| = 0.5$. The FPGA sampling frequency in this system is $f_s = 40$ MHz, and the rate decimation through the phasemeter's is $P = 40$. Substituting these parameters into equation (5.14) yields an ideal differential gain of $K_1 = 53$.

The magnitude and phase response of the analogue LPF, proportional-plus-derivative controller, and resulting compensated transfer functions were simulated in the frequency domain using equations (5.13) and (5.11) with parameters $f_s = 40$ MHz, $P = 40$ points, $R = 2.4$ kΩ, $C = 22$ nF, and differential gain $K_1 = 53$. The results of these simulations are shown in Figure 5.9, along with the measured magnitude and phase response of the analogue LPF which was recorded using a network analyser from 10 Hz to 100 kHz. The results show that the proportional-plus-differential controller effectively compensates the LPF's transfer function up to 100 kHz with a gain of $K_1 = 53$.

5.4.5 Slow controller

The slow controller worked by decreasing or increasing the laser temperature incrementally whenever it detected that the output of the fast controller exceeded a positive or negative user-defined threshold. It is important to recall here that increasing the laser temperature *decreases* its frequency within the free spectral range [106].

If the slow controller detects that the output of the fast controller has exceeded the positive threshold, then it responds by *increasing* the slave laser's temperature to reduce

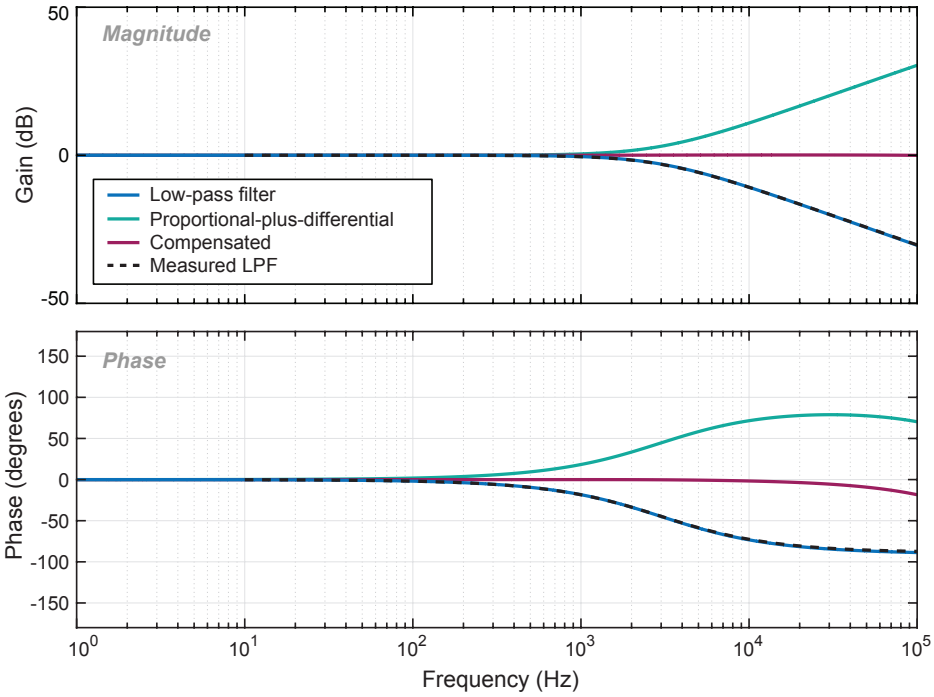

Figure 5.9: Magnitude and phase response of the analogue LPF (blue), proportional-plus-differential controller (green), and compensated transfer function (magenta). The dashed black line represents the measure magnitude and phase response of the analogue LPF.

its frequency. The fast controller responds to this disturbance by moving away from the threshold towards zero. If the fast controller exceeds the negative threshold, then the slow controller responds by decreasing the laser's temperature. In both cases the desired effect is to pull the PZT control signal away from the threshold back to zero.

The slow controller was configured to poll the fast controller output once every second to account for the slow response time ($<$ 1 second) of the laser's thermal input. In this implementation the thresholds were set to 2^{12} counts (corresponding to DAC output voltages of ± 1.25 V, or input frequency disturbances of approximately ± 3 MHz). This worked well when experimentally implemented since the effective laser frequency excursion produced by a single count at the DAC (which has a voltage resolution of $10/2^{15} = 305$ μV/count) was approximately 1 MHz/count. The slow controller was effective at preventing the fast controller from overrunning the range of the DACs, and was used to stabilise the offset phase-locked OPA for extended periods of time.

5.5 Transfer function model

The transfer function model shown in Figure 5.10 was created to analyse the offset phase-locked OPA's open-loop gain under different loop conditions (e.g., phasemeter bandwidths and fast controller gains).

The transfer functions included in the model are summarised below:

- $P_F(z)$, $P_{R1}(z)$, and $P_{R2}(z)$ represent the forward-loop transfer functions of forward and return path phasemeters defined in equation (3.42) in Section 3.3. They convert the laser's phase in radians to phase in counts.

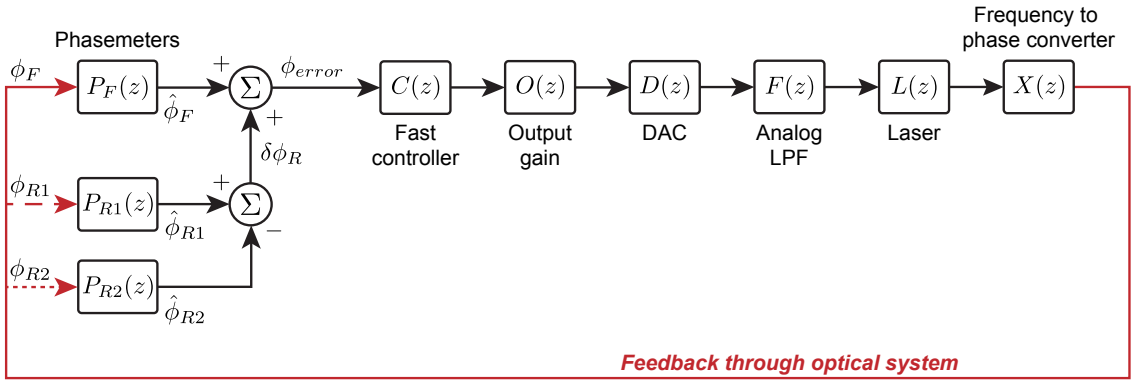

Figure 5.10: Linear transfer function model of the offset phase-locked OPA's control system.

- $C(z)$ represents the fast controller defined in equation (5.12) and converts counts to counts.
- $O(z) = 2^{-16}$ represents the frequency independent gain required to scale the 32-bit output of the controller to the 16-bit amplitude range of the DAC, converting counts to counts.
- $D(z)$ represents the transfer function of the digital-to-analogue converter (defined by equation (4.8) in Section 4.5.3) and converts counts to Volts.
- $F(z)$ represents the analogue LPF defined in equation (5.13) and converts Volts to Volts.
- $L(z)$ represents the PZT tuning coefficient of the laser described in Section 5.3.2 and converts Volts to Hz.
- $X(z)$ represents the frequency-to-phase conversion from laser frequency (Hz) into phase (radians), and is modelled in the discrete-time domain as

$$X(z) = \frac{2\pi}{f_s}\left(\frac{1}{1-z^{-1}}\right) \tag{5.15}$$

  where f_s represents the FPGA's sampling frequency.

Feedback through the optical system (highlighted red in Figure 5.10) represents the laser's phase in radians.

For simplicity we assume that the return path reference phase $\hat{\phi}_{R1}$ is zero, producing a return path phase error $\delta\phi_R$ equal to the phase measured by the second return path phasemeter, ϕ_{R2}. Under this assumption, the open-loop gain is

$$G_{OL}(z) = C(z)A(z)\left[P_F(z) - P_{R2}(z)\right] \tag{5.16}$$

where

$$A(z) = O(z)D(z)F(z)L(z)X(z) \tag{5.17}$$

represents all of the unchanging transfer functions in the system. The remaining transfer functions $P_F(z)$, $P_{R2}(z)$, and $C(z)$ have variable frequency dependent gains that can be adjusted to tune the open-loop gain.

The transfer function model in Figure 5.10 was simulated in the frequency domain using MATLAB. Initial gain parameters for the fast controller were estimated using the approach outlined in Section 3.5.2 to achieve a stable unity gain frequency of 10 kHz. Final gain parameters for the fast controllers ended up being balanced with the forward and return path phasemeter bandwidths, which depend on their respective input signal amplitudes, integral gains, and global gains as described in Section 3.5.

5.5.1 Measured vs. simulated open-loop gain

A Stanford Research Systems SR785 network analyser [107] and SR560 low-noise pre-amplifier [108] were used to measure the offset phase-locked OPA's open-loop gain using the experimental setup shown in Figure 5.11. The Mephisto NE2000 was used as the master, and the Prometheus NE20 was used as the slave.

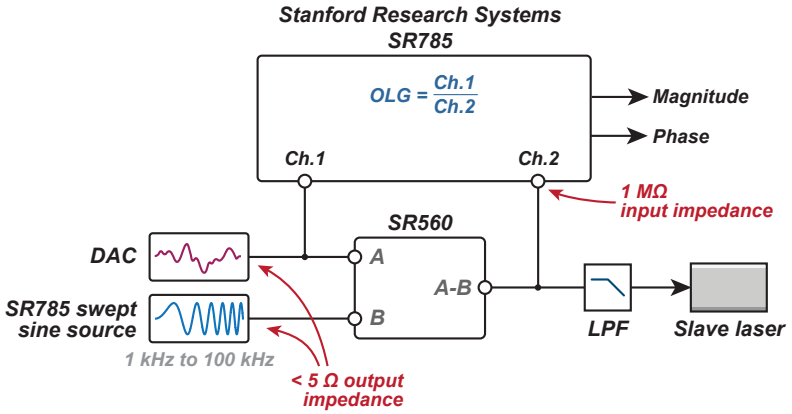

Figure 5.11: Experimental setup used to measure the open-loop gain of the offset phase-locked OPA's control system.

A 50 mV swept-sine wave with a sweep range of 1 kHz to 100 kHz was generated by the SR785 and added to the voltage signal produced by the DAC using the SR560 as a summing junction. The summed output of the SR560 is connected directly to the analogue LPF to close the loop. The output of the DAC and SR560 were connected into the SR785 using BNC tee-junctions. The voltage drop resulting from the tee-junctions is negligible due to the SR785's high input impedance (>1 MΩ) relative to the low impedance outputs of the SR785's swept-sine source and DAC (<5 Ω).

This arrangement introduces a disturbance into the loop that the control systems acts to suppress by generating a corrective voltage signal at the DAC, which was measured on channel 1 of the SR785. The suppressed signal leaving the SR560 is measured on channel 2. Dividing the correction signal by the suppressed signal produces a measurement of the system's open-loop gain.

The simulated and measured open-loop gains are shown in Figure 5.12 for high and low gain settings, which resulted in unity gain frequencies of 20 kHz and 6 kHz respectively. The noise in the measured data below 3 kHz is caused by the strong suppression of the injected disturbance as a result of high gain at low frequencies, which reduces the signal-to-noise ratio of the suppressed signal measured on channel 2.

The simulated OLG calculations assumed measured loop parameters. The PZT tuning coefficient of the Prometheus was measured to be approximately $L(z) = 1.3$ MHz/V. The sampling rate of the FPGA was $f_s = 40$ MHz. The measured input amplitudes at the forward and return path phasemeters were approximately 4000 counts and 2000 counts respectively. The integral and global gains of the forward path phasemeter were set to 2^{-5} and 80. For the return path phasemeter these were 2^{-10} and 1. Both phasemeters used 40 point decimating low-pass filters, reducing the effective sampling rate of the control system to 1 MHz.

The fast controller gains for the 20 kHz bandwidth measurement were $K_P = 2^{-6}$, $K_1 = 53$, $K_2 = 2^{-6}$, $K_3 = 2^{-7}$, and $K_G = 20$. For the 6 kHz bandwidth measurement the gains were $K_P = 2^{-6}$, $K_1 = 53$, $K_2 = 2^{-6}$, $K_3 = 2^{-7}$, and $K_G = 6$. The correlation between global

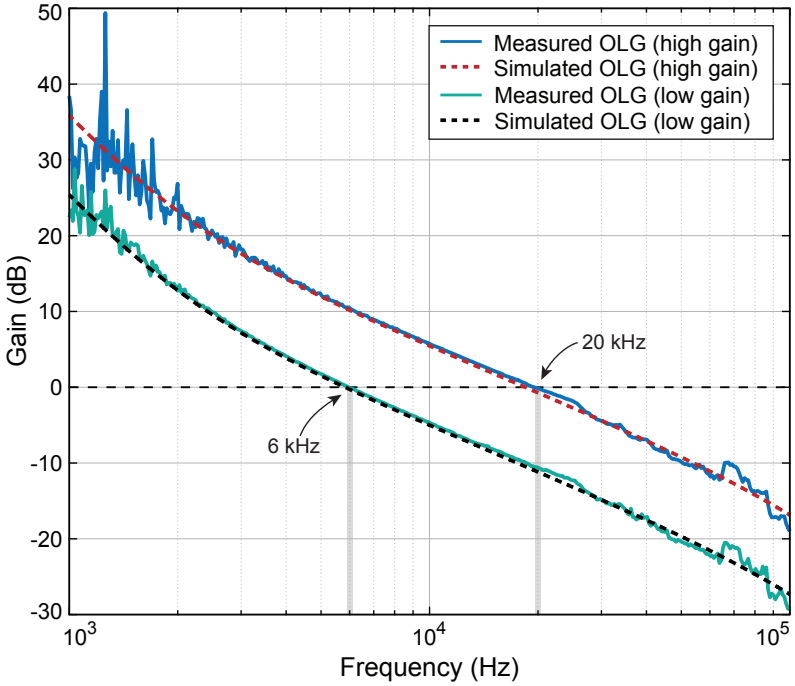

Figure 5.12: Measured (solid) and simulated (dashed) open-loop gain of the offset phase-locked OPA's control system for high and low gains.

gain and unity gain frequency was a coincidence, however it did provide a useful gauge for setting the control system's bandwidth. The measurements agree well with simulation above 2 kHz.

5.6 Results and analysis

5.6.1 Laser frequency stabilisation

The offset phase-lock controller was implemented experimentally to stabilise the relative frequency between a Mephisto NE2000 master and Prometheus NE20 slave. The bandwidth of the controller was set to approximately 10 kHz. The frequency error was recorded in-loop at the output of the phasemeter's controller prior to integration. The measured free-running and stabilised relative laser frequency noise is shown in Figure 5.13. The unstabilised frequency (blue) agrees with each laser's 10 kHz/$\sqrt{\text{Hz}}$/f frequency noise specification. The stabilised laser frequency (green) reveals excellent suppression at low frequencies.

5.6.2 RMS output phase stability

The offset phase-locked OPA's RMS output phase stability was measured using the experimental setup shown in Figure 5.5 by offsetting the relative frequencies of the slave lasers, and then measuring the phase of the resulting heterodyne beat-notes produced at an external photodetector. Whilst this measurement introduces a frequency difference between the emitters, the OPA's phase coherence is maintained. The measurement is performed out-of-loop to capture the true relative phase error between the two emitters, as neither the forward- nor return path photodetectors contain this information. The slave lasers

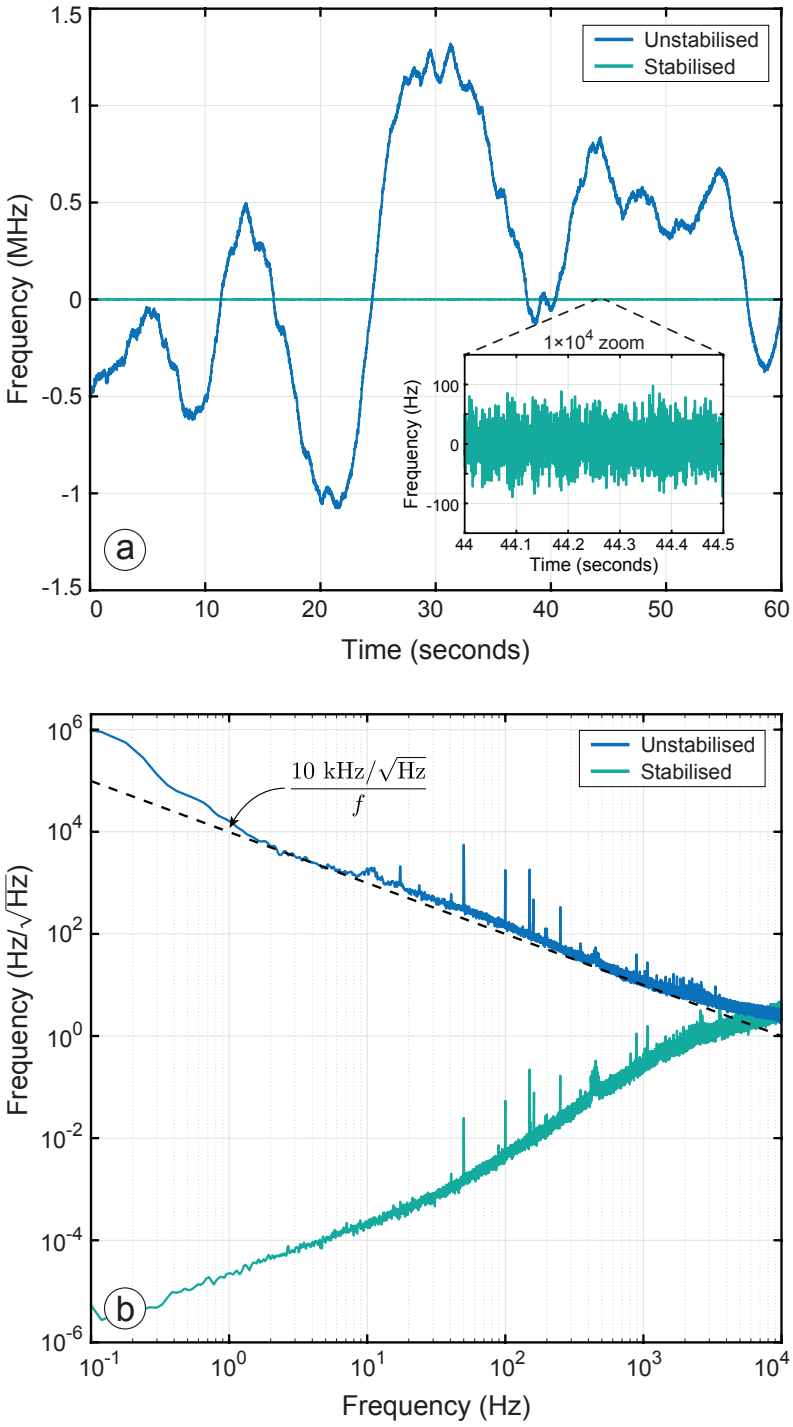

Figure 5.13: Measured free-running and stabilised laser frequencies.

were set to 4, 8, and 10 MHz relative to the master, producing heterodyne beat-notes at 2, 4, and 6 MHz at the OOL photodetector.

The OOL phase error ϕ_{OOL} was initially measured without digitally enhanced heterodyne interferometry to provide a baseline from which to gauge the effect DEHI has on ϕ_{OOL}. This was done using RF demodulation at the return path photodetector. The time series and root-power spectral density of ϕ_{OOL} measured *without* DEHI in shown in Figure 5.14. The same measurement performed *with* DEHI using 15- and 24-bit PRN codes are shown in Figures 5.15 and 5.16 respectively. The OOL phase was recorded at a rate of 32.250 kHz with second order cascaded integrator-comb anti-aliasing filters.

Two measurements are presented in each figure: 1) when the forward path phase error is locked (blue); and 2) when both the forward and return path phase errors are locked (green). The time series data when both controllers are locked has been shifted to appear at zero phase for clarity. None of the time-series data has been de-trended.

When the forward path controllers are locked (blue), the OPA's output phase fluctuates randomly due to fibre noise (caused by random fluctuations in optical path length and refractive index within the glass fibre, visible below 10 Hz) and laser frequency noise (coupled in by macroscopic optical path-length differences between the two emitters, visible between 10 and 100 Hz). The remaining uncontrolled lengths of fibre are stabilised when the return path controller is activated (as described in Section 5.2.1).

The RPSD of the fully stabilised output phase error in Figure 5.14(b) reveals the out-of-loop noise floor without DEHI. When DEHI is used there is a dramatic increase in harmonic distortion visible in Figures 5.15(b) and 5.16(b) caused by imperfect demodulation of the 15- and 24-bit PRN codes. The distortion in Figure 5.15(b) consists of peaks separated by 1.221 kHz, equal to the code-repetition rate of the $2^{15} - 1$ element long PRN code (40 MHz / $(2^{15} - 1)$ = 1.221 kHz). The same effect appears in Figure 5.16(b), which shows peaks separated by 40 MHz / $(2^{24} - 1)$ = 2.38 Hz. The controller's attempt to suppress residual PRN noise results in it being imposed onto the output phase of each emitter.

The RMS phase error for each measurement was calculated over a 15.625 kHz bandwidth up to 1 second. The RMS phase error without DEHI was calculated to be $\lambda/247$, limited by the bandwidth of the offset phase lock controller. The RMS output phase error with DEHI for 15- and 24-bit codes is $\lambda/104$ and $\lambda/78$ respectively, limited by residual PRN noise caused by imperfect demodulation.

The 'hill' visible between 100 Hz and 10 kHz in each RPSD reveals the noise-floor of the optical phase-locked loop's controller. The noise shelf present in Figure 5.14(b) below 6 Hz is typical of cyclic phase noise caused by parasitic interference at one of the photodetectors, which could have been caused by a partial reflection or interference in an orthogonal polarisation. The reduced amplitude of the noise shelf in Figure 5.14(b) indicates partial suppression by the 15-bit PRN code, suggesting that the cyclic phase noise was caused by a low amplitude reflection in optical fibre outside the modulating PRN code's range-gate (as described in Section 2.5.1). Cyclic noise did not appear to degrade the RMS output phase error, which was instead limited by residual PRN noise at higher frequencies.

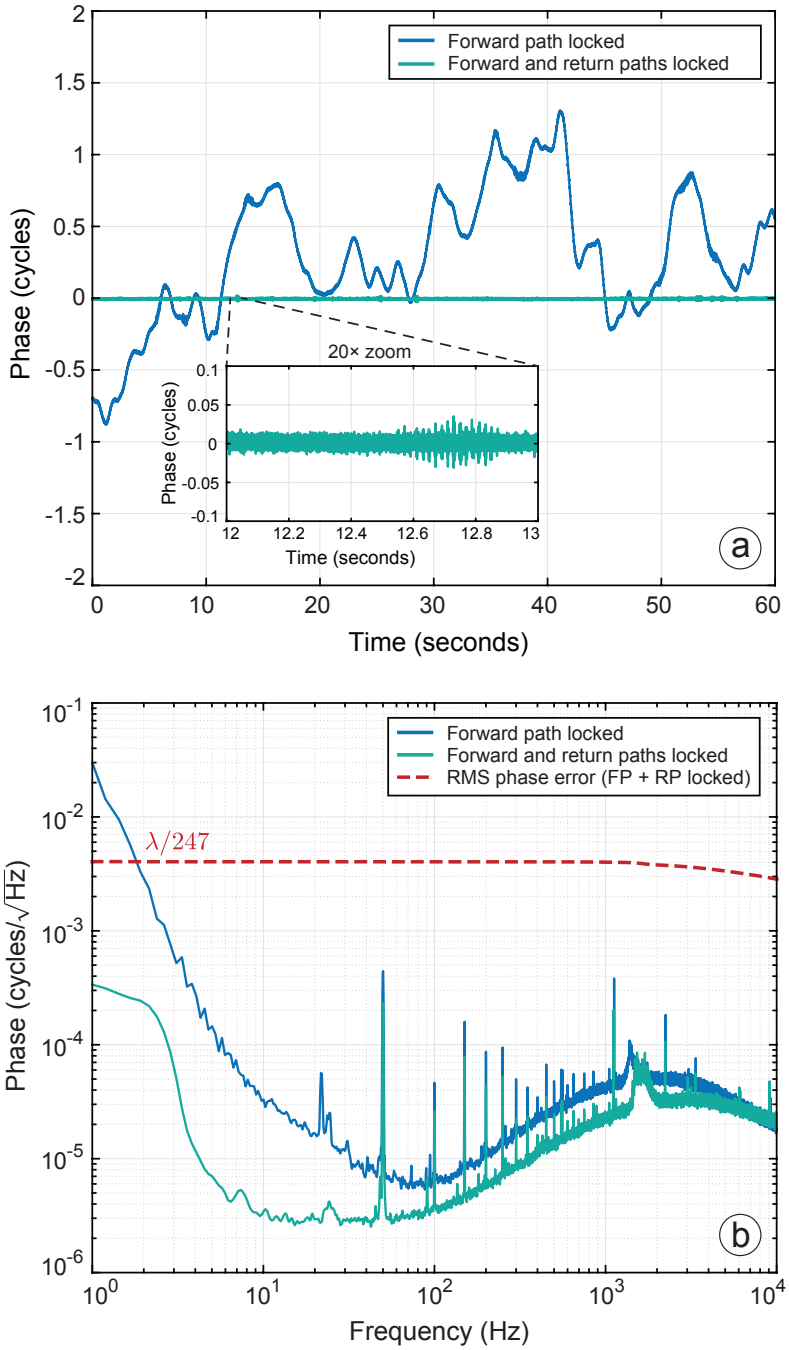

Figure 5.14: (a) Time series measurements of ϕ_{OOL} when the OPA is forward path locked (blue), and fully locked (green) without DEHI. The inset shows the zoomed-in behaviour of ϕ_{OOL} when fully locked. (b) RPSD of the measurements shown in (a); the noise-shelf at frequencies below 10 Hz is typical of cyclic phase noise introduced by parasitic interference somewhere in the optical system. The RMS phase error was calculated to be $\Phi_{\mathrm{RMS}} = \lambda/247$ at 1 Hz, limited by noise at frequencies beyond 1 kHz.

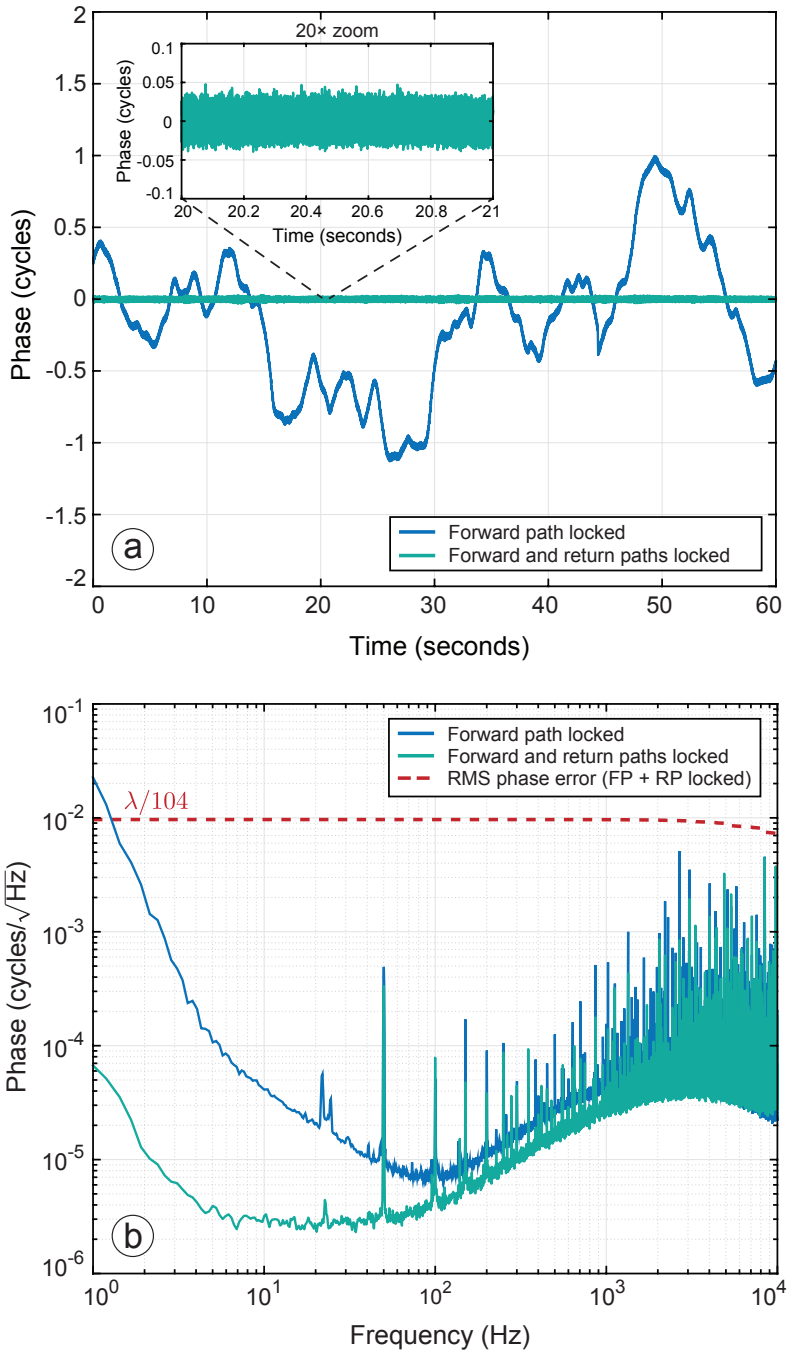

Figure 5.15: (a) Time series measurements of ϕ_{OOL} when the OPA is forward path locked (blue), and fully locked (green) using DEHI with a 15-bit PRN code. The inset shows the zoomed-in behaviour of ϕ_{OOL} when fully locked. (b) RPSD of the measurements shown in (a). The harmonic distortion visible in (b) is caused by residual PRN noise introduced by the demodulator. The RMS phase error was calculated to be $\Phi_{\mathrm{RMS}} = \lambda/104$ at 1 Hz, limited by harmonic distortion introduced by imperfect demodulation of the 10-bit PRN code.

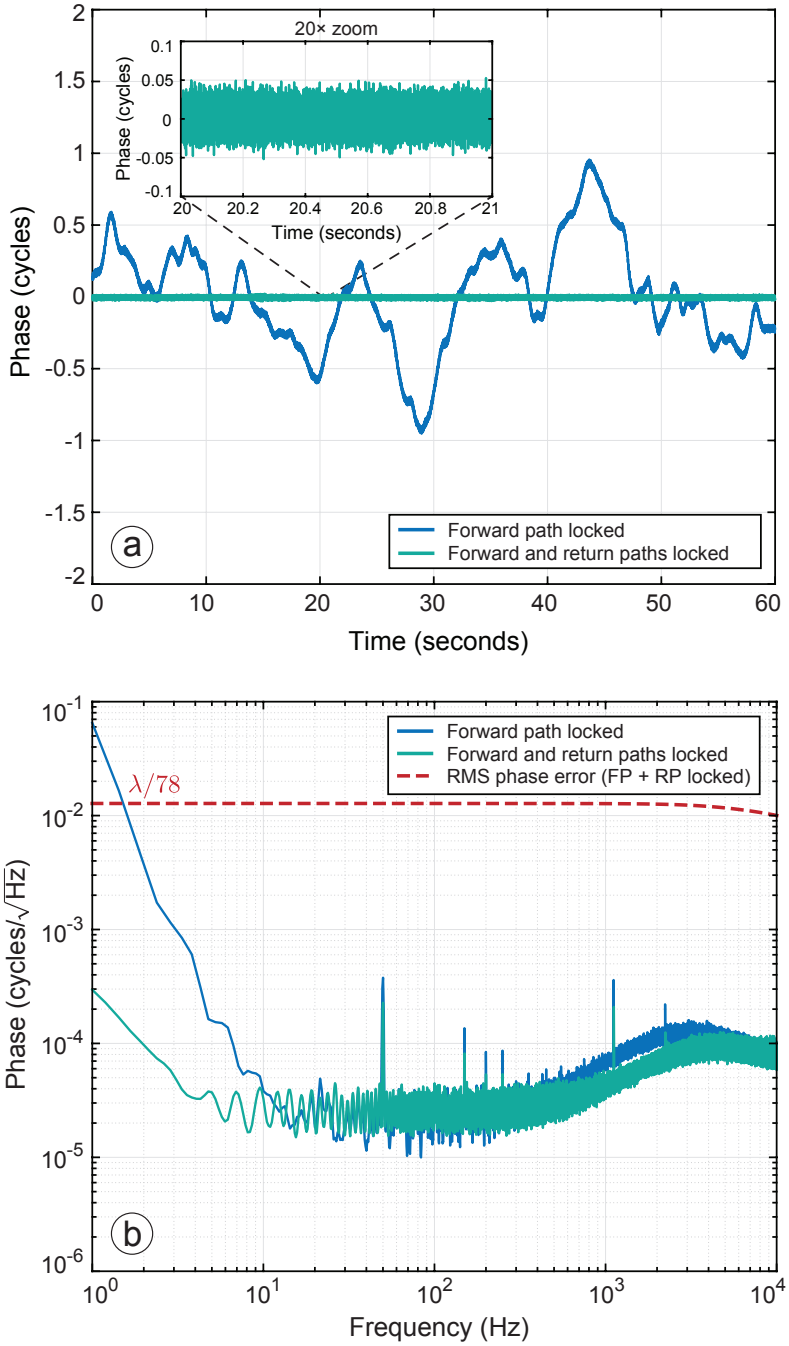

Figure 5.16: (a) Time series measurements of ϕ_{error} when the OPA is forward path locked (blue), and fully locked (green) using DEHI with a 24-bit PRN code. The inset shows the zoomed-in behaviour of ϕ_{RMS} when fully locked. (b) RPSD of the measurements shown in (a). The RMS phase error was calculated to be $\Phi_{\text{RMS}} = \lambda/78$ at 1 Hz, limited by residual PRN noise.

5.6.3 Intensity scaling and fringe visibility

The far field intensity of the offset phase-locked OPA was measured using a Thorlabs beam-profiler for one, two, and three emitter configurations at a distance of 30 ± 0.5 cm. The measured and simulated intensities are shown in Figure 5.17, normalised to the average peak intensity of a single emitter. The simulation assumed a fibre separation of 250 μm, equal emitter power, and a projection range of 30 cm.

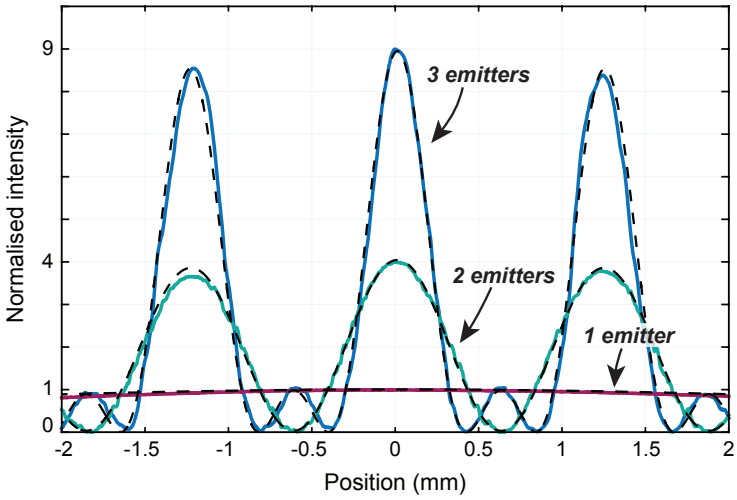

Figure 5.17: Measured (solid) and simulated (dashed) far field interference profiles for one, two, and three emitter configurations. The beam-profiler was located 30 ± 0.5 cm from the OPA.

The peak intensities of the central interference lobes in Figure 5.17 scale quadratically with the number of emitters. Fringe visibility was calculated to be 100% and 99.5% for two and three emitter configurations respectively, indicating strong far field coherence. The measured results agree well with simulation, which are represented as dashed lines.

5.6.4 Beam-steering

The distribution of optical power in the far field can be manipulated by controlling the individual output phase of emitters in the OPA. One way to do this is to adjust the controller's zero point by adding an artificial phase shift to the error signal. The controller will attempt to correct any disturbances within its bandwidth, and will thus apply an equal and opposite magnitude phase shift to suppress the error.

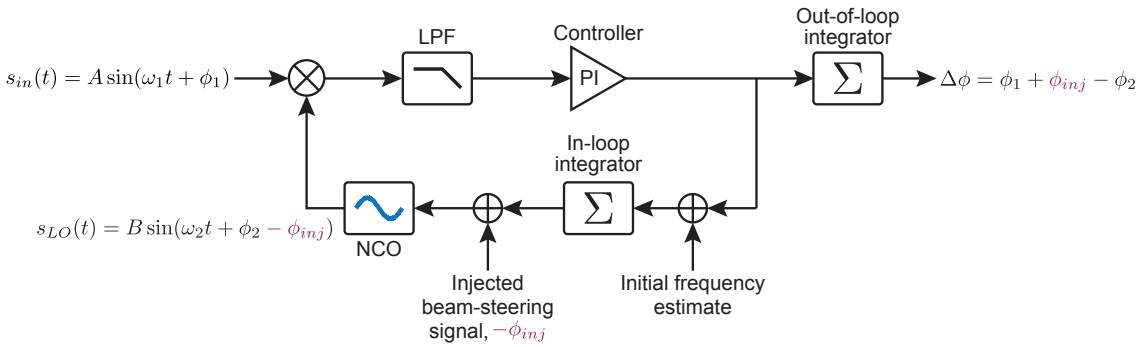

Figure 5.18: Measured magnitude response of beam-steering performed by injecting a swept sine-wave from 1 Hz up to 10 kHz. NCO stands for numerically controlled oscillator.

Beam-steering was demonstrated experimentally by adjusting the set-point of each emitter in array. This was achieved by injecting intentional phase disturbances into both the forward and return path measurements, forcing the controller to impose the injected disturbance onto the output. This is done by subtracting the injected phase disturbance ϕ_{inj} directly from the output of the phase-locked loop's in-loop integrator (as shown in Figure 5.18), resulting in an output phase error that appears indistinguishable from input signal phase. This allowed us to calculate the beam-steering magnitude response by comparing the measured phase at the OOL detector to the injected phase signal.

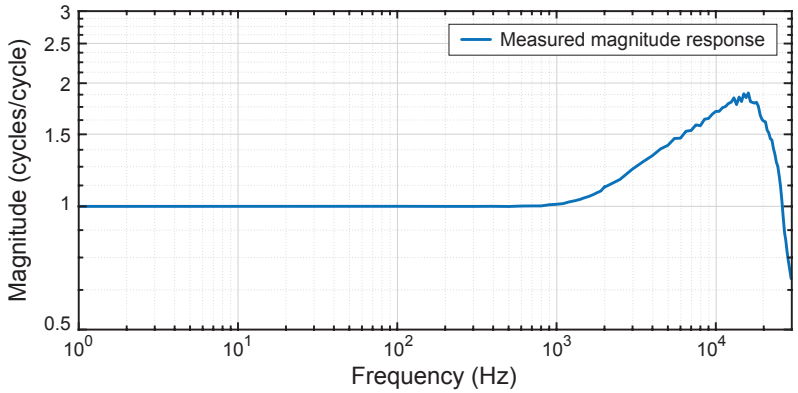

Figure 5.19: Measured magnitude response of beam-steering performed by injecting a swept sine-wave from 1 Hz up to 10 kHz directly into the forward and return path phase measurements.

Figure 5.19 shows a measurement of the beam-steering magnitude response, which was performed using a swept sine-wave (from 1 Hz up to 10 kHz) with an amplitude of 0.5 cycles. The peak that emerges at frequencies greater than 1 kHz is the result of gain peaking in the control loop, which is sensitive to the system's open-loop gain and phase margin. The transfer function shows that the OPA is able to steer the beam at up to 10 kHz. It was observed during testing that the bandwidth over which the output phase of each emitter could be controlled was proportional to the controller's unity gain frequency. The results presented here were performed with a unity gain frequency of 20 kHz.

Figure 5.20 shows a time series of the out-of-loop phase for various injected phase profiles. This test included full cycle steps (between 0 an 10 seconds), a coarse ramp (between 10 and 15 seconds), and a continuous ramp (after 25 seconds), demonstrating our ability to precisely and arbitrarily control output phase of individual emitters.

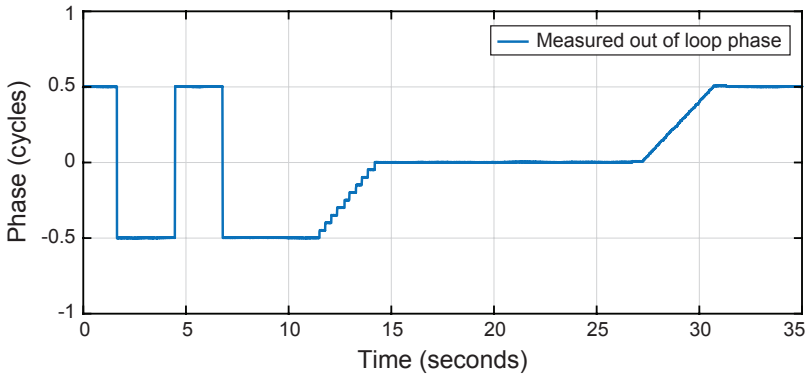

Figure 5.20: Time series measurement of various beam-steering profiles in a two emitter configuration.

5.6.5 Sensitivity to control system bandwidth

It was observed during testing that the RMS phase stability decreased at high phasemeter bandwidths. The reason for this is that higher bandwidth phasemeters measure more residual PRN noise, which is then imposed onto the control signals. This was confirmed by performing three measurements using DEHI with 15-bit PRN codes for low, medium, and high phasemeter bandwidths. The RPSDs of these measurements are shown in Figure 5.21. The RPSDs for each measurement were calculated for equal length data sets to permit meaningful comparison. The amplitude of residual PRN noise is visibly higher at high gain (magenta) than at it is at low gain (blue).

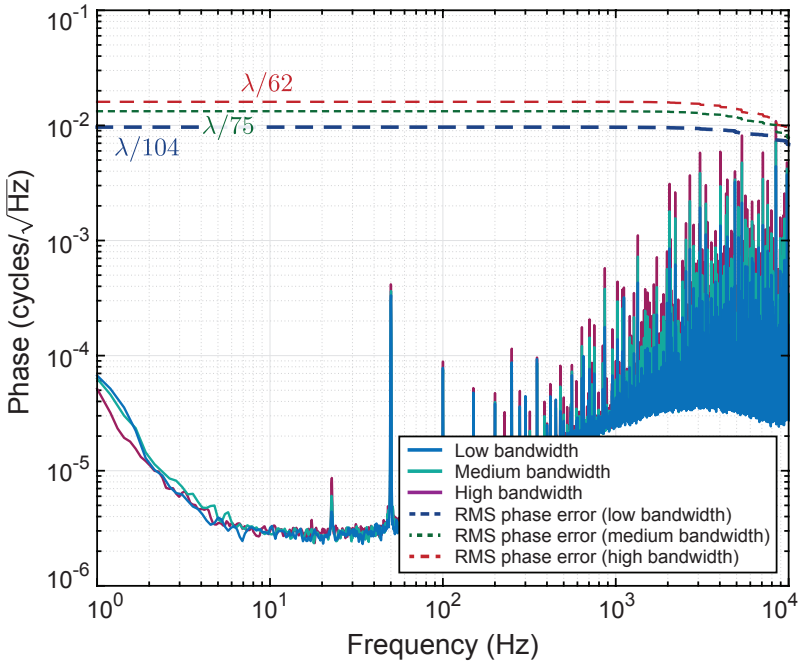

Figure 5.21: RPSDs of the OPA's measured output phase for different control bandwidths.

This observation establishes an interesting trade-off between control bandwidth (which affects the OPA's ability to perform beam-steering) and RMS output phase stability. That being said, the RMS output phase stability Φ_{RMS} at high gain is still well below the $\lambda/24$ requirement, which means it is not necessarily a difficult trade-off to make if high-bandwidth beam-steering is required.

5.6.6 Effect of PRN modulation depth on output phase stability

A significant advantage of the offset phase-locked OPA architecture over the proof-of-concept is that it enables us to perform DEHI at full π modulation depth without decohering the interfered beam. In fact, operating the OPA at modulation depths other than $\beta = \pi$ can severely degrade performance by introducing more residual PRN noise into the control system due to imperfect demodulation. This effect was demonstrated experimentally by recording the out-of-loop phase error for modulations depths $\beta = \pi$, $\pi/2$, and $\pi/4$. The results are shown in Figure 5.22.

The time series data for each measurement has been low-pass filtered to reveal underlying cyclic phase noise that grows in amplitude with increasing modulation depth. This observation supports the suggestion that the noise-shelf below 10 Hz is due to cyclic phase noise caused by parasitic interference somewhere in the interferometer. DEHI's suppression of

spurious interference diminishes with decreasing modulation depth, which explains why the amplitude of the cyclic phase noise increases when $\beta = \pi/2$ and $\pi/4$.

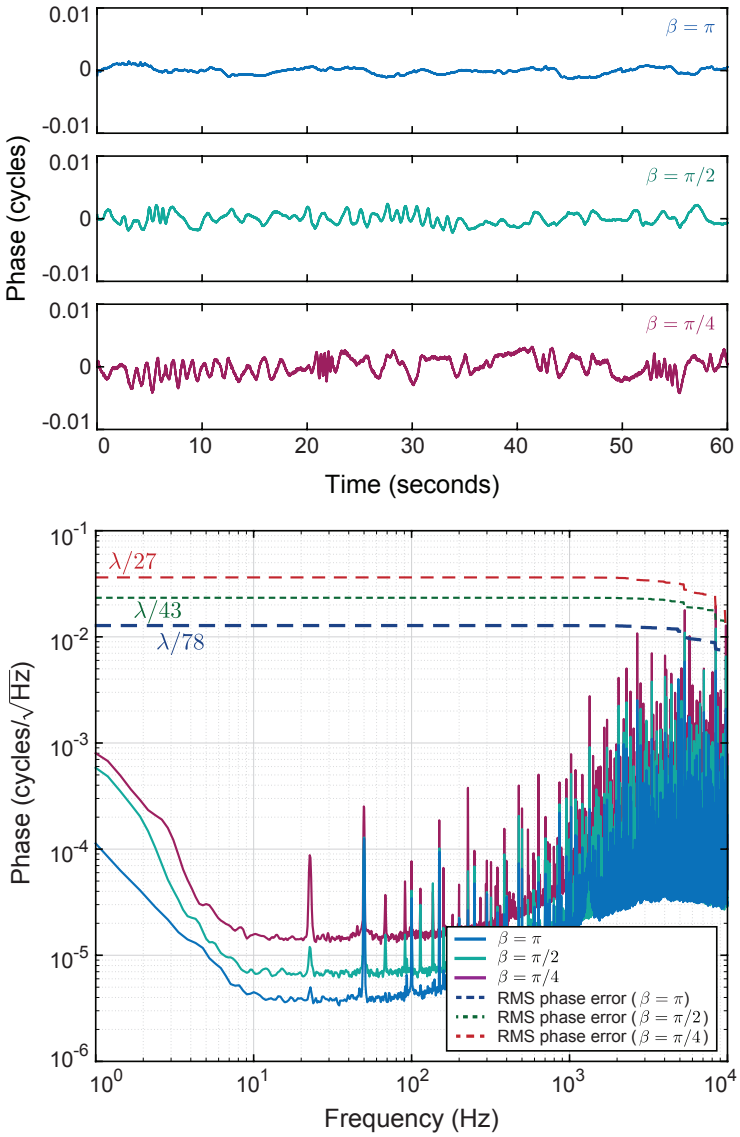

Figure 5.22: Time-series' and RPSDs of the OPA's measured output phase for different PRN modulation depths. The time measurements shown in the upper panel have been low-pass filtered to reveal the underlying cyclic phase noise.

The RPSDs in Figure 5.22 reveal an increasing white noise-floor (between 10 and 100 Hz), as well as a significant increase in harmonic distortion (visible above 1 kHz) as modulation depth decreases. Both of these effects are caused by imperfect demodulation, which reduces the signal-to-noise ratio of the demodulated signal, and introduces higher amplitude residual PRN noise into the control system. It is therefore critical that the PRN modulation depth be set to π in order to maximise Φ_{RMS}, which increases to $\lambda/27$ when $\beta = \pi/4$, close to the $\lambda/24$ performance requirement.

5.7 Summary

The offset phase-locked internally sensed offset phase-locked OPA presented in this chapter demonstrated the ability coherently combine three slave lasers with a total output power of 450 mW. This was achieved using asymmetric fibre couplers. The three slave lasers were phase-locked to a master laser at the asymmetric fibre couplers by interfering them at emitter specific forward path photodetectors. The remaining uncontrolled optical path lengths between the splitters and the output of the array were stabilised by measuring the phase of each emitter's back-reflection at a single return path photodetector. Digitally enhanced heterodyne interferometry was implemented at full π modulation depth to discriminate individual reflections in digital signal processing.

The offset phase-locked OPA's RMS output phase stability was measured to be $\Phi_{\mathrm{RMS}} = \lambda/104$ over a 16.125 kHz bandwidth, far exceeding the $\lambda/24$ design requirement. Fringe visibility of the far field interference pattern was measured to be 100% and 99.5% for two and three emitter configurations respectively, indicating extremely high coherence in the far field. Beam-steering was also demonstrated up to the 10 kHz bandwidth of the controller by injecting artificial phase disturbances directly into the forward- and return path phasemeters. The ability for the OPA to perform high-bandwidth beam-steering suggests that it may be possible to compensate for atmospheric turbulence, which is a necessary requirement for any application involving transmission through the atmosphere (especially remote manoeuvring of orbital debris).

Amplifier-compatible OPA

6.1 Introduction

The offset phase-locked OPA presented in Chapter 5 demonstrated the concept of isolating sensitive optical components from high optical powers using asymmetric fibre couplers, where the relative output phase of the OPA was stabilised by measuring each emitter's phase using a combination of forward and return path photodetectors. The design presented in this Chapter draws on a combination of techniques and ideas introduced in Chapters 4 and 5 to converge on an amplifier-compatible OPA that coherently combines multiple in-line fibre lasers in a master-oscillator power-amplifier configuration. The amplifier-compatible OPA described in this chapter can:

- Support up to 15 W of CW optical power per emitter, with the provision to support up to 100 W
- Steer the coherently combined beam at up to 10 kHz
- Actively suppress stimulated Brillouin scattering (SBS), and
- Scale to a large number of emitters.

This chapter begins with a description of the amplifier-compatible OPA in Section 6.2, followed by its implementation both optically and in digital signal processing in Sections 6.3 and 6.4 respectively. Section 6.5 describes the amplified OPA's compatibility with linewidth broadening techniques for SBS suppression. An analysis of experimental results is presented in Section 6.6.

6.1.1 Aims and requirements

The objective of the amplifier-compatible OPA was to demonstrate the ability to coherently combine three 15 W fibre amplifiers with an RMS output phase stability of $\Phi_{\mathrm{RMS}} \leq \lambda/24$ using internal sensing.

6.2 Design

The architecture of an idealised three-element amplifier-compatible internally sensed OPA is shown in Figure 6.1. Light from a free-running laser is separated into two arms. The first (upper arm in the figure) is split into three channels, all of which are connected to >100 MHz fibre waveguide electro-optic modulators for PRN phase modulation. The second arm is frequency shifted by f_h using a fibre coupled acousto-optic modulator and

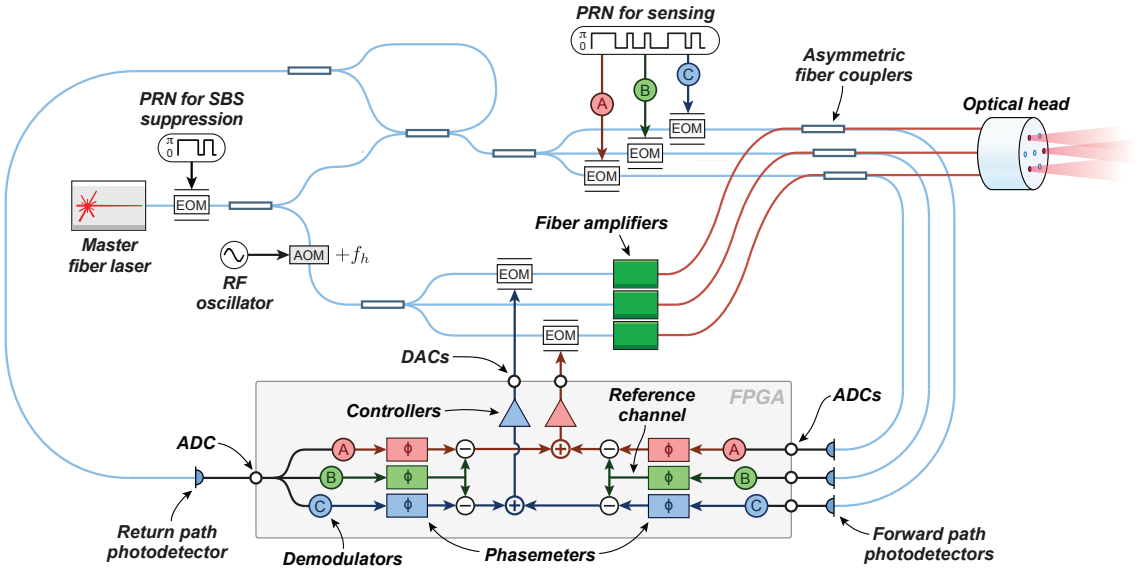

Figure 6.1: Idealised fibre amplifier-compatible internally sensed optical phased array.

then split into three channels, each containing an EOM for feedback actuation and a fibre amplifier.

Pairs of channels from the two arms are combined using commercial 99.9/0.1 asymmetric fibre couplers to form three Mach-Zehnder interferometers. The asymmetric fibre couplers are configured such that 0.1% of the amplified frequency shifted light interferes with 99.9% of the unshifted low power light at a high bandwidth forward path photodetector. The other 99.9% of the amplified frequency shifted light travels directly to the optical head where it propagates into free space. When light exits the OPA ∼4% is reflected back into the fibre due to Fresnel reflection. These reflections travel in the reverse direction back towards the asymmetric fibre couplers. 99.9% of the back-reflected light travels directly to the amplifiers where it is blocked by the amplifier's in-built optical isolator. The remaining 0.1% of the back-reflected light passes back through the first arm, passing through the first series of EOMs where it is phase modulated with PRN before being interfered with an unshifted reference signal at a high-bandwidth return path photodetector.

The signals at each detector are digitised using high bandwidth ADCs, and then digitally demodulated using delay matched versions of the PRN code belonging to each emitter in the array. Digitally implemented phasemeters are used to measure the phase of the demodulated signals to derive feedback control for active relative path length stabilisation.

The amplifier-compatible OPA uses the same forward and return path stabilisation technique as the offset phase locked OPA, except that feedback for the two control stages is actuated using a single EOM positioned before each amplifier, allowing us to use a controller that accepts a linear combination of the forward and return path error signals to stabilise the phase of each emitter. Feedback control can also be actuated using an AOM or piezo-electric fibre stretcher with only minor modification to the control system to account for the different transfer functions of both technologies.

6.2.1 Optical phase at the forward and return path photodetectors

The example two channel optical system shown in Figures 6.2 and 6.3 illustrates the various phase contributions from each of the optical path lengths travelled by the light on

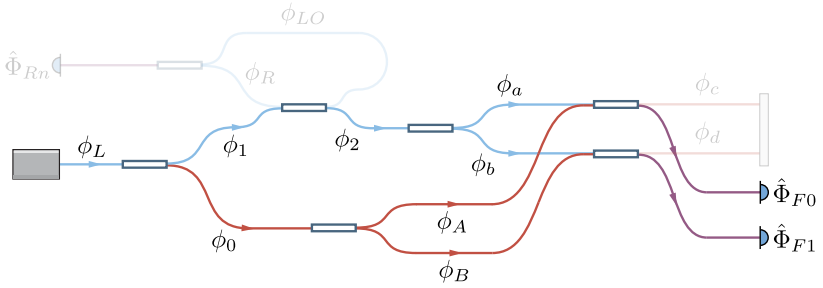

Figure 6.2: Forward path phase contributions from each unique length of fibre in the optical system. For clarity only two fibres are shown.

its way to the forward and return path photodetectors. The beat-note phase measured at the two forward path detectors is:

$$\Phi_{F0} = (\phi_0 + \phi_A) - (\phi_1 + \phi_2 + \phi_a)$$
$$\Phi_{F1} = (\phi_0 + \phi_B) - (\phi_1 + \phi_2 + \phi_b)$$

The phase of each channel measured at the return path detector is:

$$\Phi_{R0} = (\phi_0 + \phi_A + 2\phi_c + \phi_a + \phi_2 + \phi_R) - (\phi_1 + \phi_{LO}) \tag{6.1}$$
$$\Phi_{R1} = (\phi_0 + \phi_B + 2\phi_d + \phi_b + \phi_2 + \phi_R) - (\phi_1 + \phi_{LO}) \tag{6.2}$$

If we apply feedback to the stabilising EOMs to lock the relative phase of the forward paths at the asymmetric fibre couplers, then in the high-gain limit we get

$$\hat{\Phi}_{F0} = (\phi_0 + \phi_A) - (\phi_1 + \phi_2 + \phi_a) = 0$$
$$\therefore (\phi_0 + \phi_A) = (\phi_1 + \phi_2 + \phi_a) \tag{6.3}$$
$$\hat{\Phi}_{F1} = (\phi_0 + \phi_B) - (\phi_1 + \phi_2 + \phi_b) = 0$$
$$\therefore (\phi_0 + \phi_B) = (\phi_1 + \phi_2 + \phi_b) \tag{6.4}$$

but this only stabilises the relative path lengths up to the asymmetric coupler, and does not take into account the round-trip phase contributions $2\phi_c$ and $2\phi_d$.

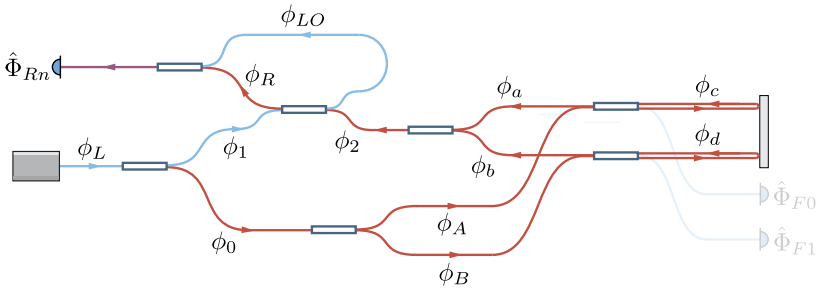

Figure 6.3: Return path phase contributions from each unique length of fibre in the optical system. For clarity only two fibres are shown.

Substituting equations (6.3) and (6.4) into equations (6.1) and (6.2) respectively provides the return path phase when the forward paths are locked:

$$\hat{\Phi}_{R0} = 2\phi_a + 2\phi_c + 2\phi_2 + \phi_R - \phi_{LO} \tag{6.5}$$
$$\hat{\Phi}_{R1} = 2\phi_b + 2\phi_d + 2\phi_2 + \phi_R - \phi_{LO} \tag{6.6}$$

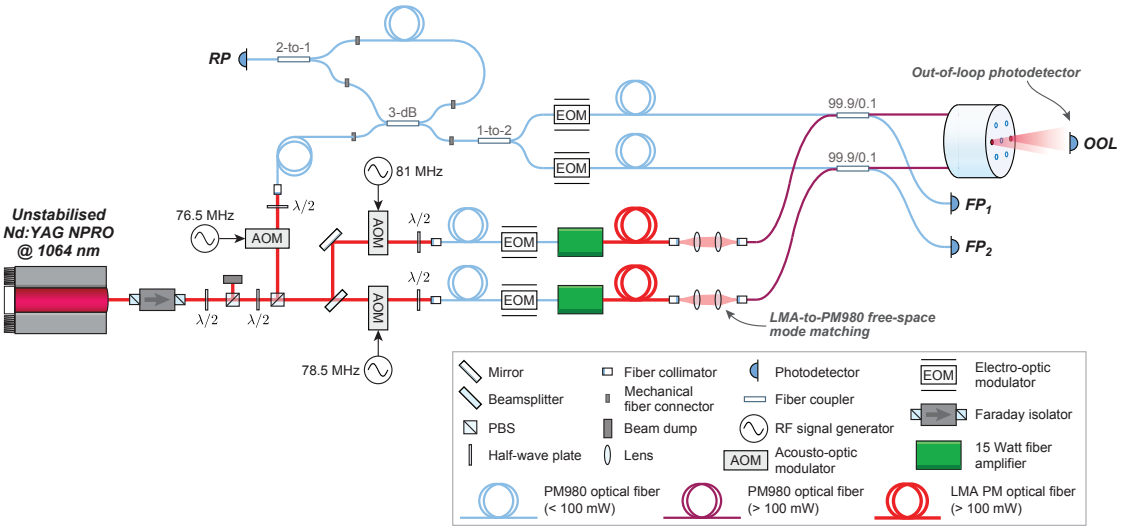

Figure 6.4: Experimental configuration of the optical system used to characterise the amplifier-compatible internally sensed OPA.

The relative phase difference of the return path interferometers is then

$$\begin{aligned}\hat{\Phi}_{RP,error} &= (2\phi_b + 2\phi_d) - (2\phi_a + 2\phi_c) \\ &= 2\phi_Y - 2\phi_X\end{aligned}$$

where

$$\begin{aligned}\phi_X &= \phi_a + \phi_c \\ \phi_Y &= \phi_b + \phi_d\end{aligned}$$

providing us with the information required to stabilise the remaining uncommon optical path lengths. By combining $\hat{\Phi}_{RP,error}$ with the forward path error signal $\hat{\Phi}_{F0}$ and $\hat{\Phi}_{F1}$ we can stabilise the output phase of both emitters, achieving our goal of stable coherent combination:

$$\begin{aligned}\hat{\Phi}_{error,YX} = 2\phi_Y - 2\phi_X = 0 \\ \therefore \phi_Y = \phi_X\end{aligned}$$

The same result applies for additional channels, for example a third channel ϕ_Z:

$$\begin{aligned}\hat{\Phi}_{error,ZX} = 2\phi_Z - 2\phi_X = 0 \\ \therefore \phi_Z = \phi_Y = \phi_X\end{aligned}$$

6.3 Optical implementation

The amplifier-compatible internally sensed OPA was demonstrated experimentally and characterised using the setup shown in Figure 6.4. Three emitters have been stabilised and coherently combined using the amplifier-compatible OPA. We characterise the performance, however, by comparing the relative phase stability of two emitters at a time.

A free-space unstabilised Nd:YAG NPRO laser operating at 1064 nm was used as the master oscillator. Free-space AOMs were placed in the local oscillator and high-power

arms to: 1) generate a heterodyne beat-note at the out-of-loop photodetector for direct measurement of the OPA's output phase stability; 2) allow the OPA to be locked *without* DEHI, providing a useful baseline from which to gauge what kind of effects DEHI has on Φ_{RMS}. Including AOMs in each path also helped prevent parasitic interference caused by the small fraction of zero-order unshifted light that coupled into the optical fibre with the first-order frequency shifted light. Note that ordinarily only a single AOM would be required in the low and high power arms located directly after the second polarising beamsplitter in Figure 6.4.

Commercial 15 Watt fibre amplifiers were used in this experiment. They use steel-braided 10.5/125 μm (core diameter/cladding diameter) large mode area (LMA) fibre to increase the SBS threshold. As our asymmetric fibre couplers (rated by the manufacturer to operate at up to 100 Watts CW optical power) use standard 6/125 μm PM980 fibre, light was coupled from the amplifier's LMA fibre to the coupler's PM980 fibre in free-space, requiring strict mode-matching optics for high-efficiency coupling. At least 90% coupling efficiency was achieved for each free-space link.

The heterodyne beat-note produced by the interference of the two emitters was detected using a commercial free-space 125 MHz bandwidth photodetector (labelled OOL in the figure), placed approximately 30 cm from the optical head. Equivalent fibre-coupled photodetectors were used to detect the interference of the return- and forward paths.

Without DEHI the light from each channel is identified at FP_1, FP_2, and RP via RF demodulation. No cross-talk occurs between the 2 MHz and 4.5 MHz beat-notes present at the RP detector, as the phasemeters apply a band-pass filter with a bandwidth of a $\leq$100 kHz centred at the specific RF frequency.

Φ_{RMS} was measured directly using the same heterodyne detection method used in Section 5.6.2 by shifting the relative frequency of two emitters in the array, and then measuring the phase of the interference beat-note produced at the out-of-loop photodetector using a phasemeter.

6.3.1 Asymmetric fibre couplers

Gooch & Housego (G&H) FPU-064M72A10 fibre couplers with 99.9/0.1 (30 dB) splitting ratio and 100 W CW optical power handling capability were used to isolate the high and low power sections of the optical system. The couplers are fabricated using PM980 fibre and designed to be passively cooled using heat-sinks. The relevant technical specifications for the G&H fibre couplers are provided in Table 6.1.

Table 6.1: Technical specifications for G&H FPU-064M72A10 99.9/0.1 fibre couplers.

Parameter	**Specification**	**Unit**
Max. input power	100	W
Min. splitting ratio	27	dB
Max. splitting ratio	33	dB
Typical splitting ratio	30	dB
Fibre type	PM980	–

The asymmetric fibre couplers were passively cooled using a heat sink, aluminium plates, and the optical bench. Thermal paste was applied to all four of the coupler's interfaces.

6.3.2 Fibre amplifiers

Nufern NUA-1064-PB-0015-B3 fibre amplifiers were used to amplify the optical power of each emitter up to 15 W. The relevant technical specifications for the Nufern amplifiers are provided in Table 6.2.

Table 6.2: Technical specifications for the Nufern 15 Watt NUA-1064-PB-0015-B3 fibre amplifiers.

Parameter	**Specification**	**Unit**
Output power	0.5 to 15.0	W
Input power	1.0 to 15.0	mW
Operating wavelength	1064	nm
Linewidth	≤ 1	kHz
Mode	TEM_{00}	–
Mode of operation	CW	–
Fibre type	Nufern FUD-8015	–
Polarization	Linear	–
Output power stability	≤ 3	%
Polarization extinction ratio	≥ 15	dB

A photograph of the three Nufern fibre amplifiers used in this experiment is shown in Figure 6.5.

6.3.3 Active water cooling of fibre amplifiers

The fibre amplifiers require active water cooling to maintain an operating temperature between 18°C and 23°C. Each amplifier is supplied with an aluminium cooling plate with embedded copper pipes to allow chilled water to be pumped through them. The cooling plates were mounted to the base of the amplifier units using 5/32" counter-sunk screws. No thermal paste was required at the interface between the amplifier and cold plate as existing heat transfer was sufficient.

A ThermoCube 250 W water chiller was used to cool the amplifiers in series. The decision to cool the amplifiers in series simplified the plumbing and was found to adequately cool each amplifier to within their acceptable operating range at full power. The temperature monitors on all three amplifiers indicated that they were each operating at the same temperature to within ±0.5°C.

The inlets to the cold plates are 3/8" outer diameter. The outlet and inlet on the water chiller is 1/4" outer diameter. To connect the cooling plates to the water chiller we converted the 3/8" pipe on the plates to 1/4" using standard brass nut and olive assemblies. 1/4" Teflon tubing was used to interface the water chiller with the cooling plates, requiring plastic olives (in lieu of standard brass olives) to prevent damage to the piping.

Figure 6.5: Photograph of the three Nufern NUA-1064-PB-0015-B3 fibre amplifiers used in the amplifier-compatible OPA.

6.3.4 LMA-to-PM980 free-space mode matching

The amplifier uses steel braided LMA (10.5/125 μm core/cladding diameter) fibre, preventing us from fusion splicing the amplifier directly to the asymmetric fibre couplers which use PM980 fibre. Even if the LMA fibre could be fusion spliced to the PM980 fibre, their different core diameters and associated mode field diameters mean that not all of the light guided through the core of the LMA fibre would couple into the core of the PM980 fibre due to mode mismatch. The fraction of light that is lost instead couples into cladding modes which do not necessarily contribute to total output power as they are strongly attenuated by bend losses, and are not polarization maintaining.

The launch efficiency η from one fibre to another can be estimated using the model

$$\eta = \frac{4\omega_1^2\omega_2^2}{\left(\omega_1^2+\omega_2^2\right)^2}\exp\left(-\frac{2\delta^2}{\omega_1^2+\omega_2^2}\right) \tag{6.7}$$

where ω_1 and ω_2 represent the mode field diameters of the abutting fibres, and δ is the axial offset of the two fibres [109]. For equal core sizes ($\omega_1 = \omega_2$) and zero transverse offset ($\delta = 0$) the efficiency is 100%. The coupling efficiency between LMA fibre ($\omega_1 = 11.4$ μm) and PM980 fibre ($\omega_2 = 6.6$ μm) assuming no transverse offsets is 75.2%. In an ideal system the asymmetric fibre couplers would be fabricated using LMA fibre and fusion spliced directly to the amplifiers to maximise coupling efficiency. Gooch & Housego were contacted about the possibility of fabricating LMA asymmetric fibre couplers and stated that it would be possible since they have already done it for another customer [110].

In this system light was coupled from the LMA fibre amplifiers to the PM980 asymmetric fibre couplers in free space. Appropriate mode-matching lenses were selected by charac-

terising the output modes of the LMA and PM980 fibres, both of which were collimated using Thorlabs 11.17 mm focal length anti-reflection coated fibre collimators.

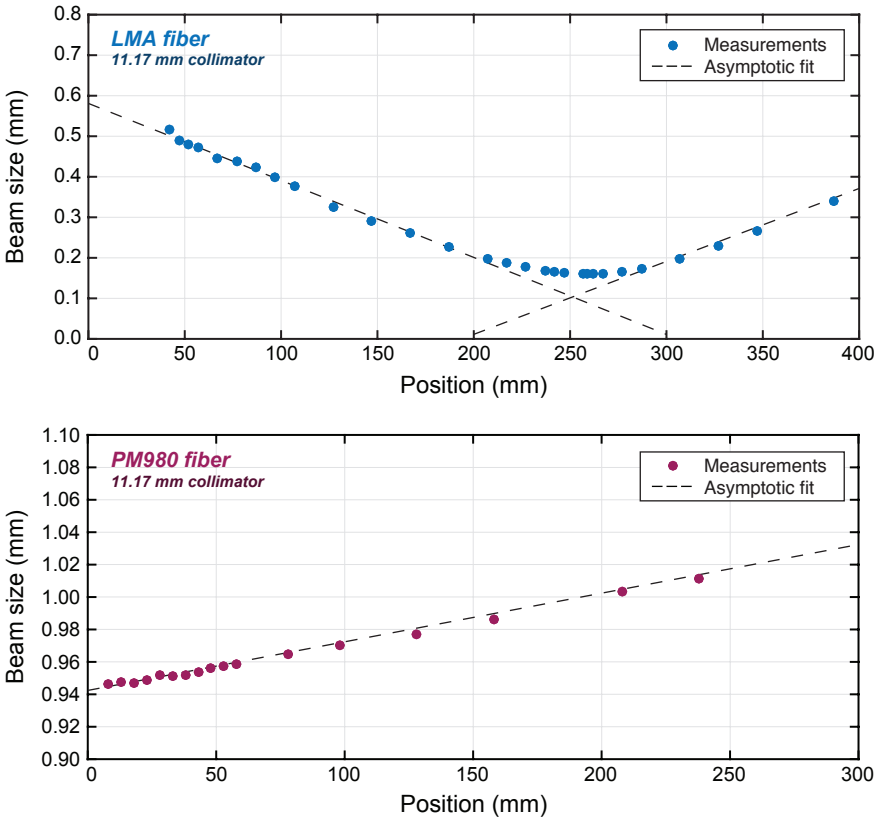

Figure 6.6: Measured beam widths of the collimated output of the LMA fibre (top) and PM980 fibre (bottom). The collimated LMA fibre output has a focal point at approximately 250 mm from the collimator, caused by a longitudinal error of approximately 500 μm between the fibre end-face to the collimating lens.

A Thorlabs BP209-IR/M beam-profiler was used to measure the beam-waist of the collimated outputs of the two fibre types as a function of propagation distance. These measurements are shown in Figure 6.6. Where the beam leaving the PM980 fibre diverges consistently, the beam leaving the LMA fibre appears to converge to a point 250 mm away from the collimator. This indicated that the distance from the LMA fibre end to the collimator was not equal to its specified focal length. This was confirmed using a beam propagation model which suggested that the actual distance between the fibre end and collimator was closer to 11.67 μm, approximately 500 μm longer than the collimator's 11.17 mm focal length. It was discovered that Nufern adds a 500 μm $\pm$ 50 μm coreless fibre end-cap to the ends of their fibre amplifiers to protect them from core damage at high optical power. An illustration of the coreless fibre end-cap is shown in Figure 6.7.

A free-space mode-matching program written in MATLAB was used to calculate various lens configurations that could be used to maximise mode-matching efficiency between the LMA and PM980 fibres. The simplest configuration required two commercially available lenses: one with a 100 mm focal length, the other with 150 mm. The lens arrangement and corresponding beam profile for the selected mode-matching configuration is shown in Figure 6.8. Other configuration were possible, but it was decided that they would either be too sensitive to positioning errors, or would be impractical to implement because they would require too much space on the optical bench ($>$ 1 m).

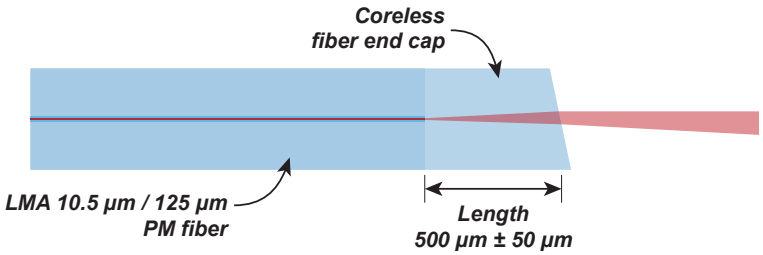

Figure 6.7: Illustration of the 500 μm coreless fibre end cap attached to the output of Nufern's NUA-1064-PB-0015-B3 fibre amplifiers.

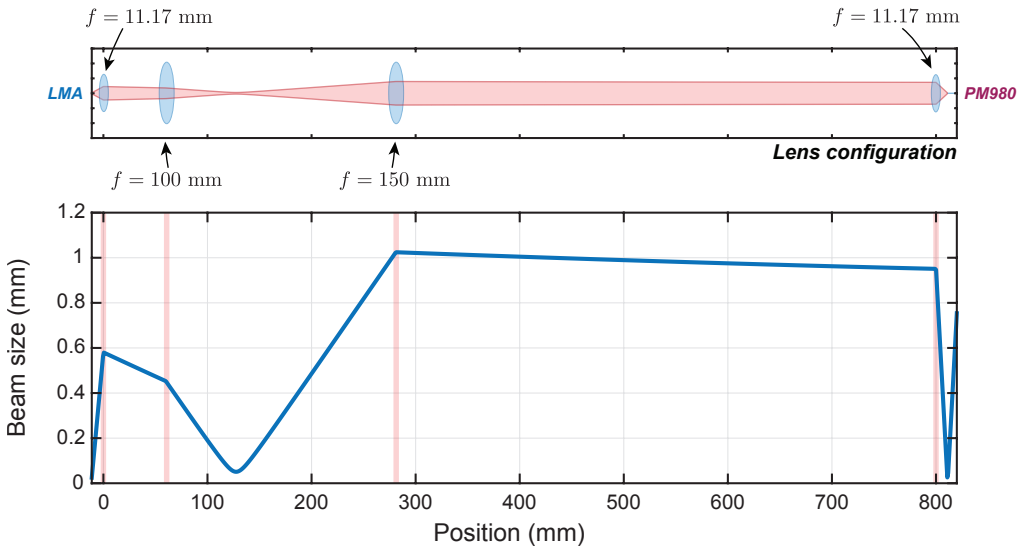

Figure 6.8: (Upper) Lens configuration used to mode-match between LMA and PM980 fibre in free-space. The amplifier's output beam is on the left. (Lower) Beam width as a function of distance.

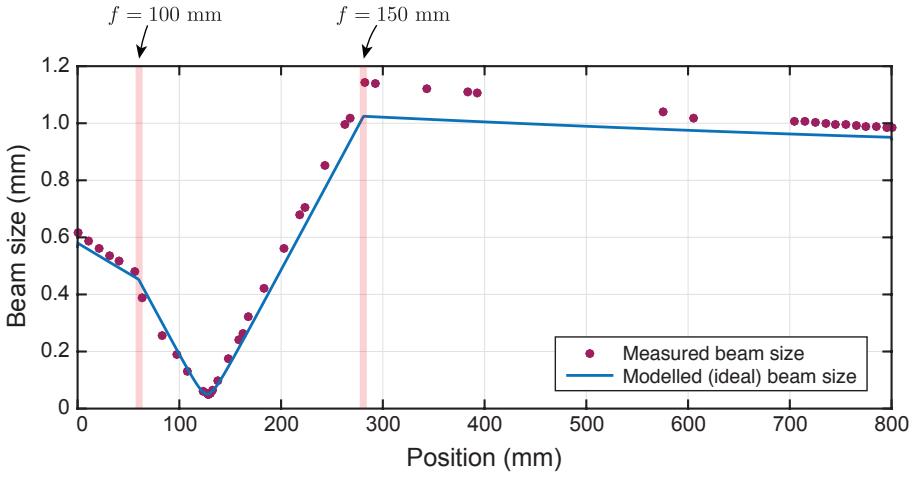

Figure 6.9: Comparison between measured and modelled beam widths.

The 100 mm lens was located 60 mm from the LMA collimator, and the 150 mm lens was positioned 280 mm from the LMA collimator. The beam-profiler was used to measure the beam-waist at different locations during the placement of the lenses to compare with the theoretical predictions of the model. The measured beam size is shown alongside the ideal beam size in Figure 6.9. The slight discrepancies between the model and measured results are caused by slight errors in lens position, and potentially the lens' focal length. There were minor discrepancies in mode-matching for each free-space link, requiring minor adjustments (at most 5 mm) to the positions of the lenses in each path. When all three of the free-space links were installed, the NewPort steering mounts used to hold the collimators were used to maximise power transfer from the amplifier into the asymmetric fibre couplers by maximising the total power at the output of the OPA. All three free-space links achieved at least 90% efficiency, the highest being 95%.

The need for free-space coupling can be eliminated in the future by fusion splicing large mode area asymmetric fibre couplers directly to the amplifiers. Moving to large mode area fibre couplers may also overcome the 100 W per channel limit.

6.4 Digital signal processing

The FPGA implemented measurement and feedback control system used to generate the feedback signals required to stabilise the relative output phase of the OPA is shown in Figure 6.10. 16-bit ADCs are used to digitise the output of each photodetector at 40 MHz. The digitised signals are demodulated with delay matched versions of the PRN codes used to phase modulate the light in each channel of the array. The phase of each demodulated signal is measured using a phasemeter. The phase of the forward path reference emitter (highlighted blue in Figure 6.10) is subtracted from channels 2 and 3, and combined with those produced by the return path phasemeters (highlighted green and magenta) to generate the full phase error signals. The error signals are passed through PI controllers to generate the feedback signals needed to stabilise the relative path-lengths of the three channels. The feedback signals are passed through wrapping algorithms to ensure they do not exceed the operating range of the EOMs (see Table 4.2). The wrapped feedback signals are converted into analogue voltages using 1 MS/s DACs.

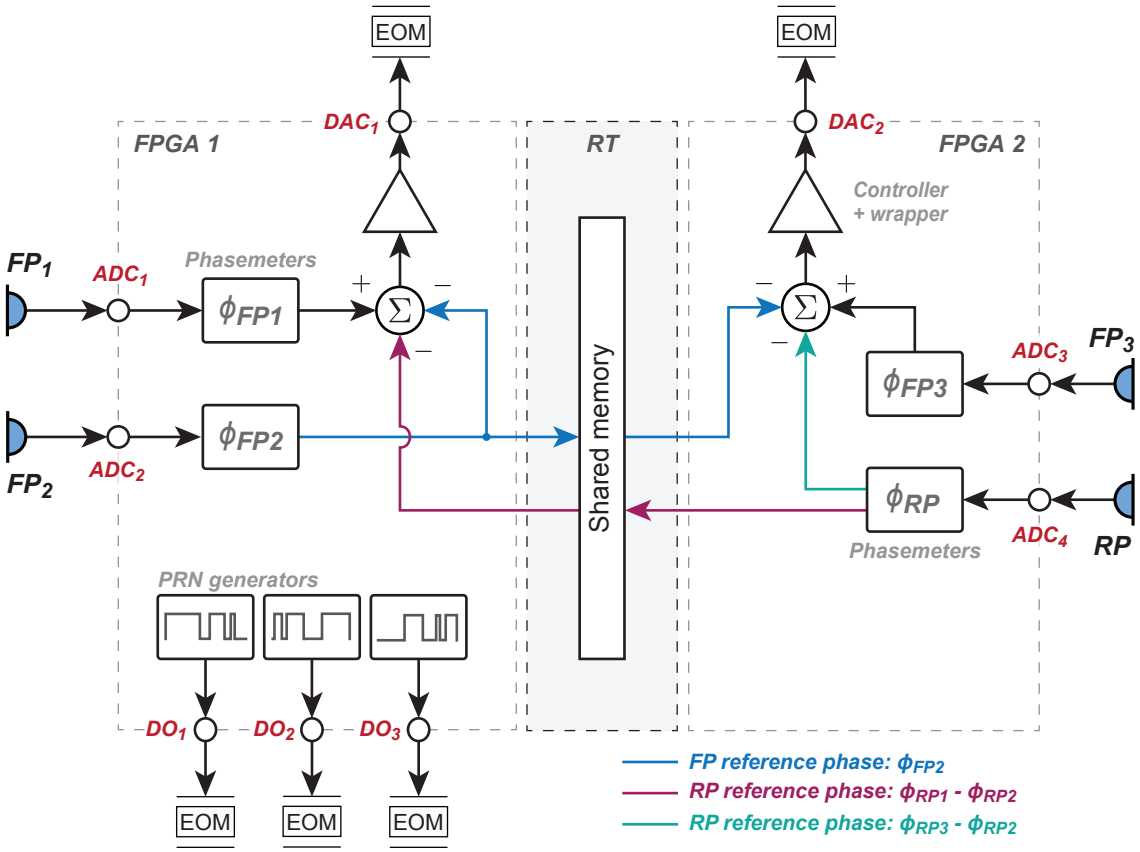

Figure 6.10: Digital signal processing architecture for the amplifier-compatible OPA showing the shared allocation of DSP resources between two NI PXI-7852R FPGAs. The FPGA's clocks are phase-locked to a common 10 MHz reference. Data is transferred from one FPGA to the other using the real-time (RT) computer's shared memory.

6.4.1 Dual FPGAs

Two NI PXI-7852R FPGA platforms were required to implement the amplifier-compatible OPA's measurement and feedback control system. The allocation of DSP resources is shown in Figure 6.10, showing the two PXI-7852R FPGAs (identified as $FPGA_1$ and $FPGA_2$); the real-time computer (identified by RT); and the DSP system's external interfaces (ADCs, DACs, and DOs). Both FPGAs were mounted in the same chassis, which automatically phase locks their clocks to a common 10 MHz back-plane. The PXI chassis includes eight trigger lines that can be configured to synchronise events (e.g., resets) across different platforms. The trigger lines were also used to synchronise the PRN generators on the two FPGAs.

$FPGA_1$ was used to measure the phase of forward path channels one and two (FP_1 and FP_2); derive feedback control signals required to stabilise channel 1 relative to channel 2; and generate PRN codes for DEHI. $FPGA_1$ has two ADC inputs (ADC_1 and ADC_2), one DAC output (DAC_1); and three digital outputs (DO_1, DO_2, and DO_3).

$FPGA_2$ was used to measure the phase of the third forward path channel (FP_3) and the three return path (RP) signals, and to generate feedback to stabilise channel 3 relative to channel 2. $FPGA_2$ has two ADC inputs (ADC_3 and ADC_4), and one DAC output (DAC_2).

Data was transferred between $FPGA_1$ and $FPGA_2$ using the RT's shared memory in the same way as the offset phase-locked OPA (see Section 5.4.1).

6.4.2 Controllers

The amplifier-compatible OPA's controller (shown in Figure 6.11) consists of two PI controllers connected in series to increase gain at low frequencies. The pre-gain K_P reduces the amplitude of the incoming error signal to prevent saturation at the first stage integrator. The proportional arms in both the first and second stage PI controllers have unity gain, where the overall gain of the loop is instead scaled using by K_P and the global gain K_G at the output. The second stage integral gain K_2 is situated before the integrator to prevent saturation caused by integrator wind-up during initial lock acquisition.

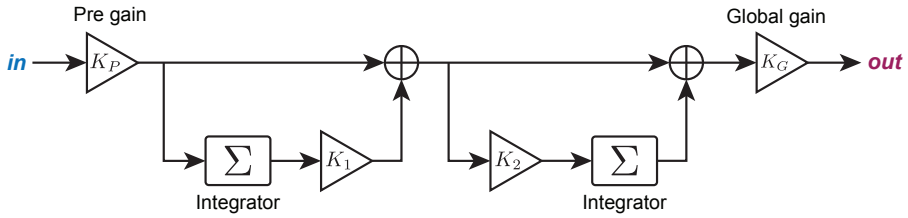

Figure 6.11: Controller architecture used in the amplifier-compatible OPA.

With the exception of K_G, all gains in the controller (K_P, K_1, and K_2) are inverse powers of 2 (e.g., $K_1 = 2^{-3}$) implemented using bit-shift operators. The global gain is implemented using a dedicated multiplier to provide greater flexibility in tuning the controller's bandwidth.

Gains were selected using a transfer function model of the feedback loop, similar to the one described in Section 5.5. Typical gains used in this system were $K_P = 2^{-2}$, $K_1 = 2^{-4}$, $K_2 = 2^{-5}$, and $K_G = 20$. The global gain was adjusted to tune the controller's closed-loop bandwidth. The remaining gains (K_P, K_1, and K_2) were not changed once they were set.

6.5 Linewidth broadening for SBS suppression

The primary motivation behind CBC is to overcome the limitations caused by non-linear effects like stimulated Brillouin scattering. Whilst OPAs overcome these limits by combining the power of multiple lasers, there are obvious advantages in maximizing the amount of power that can be delivered by each element in the array. An approximate expression for the SBS threshold power is [111]

$$P_{th} \approx 21 \frac{bA_e}{g_B L_e} \left(1 + \frac{\Delta\nu_L}{\Delta\nu_B}\right)$$

where b is a number between 1 and 2 depending on the polarization state of the laser, A_e is the effective mode area, g_B is the peak Brillouin gain for the dominant acoustic mode, L_e is the effective transmission length of the fibre, $\Delta\nu_L$ is the line width of the laser, and $\Delta\nu_B$ is the SBS interaction bandwidth. Options to increase P_{th} are to decrease the effective transmission length L_e, increase the effective mode area A_e, and increase the spectral line width of the laser $\Delta\nu_L$. Assuming that L_e can be minimised by design, and that large mode area optical fibre can be used to increase A_e, the remaining strategy is to broaden the linewidth of the laser, which can be achieved by modulating its phase with high-frequency pseudo-random noise to 'spread' the energy of the carrier frequency [61, 112, 113].

This phased array architecture is compatible with linewidth broadening techniques for SBS suppression. To maintain coherent combination, the modulation must be common to all arms of the interferometer. This modulation can be applied using an EOM located directly after the master laser as shown in Fig. 6.1. Any macroscopic path length differences in the optical system will degrade the coherence of the combined beam, the extent of which depends on the root mean square (RMS) path length difference ΔL_{RMS} between channels and the specific auto-correlation properties of the code used. The binary auto-correlation of the maximal-length sequences used in this system is

$$A(\tau) = \begin{cases} 2^N - 1 & \text{for } \tau = 0, L, 2L, \cdots \\ -1 & \text{for any other } \tau \end{cases}$$

where N is the order of the sequence, $L = 2^N - 1$ is the sequence length, and τ is the delay relative to itself. Given a code modulation frequency f_{chip}, the wavelength of a single chip in the PRN code travelling in a medium with refractive index n is $L_{chip} = c/(nf_{chip})$ where c represents the speed of light in vacuum. Coherence degrades approximately linearly as the RMS path length error ΔL_{RMS} approaches $\pm L_{chip}$. Measurements performed by Anderson et al. show that fringe visibility degrades proportional to a sinc function as the optical path length error increases, and re-coheres every code length due to the periodic nature of PRN [112]. Consequently, any path length errors greater than L_{chip} result in incoherent combination. The faster the modulation f_{chip}, the tighter the tolerance on ΔL_{RMS}. The maximum tolerable RMS path length error ΔL_{RMS} in our system is at worst 10% of L_{chip}. For a 1 GHz code this is 2 cm; a 1 GHz PRN modulation can potentially increase the SBS threshold by a factor of ∼10 [113].

6.5.1 Experimental demonstration of SBS suppression using PRN

The experiment shown in Figure 6.12 was used to investigate SBS suppression via linewidth broadening. The light entering the fibre amplifier was phase modulated with a $2^{15} - 1$ chip long PRN code at a chip frequency of 40 MHz, 80 MHz, and 120 MHz. A half-wave

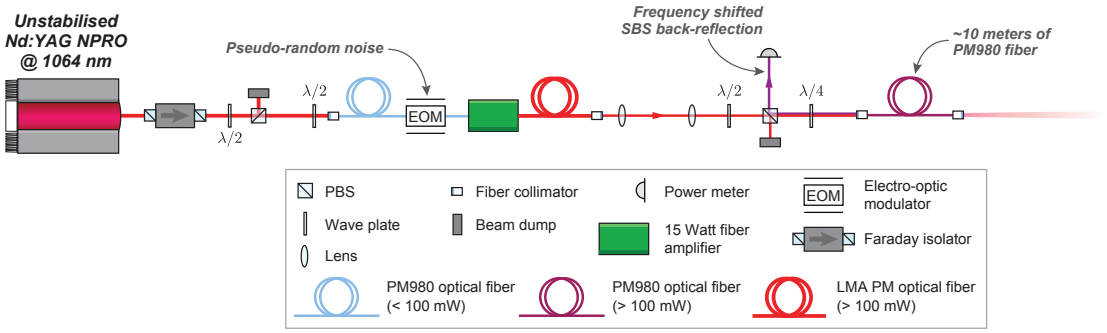

Figure 6.12: Experimental setup used to measure back-reflected power caused by SBS.

plate was used to rotate the polarization of the amplifier's into a horizontal planar state for maximum transmission through a PBS.

Because SBS exhibits the same behaviour as light reflected off a conventional mirror, it's polarization state does not change. The reflection of vertically polarised light remains vertically polarised. This property extends to circularly polarised light as well, where left-hand circularly polarised light induces a right-hand circularly polarised back-reflection (where the change in handedness is purely a convention due to the change in direction of the back-reflection relative to the forward propagating light). This property can be exploited to optically isolate the back-reflected SBS using a Faraday rotator, or alternatively using a PBS and quarter-wave plate as shown in Figure 6.13.

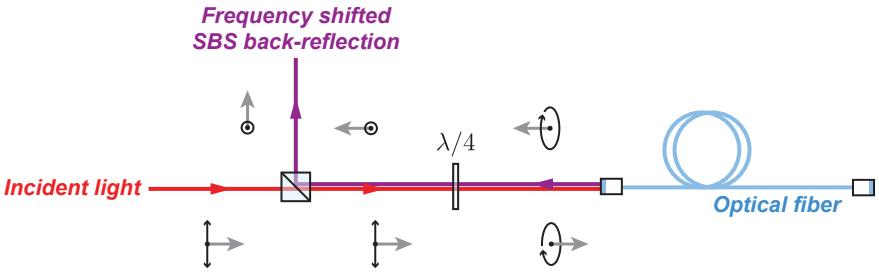

Figure 6.13: Illustration of polarising beam splitter and quarter-wave plate apparatus used to measure back-reflected optical power of SBS.

In this experiment we used the PBS and a single quarter-wave plate to isolate the back-reflection since we did not have access to free-space Faraday isolators. The light exiting the PBS thus passes through a quarter-wave plate to induce circular polarisation. When the SBS back-refection encounters the quarter-wave plate on its way back to the PBS, its polarization is transformed back into a planar state—this time with vertical orientation—and is mostly reflected by the PBS onto a commercial power meter. We were unable to simultaneously monitor the optical power of the back-reflection and output of the 10 meter length of PM980 fibre as we only had access to a single power meter.

The specific isolation of the quarter-wave plate and PBS combination was not measured, allowing an unknown amount of light to leak back into the amplifiers. Whilst this prevented us from being able to accurately measure the absolute back-reflected power, we were still able to measure the effect that linewidth broadening had on it relative to the total back-reflected power without any suppression.

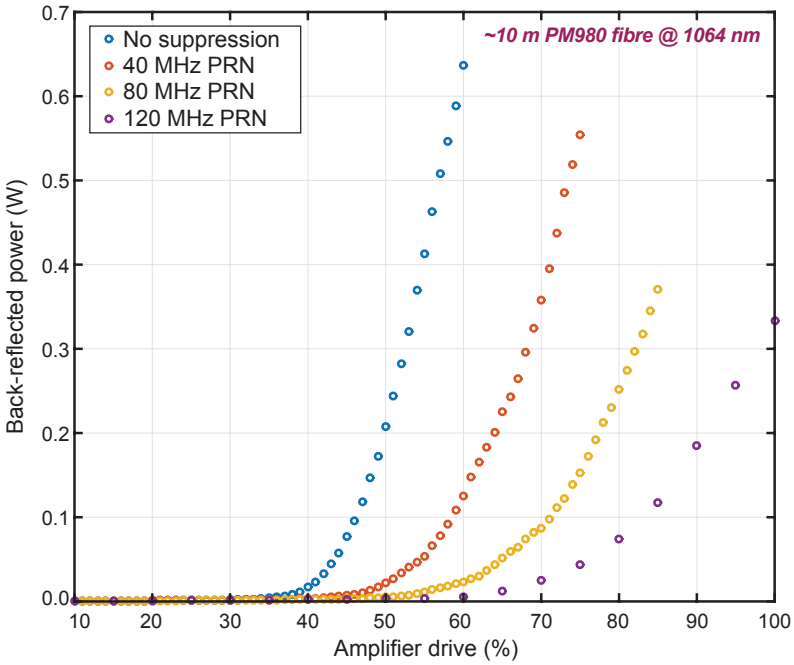

Figure 6.14: Measured back-reflected optical power at different PRN phase modulation frequencies.

Figure 6.14 shows the measured back-reflected power as a function of the amplifier's drive power for various PRN modulation frequencies.

Automatic shut-downs caused by excessive back-reflected power prevented the amplifier from being operated beyond 60% drive power. With 120 MHz PRN modulation the amplifier could be safely operated up to 100% drive power. The SBS threshold with 120 MHz PRN modulation is approximately twice that of no PRN modulation.

Temperature dependence of the amplifier's Faraday isolator

Back-reflections entering the output of the amplifier in the reverse direction are continuously measured at the input of the amplifier's isolator. When the amplifiers were first activated after being idle, it was found that the measured back reflection was initially high, before decreasing gradually over time. This meant that the amplifiers could not be driven to their full capabilities within approximately 30 minutes of activation since they would automatically shut down due to the excessive back-reflections.

While operating at a constant output power, the only parameter observed to be changing over time was the amplifier's internal temperature. The amplifier's temperature at activation was typically 15°C; notably cooler than the temperature of the room (18°C) since the water coolers were operating even when the amplifiers were inactive. After approximately 30 minutes of continuous operation, the amplifier's temperature was found to gradually increase to a steady state level of 19.3°C, and the measured back-reflection stopped decreasing. This suggested a relationship between the isolator's ability to suppress back-reflections and its temperature. This was confirmed by Nufern, who stated that the isolators used in their amplifiers are specified to operate between 20°C and 23°C to achieve typical isolations between 32 and 40 dB. When the isolator is operated below the

specified operating temperature, the Verdet constant of the Faraday rotator[1] increases, strengthening the effect a magnetic field has on the crystal. This diminishes the effectiveness of the isolator which now rotates the polarization of any back-reflected light by more than 90 degrees. This prompted us to increase the set temperature of the water cooling system to ensure that all three amplifiers were operating between 20°C and 23°C. Whilst it is possible to dynamically control the cooling power of the water cooling system, it was simpler in this case to wait until the amplifier's reached steady state temperature before beginning any high-power experiments.

6.6 Results and analysis

6.6.1 Output phase stability

The OOL phase was recorded at a rate of 31.250 kHz to preserve high-frequency harmonics. The time-series and root-power spectral density of the OOL measurements are shown *without* DEHI in Figure 6.15, where the phase of each emitter was isolated using RF demodulation at their respective heterodyne frequencies. The same measurements performed *with* DEHI are shown in Figure 6.16. Three measurements are presented in each figure: 1) when the OPA is unlocked (green); 2) when the forward path controllers are locked (blue); and 3) when both the forward and return path controllers are locked (magenta). The time series of Φ_{RMS} when both controllers are locked has been shifted to appear at zero phase for clarity. None of the time-series data has been de-trended.

While unlocked, the OPA's output phase drifts around due to fibre noise and laser frequency noise. The harmonics visible in the unlocked RPSD's between 10 Hz and 1 kHz originate from acoustic noise in the lab. Much of the acoustic noise is suppressed when the forward path controllers are engaged, which locks the relative phase of the amplified (lower) and local oscillator (upper) paths of the interferometer at their point of recombination at the asymmetric fibre couplers.

When the forward paths are locked the OOL phase error is limited primarily by fibre noise from the uncontrolled lengths of fibre between the asymmetric couplers and the optical head. The return path controller stabilises these uncontrolled lengths of fibre (magenta).

The RPSD of the fully stabilised output phase error shown in Figure 6.15(b) reveals a white noise floor of 10 μcycles/$\sqrt{\text{Hz}}$ at frequencies between 10 Hz and 10 kHz. Whilst this noise-floor is higher than that of the forward path locked RPSD (blue), the RPSD demonstrates the RP controller's suppression of low-frequency fibre noise (at frequencies lower than 10 Hz). The 10 μcycles/$\sqrt{\text{Hz}}$ noise-floor of the fully locked RPSD is due to the low SNR of the RP signal, which limits the phasemeter's precision.

The RMS phase error for each measurement was calculated over a bandwidth of 15.625 kHz for a measurement period of 15 minutes. The RMS phase error without DEHI was calculated to be $\lambda/206$, limited at longer time scales by a low-frequency shelf at frequencies lower than 10 Hz. The RMS phase error with DEHI (with a $2^{24}-1$ bit long PRN code) is $\lambda/194$, limited primarily by harmonic distortion introduced by the imperfect demodulation of PRN in the controller. The harmonic distortion present in the DEHI measurements consists of peaks separated by 2.38 Hz, equal to the code-repetition rate of the $2^{24}-1$ element long PRN code: 40 MHz / $(2^{24}-1) = 2.38$ Hz. The amplifier compatible OPA's RMS phase error is lower than that the offset phase-locked OPA due to the increased bandwidth of the controller.

[1] Nufern's 15 Watt 1064 nm fibre amplifiers use terbium gallium garnet isolators.

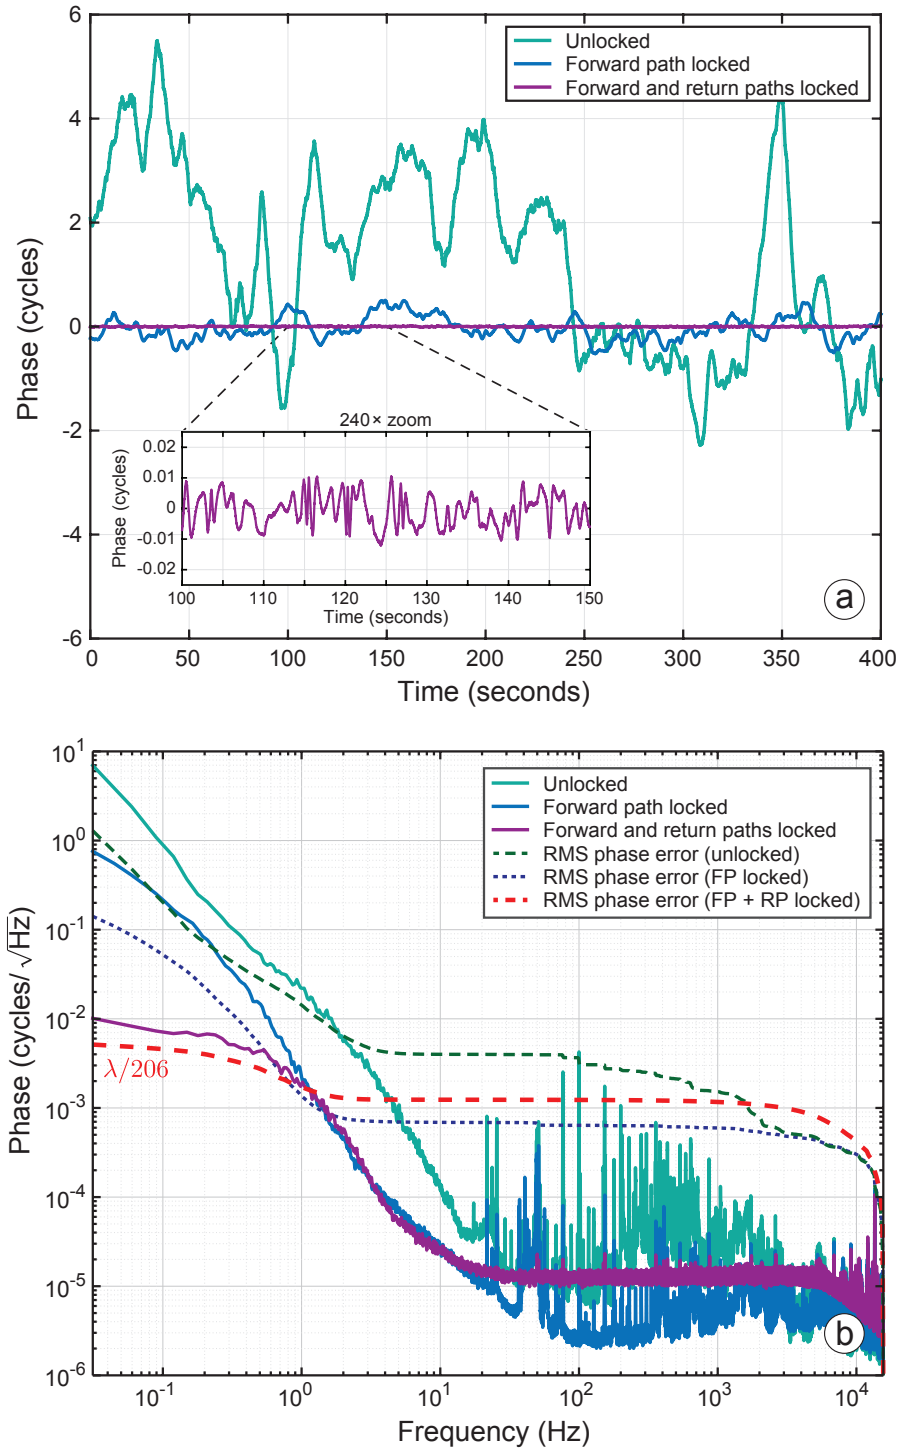

Figure 6.15: (a) Time series measurements of Φ_{RMS} when the OPA is unlocked (green), forward path locked (blue), and fully locked (magenta) without DEHI. The inset shows the zoomed-in behaviour of Φ_{RMS}. (b) RPSD of the measurements shown in (a); the noise-shelf at frequencies below 10 Hz is typical of cyclic phase noise introduced by parasitic interference somewhere in the optical system.

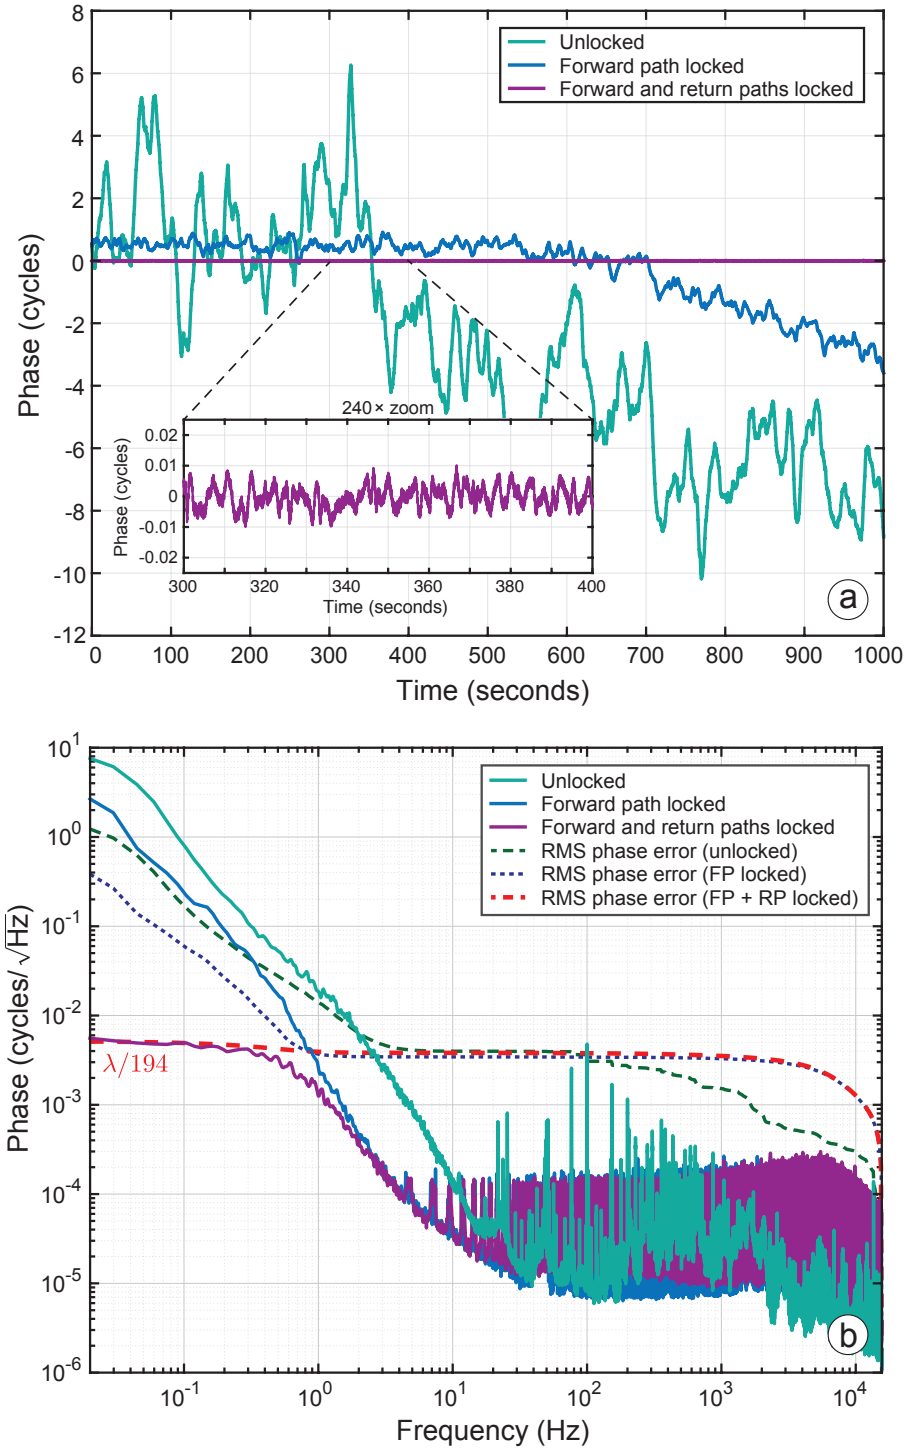

Figure 6.16: (a) Time series measurements of Φ_{RMS} when the OPA is unlocked (green), forward path locked (blue), and fully locked (magenta) *with* DEHI. The inset shows the zoomed-in behaviour of Φ_{RMS}. (b) RPSD of the measurements shown in (a). The harmonic distortion visible on the forward and return path locked RPSD in (b) is caused by residual PRN noise introduced by the demodulator.

Cyclic phase noise

The noise-shelf present below 10 Hz in both the fully locked traces is typical of cyclic phase noise, suggesting parasitic interference somewhere in the optical system. Because the noise is only revealed when both the forward and return path controllers are locked, it was difficult to determine whether or not the parasitic interference was occurring at the out of loop photodetector, or alternatively at one of the forward or return path photodetectors.

Whilst parasitic interference is typically caused by partial reflections, it can also be caused by interference of light in an orthogonal polarisation. Other possible causes include the interference between cladding modes introduced by non-perfect coupling into optical fibre, which experience slightly different optical path lengths—and thus phase—to modes guided in the core; harmonic aliasing, caused by the discrete-time nature of the digital signal processing system; and RF coupling between AOMs.

Quarter-wave plates were used to compensate for any polarisation ellipticity caused by the laser or upstream optics, and a polariser was placed in front of the OOL photodetector to reject orthogonal polarizations. Despite improving stability over very long time scales (>10,000 seconds), it did not significantly reduce the magnitude of the cyclic phase noise. The cyclic phase noise was, however, observed to be consistently lower when using DEHI, suggesting partial suppression of the parasitic interference by DEHI's range-gate, and supporting the idea that the cyclic phase noise was being caused by a partial reflection.

The potential interference between cladding modes was managed by inducing strong bend losses in each length of optical fibre by tightly wrapping short segments around one inch cylindrical optics posts. This did not result in any noticeable change in the magnitude of the cyclic noise.

Aliased harmonics were suppressed by placing 100 MHz cut-off frequency low-pass filters at the output of each photodetector. It was anticipated that this would affect the demodulation of the PRN modulated signals, however it had no noticeable effect on phasemeter performance. Including LPFs at each photodetector did not affect the cyclic phase noise.

The locations in the optical system that exacerbate the cyclic phase noise were anywhere between the: asymmetric fibre coupler and the optical head; asymmetric fibre coupler and return path PRN modulating EOMs; and the return path EOMs and recombining 3-dB fibre couplers. These observations indicate a partial reflection somewhere between the asymmetric fibre coupler and optical head, which would experience a different optical path length on its path back towards the return path photodetectors. If this is the case, then a potential next step would be to force the cyclic phase noise into a higher frequency using a PZT fibre stretcher (or any device capable of modulating the length of fibre, e.g., heating element) to reveal the underlying noise that is currently hidden by the parasitic interference [114].

Whilst the noise-shelf caused by cyclic phase noise is the highest magnitude feature in the RPSD, the results show that the it does not contribute that significantly to Φ_{RMS}. In the DEHI measurements, Φ_{RMS} is instead limited by residual phase noise introduced by imperfect demodulation of PRN.

6.6.2 Beam-steering

Dynamic beam-steering was demonstrated experimentally by adding a swept sine-wave (from 1 Hz up to 10 kHz) with an amplitude of 0.25 cycles to both the forward- and return path phase measurements (as described in Section 5.6.4). This is equivalent to injecting

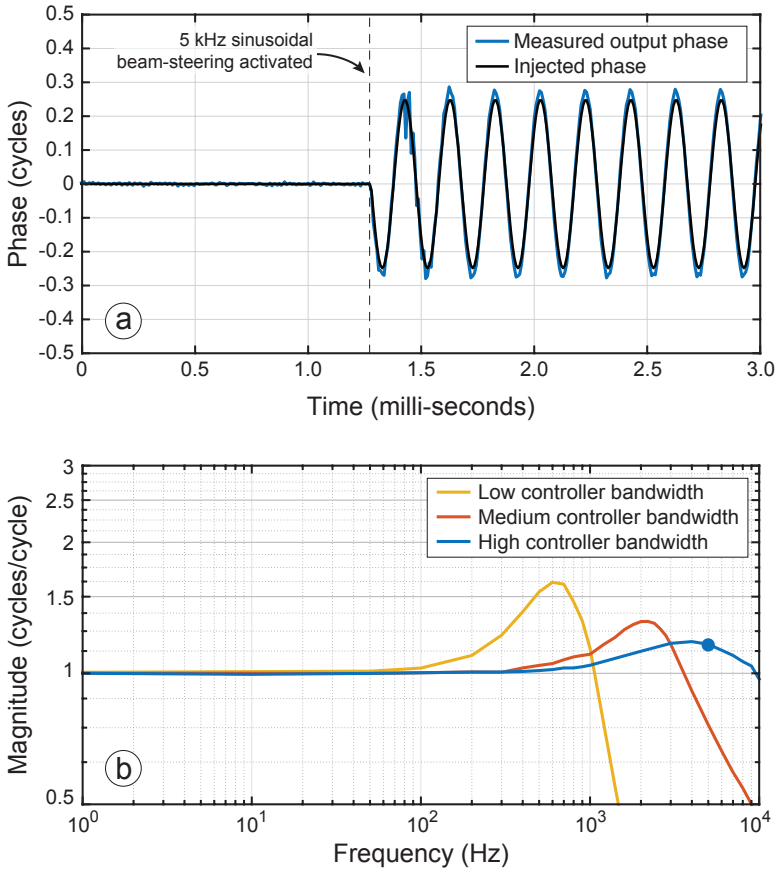

Figure 6.17: (a) Time series measurement of 5 kHz sinusoidal beam-steering with a high controller bandwidth. (b) Measured magnitude response of beam-steering for different controller bandwidths. The 5 kHz tone shown in (a) is identified by the blue circle in (b).

a signal into the error point of the control loop; the control system then imposes this signal on the output. By comparing the injected signal to the measured phase at the OOL detector, we were able to calculate the beam-steering forward-loop magnitude response. The OOL phase was measured using the setup shown in Figure 6.4. The bandwidth of the phasemeter used to measure the phase of the OOL signal was set at 100 kHz.

Figure 6.17 shows a time series of the output and input phase measurements (upper panel), as well as the calculated magnitude response of the beam-steering for different controller bandwidths (lower panel). The results reveal the strong influence the controller bandwidth has over the range of frequencies the beam can be steered: higher controller bandwidths enable the controller to 'correct' higher frequency disturbances. At high controller bandwidths the beam can be steered by up to 10 kHz. The peaks clearly visible in all three magnitude response traces is caused by gain-peaking in the control loops.

The beam-steering bandwidth can in principle be extended by simultaneously injecting phase disturbances at frequencies higher than the controller bandwidth at the *output* of the controller, with appropriate blend filtering between the signals injected at the controller's input and output. These higher-frequency signals will not be suppressed by the controller as they exceed its bandwidth.

Whilst this demonstration of beam-steering was performed with two emitters, it is simple to extend the capability to an array of emitters with any topology. McManamon et al. provide a detailed explanation of this in [97].

6.7 Summary

The amplifier-compatible OPA presented in this chapter successfully demonstrates internal sensing's compatibility with high power fibre amplifiers. The proposed amplifier-compatible OPA can be scaled to a large number of emitters, each capable of delivering up to 100 W of continuous wave optical power limited by the damage threshold of the commercial asymmetric fibre couplers. An experimental demonstration of the amplifier-compatible OPA validated the concept, which was shown to stabilise the relative output phase of the array with a state-of-the-art RMS output phase error of $\lambda/194$ using digitally enhanced heterodyne interferometry, significantly better than the $\lambda/24$ requirement. As the phase actuation is provided by high-bandwidth waveguide-based electro-optic modulators, it is possible to dynamically manipulate the distribution of optical power up to the bandwidth of the control system, enabling agile beam forming and steering. This system is also compatible with pseudo-random noise based linewidth broadening techniques for the suppression of stimulated Brillouin scattering.

Conclusions and future work

7.1 Summary

This thesis presented the iterative design and experimental implementation of an internally sensed optical phased array for ground-based space debris laser ranging and manoeuvring systems.

One of the greatest challenges with manoeuvring orbital space debris via radiation pressure is object illumination, requiring large telescopes equipped with extremely high-power lasers and adaptive optics capabilities. Whilst it may soon be possible to demonstrate remote manoeuvring using commercial 15 kW fiber lasers, it is important to explore other ways to scale optical power even further, without sacrificing the high beam quality required for near diffraction limited propagation through the atmosphere. We have presented a technique that aims to do this by coherently combining multiple high-power sources using an internally sensed OPA.

The internal sensing concept was first validated using a proof-of-concept experiment to coherently combine three low power emitters. This experiment demonstrated the effective use of digitally enhanced heterodyne interferometry and high-precision phasemeters to isolate and measure the output phase of individual emitters at a single photodetector. Successful as it was, the experiment highlighted two crucial limitations that would prevent it from being effective in its proposed application. The first was its incompatibility with in-line fiber amplifiers, which effectively limited the power handling capabilities of each emitter to the damage threshold of sensitive optical elements. The second was the need to implement DEHI using low modulation depth pseudo-random noise codes to preserve far-field coherence, which as a consequence severely limited the control system's bandwidth.

These limitations were overcome by including asymmetric fiber couplers in each emitter, allowing high-intensity sources to be coupled in to the output of the OPA without exposing low damage threshold components to high optical power. This modification to the architecture also permitted the use of full π modulation depth PRN codes, which improved the SNR of the decoded signals and consequently the control system's bandwidth.

This high-power compatible internally sensed optical phased array was first demonstrated experimentally by coherently combining three slave lasers. This involved offset phase-locking each slave to a single master laser at the asymmetric fiber couplers by interfering them at separate forward path photodetectors. The remaining uncontrolled lengths of optical fiber were stabilised by measuring the phase of each emitter's back-reflection at a single return-path photodetector.

The ability to arbitrarily set the relative frequencies of the slave lasers allowed us to measure the offset phase-locked OPA's RMS output phase stability using heterodyne de-

tection. Whilst this measurement introduced a frequency difference between emitters, it did so whilst maintaining the OPA's output phase coherence. These measurements revealed excellent RMS output phase stability at $\Phi_{\mathrm{RMS}} = \lambda/104$, well below the $\lambda/24$ design requirement. An analysis of the measurement's root power spectral density showed that Φ_{RMS} was limited by residual PRN noise introduced by imperfect demodulation. Beam-steering at up to 10 kHz was demonstrated by injecting a swept sine-wave into both the forward- and return-path phase measurements.

A more efficient way to scale the optical power handling capabilities of the OPA is to use in-line fiber amplifiers. The internally sensed OPA's compatibility with fiber amplifiers was demonstrated by coherently combining three commercial 15 W fiber amplifiers, drawing on a combination of techniques and ideas introduced in the proof-of-concept and offset phase-locked OPA experiments. Implementing the same measurement scheme used to characterise the offset-phase locked OPA's performance, the amplifier-compatible internally sensed OPA was shown to stabilise the relative output phase of the array with an unprecedented RMS output phase error of $\Phi_{\mathrm{RMS}} = \lambda/194$, limited by residual PRN noise. This system was also used to perform beam-steering up to 10 kHz.

The final amplifier-compatible OPA design presented in this thesis is potentially scalable to a large number of emitters, each capable of supporting optical powers up to damage threshold of the commercial asymmetric fiber couplers. The system is also compatible with PRN based linewidth broadening techniques for the suppression of stimulated Brillouin scattering, potentially allowing it to support higher optical powers than would otherwise be possible.

7.2 Future work

The internally sensed OPA presented in this thesis is still in the early stages of development, and will require a number of improvements before it can be implemented in future ground-based space debris laser ranging and manoeuvring systems. These improvements include: maximising the power handling capabilities of each emitter; scaling the OPA to tens or hundreds of emitters; solving the π phase ambiguity caused by double-passing the same length of optical fiber; and demonstrating the ability to compensate for atmospheric turbulence.

The power handling capabilities of each emitter is currently limited by the damage threshold of the asymmetric fiber couplers, which in this implementation is 100 W. This threshold can in principle be extended to hundreds of Watts by fabricating the asymmetric fiber couplers using large mode area optical fiber instead of standard PM980 fiber [110]. This would also allow the amplifiers to be fusion spliced directly to the asymmetric fiber couplers, which would maximise coupling efficiency, and forgo the need for free-space mode-matching. A large mode area fiber coupler would also require the optical head to be assembled using the same fiber. One of the complications with this modification, however, would be interfacing it with waveguide-based electro-optic modulators which are typically fabricated using standard PM980 fiber.

The individual power handling capacity of each emitter may also be extended using linewidth broadening techniques to suppress stimulated Brillouin scattering. This could be demonstrated experimentally by modulating the phase of the master laser with pseudo-random noise as shown in Figure 6.1, requiring all optical path-lengths in the system to be matched in order to preserve far-field coherence.

The internally sensed OPA's ability to scale optical power requires it to support potentially hundreds of emitters. This depends largely on digitally enhanced heterodyne interferometry's multiplexing capabilities, which until this point has not been demonstrated for more than four interferometric signals [100]. Significant work is thus needed in order to establish DEHI's limitations in this regard.

Scaling to large numbers of emitters will also require one or multiple one-to-many splitters, which may be fabricated using three-dimensional pulsed laser waveguide inscription (e.g., [115, 116]).

A consequence of the internally sensed architecture is an ambiguity in phase due to the light's double pass through fiber. This π phase ambiguity may be solved by performing an additional coarse phase measurement through a different return-path. If the two emitters are out of phase, then they will interfere destructively. If they are in-phase then they will interfere constructively. The *true* relative phase of the two emitters can therefore be inferred by measuring the intensity of their interference.

With these improvements, the amplifier-compatible OPA's scalability, high-power handling characteristics, excellent output phase coherence, and high-bandwidth beam-steering capabilities will make it a promising technology for future ground-based space debris laser ranging and manoeuvring systems.

References

[1] Satellite Industry Association, "State of the satellite industry report," Report, 2015. [Online]. Available: http://www.sia.org/wp-content/uploads/2015/06/Mktg15-SSIR-2015-FINAL-Compressed.pdf

[2] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: the creation of a debris belt," *Journal of Geophysical Research: Space Physics*, vol. 83, no. A6, pp. 2637–2646, 1978.

[3] P. Anz-Meador and D. Shoots, "Monthly number of objects in earth orbit by object type," *Orbital Debris Quarterly News*, vol. 20, no. 1 - 2, p. 14, 2016.

[4] C. Stokely and M. Matney, "Haystack radar observations of debris from the Fengyun-1C antisatellite test," *Orbital Debris Quarterly News*, vol. 12, no. 3, pp. 7–8, 2008.

[5] J. Liou and D. Shoots, "Satellite collision leaves significant debris clouds," *Orbital Debris Quarterly News*, vol. 13, no. 2, pp. 1–2, 2009.

[6] M. Matney, "Small debris observations from the iridium 33/cosmos-2251 collision," *Orbital Debris Quarterly News*, vol. 14, no. 2, pp. 6–8, 2010.

[7] P. Anz-Meador and D. Shoots, "International space station performs fourth and fifth debris avoidance maneuvers of 2014," *Orbital Debris Quarterly News*, vol. 19, no. 1, pp. 1–2, 2015.

[8] National Aeronautics Space Administration (NASA) Orbital Debris Program Office, "Orbital debris frequently asked questions," Website, 2012. [Online]. Available: http://www.orbitaldebris.jsc.nasa.gov/faqs.html#4

[9] C. H. Smith, "The EOS space debris tracking system," *Proceedings of 2006 AMOS Technical Conference, Kihei, HI*, pp. 719–728, 2006.

[10] J. Liou and S. Portman, "Chinese anti-satellite test creates most severe orbital debris cloud in history," *Orbital Debris Quarterly News*, vol. 11, no. 2, pp. 2–3, 2007.

[11] B. Greene, Y. Gao, and C. Moore, "Laser tracking of space debris," *13th International Workshop on Laser Ranging Instrumentation, Washington DC*, 2002.

[12] J. Liou and S. Portman, "United nations adopts space debris mitigation guidelines," *Orbital Debris Quarterly News*, vol. 11, no. 2, pp. 1–2, 2007.

[13] J. Mason, J. Stupl, W. Marshall, and C. Levit, "Orbital debris–debris collision avoidance," *Advances in Space Research*, vol. 48, no. 10, pp. 1643–1655, 2011.

[14] C. R. Phipps, "A laser-optical system to re-enter or lower low earth orbit space debris," *Acta Astronautica*, vol. 93, pp. 418–429, 2014.

[15] J. Stupl, N. Faber, C. Foster, F. Y. Yang, B. Nelson, J. Aziz, A. Nuttall, C. Henze, and C. Levit, "Lightforce photon-pressure collision avoidance: Updated efficiency analysis utilizing a highly parallel simulation approach," DTIC Document, Report, 2014.

[16] L. DeLuca, F. Bernelli, F. Maggi, P. Tadini, C. Pardini, L. Anselmo, M. Grassi, D. Pavarin, A. Francesconi, and F. Branz, "Active space debris removal by a hybrid propulsion module," *Acta Astronautica*, vol. 91, pp. 20–33, 2013.

[17] D. Richardson, J. Nilsson, and W. Clarkson, "High power fiber lasers: current status and future perspectives [invited]," *JOSA B*, vol. 27, no. 11, pp. B63–B92, 2010.

[18] IPG Photonics Corporation, "YLS-100000 High Power Ytterbium Fiber Laser," Technical Specification, 22 June 2016 2014. [Online]. Available: http://new.ipgphotonics.com/47/FileAttachment/YLS-100000+Datasheet.pdf

[19] S. J. McNaught, C. P. Asman, H. Injeyan, A. Jankevics, A. M. Johnson, G. C. Jones, H. Komine, J. Machan, J. Marmo, and M. McClellan, "100-kW coherently combined Nd: YAG MOPA laser array," in *Frontiers in Optics*. Optical Society of America, Conference Proceedings, p. FThD2.

[20] H. Injeyan and G. Goodno, *High Power Laser Handbook*. McGraw-Hill Education, 2011.

[21] V. Daneu, A. Sanchez, T. Fan, H. Choi, G. Turner, and C. Cook, "Spectral beam combining of a broad-stripe diode laser array in an external cavity," *Optics Letters*, vol. 25, no. 6, pp. 405–407, 2000.

[22] E. J. Bochove, "Theory of spectral beam combining of fiber lasers," *IEEE journal of quantum electronics*, vol. 38, no. 5, pp. 432–445, 2002.

[23] S. J. Augst, A. Goyal, R. Aggarwal, T. Fan, and A. Sanchez, "Wavelength beam combining of ytterbium fiber lasers," *Optics Letters*, vol. 28, no. 5, pp. 331–333, 2003.

[24] C. E. Hamilton, S. C. Tidwell, D. Meekhof, J. Seamans, N. Gitkind, and D. D. Lowenthal, "High-power laser source with spectrally beam-combined diode laser bars," *Lasers and Applications in Science and Engineering*, pp. 1–10, 2004.

[25] B. Chann, R. Huang, L. Missaggia, C. Harris, Z. Liau, A. Goyal, J. Donnelly, T. Fan, A. Sanchez-Rubio, and G. Turner, "Near-diffraction-limited diode laser arrays by wavelength beam combining," *Optics Letters*, vol. 30, no. 16, pp. 2104–2106, 2005.

[26] P. Sprangle, A. Ting, J. Penano, R. Fischer, and B. Hafizi, "Incoherent combining and atmospheric propagation of high-power fiber lasers for directed-energy applications," *IEEE Journal of quantum electronics*, vol. 45, no. 2, pp. 138–148, 2009.

[27] C. Wirth, O. Schmidt, I. Tsybin, T. Schreiber, R. Eberhardt, J. Limpert, A. Tünnermann, K. Ludewigt, M. Gowin, and E. Ten Have, "High average power spectral beam combining of four fiber amplifiers to 8.2 kW," *Optics Letters*, vol. 36, no. 16, pp. 3118–3120, 2011.

[28] S. Augst, S. Redmond, C. Yu, D. Ripin, T. Fan, G. Goodno, P. Thielen, J. Rothenberg, and A. Sanchez-Rubio, "Coherent and spectral beam combining of fiber lasers," in *SPIE LASE*. International Society for Optics and Photonics, Conference Proceedings, pp. 823 704–823 704–10.

[29] Y. Xiao, F. Brunet, M. Kanskar, M. Faucher, A. Wetter, and N. Holehouse, "1-kilowatt cw all-fiber laser oscillator pumped with wavelength-beam-combined diode stacks," *Optics Express*, vol. 20, no. 3, pp. 3296–3301, 2012.

[30] E. Honea, R. S. Afzal, M. Savage-Leuchs, N. Gitkind, R. Humphreys, J. Henrie, K. Brar, and D. Jander, "Spectrally beam combined fiber lasers for high power, efficiency, and brightness," *SPIE LASE*, pp. 860 115–860 115–5, 2013.

[31] J. R. Leger, "External methods of phase locking and coherent beam addition of diode lasers," *Surface Emitting Semiconductor Lasers and Arrays*, pp. 379–433, 1993.

[32] W. M. Neubert, K. H. Kudielka, W. R. Leeb, and A. L. Scholtz, "Experimental demonstration of an optical phased array antenna for laser space communications," *Applied Optics*, vol. 33, no. 18, pp. 3820–3830, 1994.

[33] J. Anderegg, S. J. Brosnan, M. E. Weber, H. Komine, and M. G. Wickham, "8-w coherently phased 4-element fiber array," in *High-Power Lasers and Applications*. International Society for Optics and Photonics, Conference Proceedings, pp. 1–6.

[34] S. J. Augst, T. Fan, and A. Sanchez, "Coherent beam combining and phase noise measurements of ytterbium fiber amplifiers," *Optics Letters*, vol. 29, no. 5, pp. 474–476, 2004.

[35] H. Bruesselbach, D. C. Jones, M. S. Mangir, M. Minden, and J. L. Rogers, "Self-organized coherence in fiber laser arrays," *Optics Letters*, vol. 30, no. 11, pp. 1339–1341, 2005.

[36] H. Bruesselbach, S. Wang, M. Minden, D. C. Jones, and M. Mangir, "Power-scalable phase-compensating fiber-array transceiver for laser communications through the atmosphere," *JOSA B*, vol. 22, no. 2, pp. 347–353, 2005.

[37] T. Y. Fan, "Laser beam combining for high-power, high-radiance sources," *IEEE Journal of selected topics in Quantum Electronics*, vol. 11, no. 3, pp. 567–577, 2005.

[38] J. Anderegg, S. Brosnan, E. Cheung, P. Epp, D. Hammons, H. Komine, M. Weber, and M. Wickham, "Coherently coupled high-power fiber arrays," in *Lasers and Applications in Science and Engineering*. International Society for Optics and Photonics, Conference Proceedings, pp. 61 020U–61 020U–5.

[39] G. Goodno, H. Komine, S. McNaught, S. Weiss, S. Redmond, W. Long, R. Simpson, E. Cheung, D. Howland, and P. Epp, "Coherent combination of high-power, zigzag slab lasers," *Optics Letters*, vol. 31, no. 9, pp. 1247–1249, 2006.

[40] B. He, Q. Lou, J. Zhou, J. Dong, Y. Wei, D. Xue, Y. Qi, Z. Su, L. Li, and F. Zhang, "High power coherent beam combination from two fiber lasers," *Optics Express*, vol. 14, no. 7, pp. 2721–2726, 2006.

[41] T. Shay, V. Benham, J. Baker, B. Ward, A. D. Sanchez, M. A. Culpepper, D. Pilkington, J. Spring, D. J. Nelson, and C. A. Lu, "First experimental demonstration of self-synchronous phase locking of an optical array," *Optics Express*, vol. 14, no. 25, pp. 12 015–12 021, 2006.

[42] T. M. Shay, "Theory of electronically phased coherent beam combination without a reference beam," *Optics Express*, vol. 14, no. 25, pp. 12 188–12 195, 2006.

[43] C. Yu, J. Kansky, S. Shaw, D. Murphy, and C. Higgs, "Coherent beam combining of large number of PM fibres in 2-D fibre array," *Electronics Letters*, vol. 42, no. 18, p. 1, 2006.

[44] G. D. Goodno, C. P. Asman, J. Anderegg, S. Brosnan, E. C. Cheung, D. Hammons, H. Injeyan, H. Komine, W. H. Long, and M. McClellan, "Brightness-scaling potential of actively phase-locked solid-state laser arrays," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 3, no. 13, pp. 460–472, 2007.

[45] E. C. Cheung, J. G. Ho, G. D. Goodno, R. R. Rice, J. Rothenberg, P. Thielen, M. Weber, and M. Wickham, "Diffractive-optics-based beam combination of a phase-locked fiber laser array," *Optics Letters*, vol. 33, no. 4, pp. 354–356, 2008.

[46] R. Xiao, J. Hou, M. Liu, and Z. Jiang, "Coherent combining technology of master oscillator power amplifier fiber arrays," *Optics Express*, vol. 16, no. 3, pp. 2015–2022, 2008.

[47] N. Satyan, W. Liang, A. Kewitsch, G. Rakuljic, and A. Yariv, "Coherent power combination of semiconductor lasers using optical phase-lock loops," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 15, no. 2, pp. 240–247, 2009.

[48] T. Shay, J. Baker, A. D. Sanchez, C. Robin, C. Vergien, C. Zeringue, D. Gallant, C. A. Lu, B. Pulford, and T. Bronder, "High-power phase locking of a fiber amplifier array," *SPIE LASE: Lasers and Applications in Science and Engineering*, pp. 71 951M–71 951M–8, 2009.

[49] B. Wang, E. Mies, M. Minden, and A. Sanchez, "All-fiber 50 W coherently combined passive laser array," *Optics Letters*, vol. 34, no. 7, pp. 863–865, 2009.

[50] G. D. Goodno, S. J. McNaught, J. E. Rothenberg, T. S. McComb, P. A. Thielen, M. G. Wickham, and M. E. Weber, "Active phase and polarization locking of a 1.4 kW fiber amplifier," *Optics Letters*, vol. 35, no. 10, pp. 1542–1544, 2010.

[51] R. Uberna, A. Bratcher, and B. G. Tiemann, "Coherent polarization beam combination," *IEEE Journal of Quantum Electronics*, vol. 46, no. 8, pp. 1191–1196, 2010.

[52] M. A. Vorontsov, S. L. Lachinova, L. A. Beresnev, and T. Weyrauch, "Obscuration-free pupil-plane phase locking of a coherent array of fiber collimators," *JOSA A*, vol. 27, no. 11, pp. A106–A121, 2010.

[53] Y. Ma, X. Wang, J. Leng, H. Xiao, X. Dong, J. Zhu, W. Du, P. Zhou, X. Xu, and L. Si, "Coherent beam combination of 1.08 kW fiber amplifier array using single frequency dithering technique," *Optics Letters*, vol. 36, no. 6, pp. 951–953, 2011.

[54] X. Wang, P. Zhou, Y. Ma, J. Leng, X. Xu, and Z. Liu, "Active phasing a nine-element 1.14 kW all-fiber two-tone MOPA array using SPGD algorithm," *Optics Letters*, vol. 36, no. 16, pp. 3121–3123, 2011.

[55] C. Yu, S. Augst, S. Redmond, K. Goldizen, D. Murphy, A. Sanchez, and T. Fan, "Coherent combining of a 4 kW, eight-element fiber amplifier array," *Optics Letters*, vol. 36, no. 14, pp. 2686–2688, 2011.

[56] P. Ma, P. Zhou, R. Su, Y. Ma, and Z. Liu, "Coherent polarization beam combining of eight fiber lasers using single-frequency dithering technique," *Laser Physics Letters*, vol. 9, no. 6, p. 456, 2012.

[57] B. Pulford, "LOCSET phase locking: operation, diagnostics, and applications," 2012.

[58] D. J. Bowman, M. J. King, A. J. Sutton, D. M. Wuchenich, R. L. Ward, E. A. Malikides, D. E. McClelland, and D. A. Shaddock, "Internally sensed optical phased array," *Optics Letters*, vol. 38, no. 7, pp. 1137–1139, 2013.

[59] L. E. Roberts, R. L. Ward, A. J. Sutton, R. Fleddermann, G. de Vine, E. A. Malikides, D. M. Wuchenich, D. E. McClelland, and D. A. Shaddock, "Coherent beam combining using a 2D internally sensed optical phased array," *Applied Optics*, vol. 53, no. 22, pp. 4881–4885, 2014.

[60] T. Weyrauch, M. Vorontsov, J. Mangano, V. Ovchinnikov, D. Bricker, E. Polnau, and A. Rostov, "Deep turbulence effects mitigation with coherent combining of 21 laser beams over 7 km," *Optics Letters*, vol. 41, no. 4, pp. 840–843, 2016.

[61] A. Flores, T. Ehrehreich, R. Holten, B. Anderson, and I. Dajani, "Multi-kW coherent combining of fiber lasers seeded with pseudo random phase modulated light," *SPIE LASE*, pp. 97 281Y–97 281Y–6, 2016.

[62] T. Wagner, "Fiber laser beam combining and power scaling progress: Air force research laboratory laser division," *SPIE LASE*, pp. 823 718–823 718–9, 2012.

[63] R. J. Mailloux, *Phased array antenna handbook*. Artech House Boston, 2005, vol. 2.

[64] D. Shaddock and R. Ward, "Personal discussion on coherent beam combining requirements," 2014.

[65] G. D. Goodno, C.-C. Shih, and J. E. Rothenberg, "Perturbative analysis of coherent combining efficiency with mismatched lasers," *Optics Express*, vol. 18, no. 24, pp. 25 403–25 414, 2010.

[66] D. A. Shaddock, "Digitally enhanced heterodyne interferometry," *Optics Letters*, vol. 32, no. 22, pp. 3355–3357, 2007.

[67] G. De Vine, D. S. Rabeling, B. J. Slagmolen, T. T. Lam, S. Chua, D. M. Wuchenich, D. E. McClelland, and D. A. Shaddock, "Picometer level displacement metrology with digitally enhanced heterodyne interferometry," *Optics Express*, vol. 17, no. 2, pp. 828–837, 2009.

[68] B. E. Saleh and M. C. Teich, *Fundamentals of photonics*. Wiley New York, 1991, vol. 22.

[69] M. T. Hsu, I. C. Littler, D. A. Shaddock, J. Herrmann, R. B. Warrington, and M. B. Gray, "Subpicometer length measurement using heterodyne laser interferometry and all-digital RF phase meters," *Optics Letters*, vol. 35, no. 24, pp. 4202–4204, 2010.

[70] M. George and P. Alfke, "Linear feedback shift registers in virtex devices," *Xilinx application note XAPP210*, 2007.

[71] New Wave Instruments, "Linear Feedback Shift Registers," Website, 2005. [Online]. Available: http://www.newwaveinstruments.com/resources/articles/m_sequence_linear_feedback_shift_register_lfsr.htm

[72] P. P. Chu and R. E. Jones, "Design techniques of FPGA based random number generator," in *Military and Aerospace Applications of Programmable Devices and Technologies Conference*, vol. 1, no. 999, 1999, pp. 28–30.

[73] Xilinx Inc., "Field Programmable Gate Array (FPGA): What is an FPGA?" Report, 2016. [Online]. Available: http://www.xilinx.com/training/fpga/fpga-field-programmable-gate-array.htm

[74] J. F. Wakerly, *Digital design*. Prentice Hall, 2000, vol. 3.

[75] D. Shaddock, B. Ware, P. Halverson, R. Spero, and B. Klipstein, "Overview of the lisa phasemeter," *Laser Interferometer Space Antenna (AIP Conference Proceedings Volume 873)*, vol. 873, pp. 654–660, 2006.

[76] O. Gerberding, B. Sheard, I. Bykov, J. Kullmann, J. J. E. Delgado, K. Danzmann, and G. Heinzel, "Phasemeter core for intersatellite laser heterodyne interferometry: modelling, simulations and experiments," *Classical and Quantum Gravity*, vol. 30, no. 23, p. 235029, 2013.

[77] O. Gerberding, "Phase readout for satellite interferometry," Thesis, 2014.

[78] K. Danzmann and A. Rüdiger, "LISA technology—concept, status, prospects," *Classical and Quantum Gravity*, vol. 20, no. 10, p. S1, 2003. [Online]. Available: http://stacks.iop.org/0264-9381/20/i=10/a=301

[79] A. V. Oppenheim and R. W. Schafer, *Discrete-time signal processing*. Pearson Higher Education, 2010.

[80] E. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 2, pp. 155–162, 1981.

[81] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 5th ed. Pearson Prentice-Hall, 2006.

[82] W. T. Scherer, T. A. Pomroy, and D. N. Fuller, "The triangular density to approximate the normal density: decision rules-of-thumb," *Reliability Engineering & System Safety*, vol. 82, no. 3, pp. 331–341, 2003.

[83] D. Montgomery and G. Runger, *Applied Statistics and Probability for Engineers, 3rd Edition*. John Wiley & Sons, Inc., 2003.

[84] S. K. Garakoui, E. A. Klumperink, B. Nauta, and F. E. van Vliet, "Phased-array antenna beam squinting related to frequency dependency of delay circuits," *Microwave Conference (EuMC), 2011 41st European*, pp. 1304–1307, 2011.

[85] Optiphase, "PZ1 Data Sheet (Rev F)," 2013. [Online]. Available: http://www.optiphase.com/data_sheets/PZ1_Data_Sheet_Rev_F.pdf

[86] Coherent, Inc., "Prometheus Ultra-Narrow Linewidth CW DPSS Green Laser," 2013. [Online]. Available: http://www.coherent.com/downloads/Prometheus_DS_1013revA_3.pdf

[87] Isomet Corp., "1205-843C NIR Acousto-Optic Modulator," 2014. [Online]. Available: http://www.isomet.com/AO_Pdf/1205C-843.pdf

[88] Keysight Technologies, "Keysight 33250A Function/Arbitrary Waveform Generator," 2015. [Online]. Available: http://literature.cdn.keysight.com/litweb/pdf/5968-8807EN.pdf?id=1000034211:epsg:dow

[89] New Focus, Inc., "User's Guide: 125-MHz Photoreceivers Models 1801 and 1811," 2001.

[90] National Instruments Corporation, "NI PXIe-1082 User Manual," 2010. [Online]. Available: http://www.ni.com/pdf/manuals/372752b.pdf

[91] National Instruments Corporation, "NI R Series Multifunction RIO Specifications," 2009. [Online]. Available: http://www.ni.com/pdf/manuals/373712f.pdf

[92] National Instruments Corporation, "NI PXIe-8133 User Manual," 2012. [Online]. Available: http://www.ni.com/pdf/manuals/372870d.pdf

[93] Maxim Integrated Products, "MAX1165 Evaluation Kit," 2002. [Online]. Available: https://datasheets.maximintegrated.com/en/ds/MAX1165EVKIT.pdf

[94] National Instruments Corporation, "Choosing Pull-Up Resistor Values for R Series Devices Digital Output Lines," Website, 2011. [Online]. Available: http://digital.ni.com/public.nsf/allkb/D58C269F7771770E86256F490082E6C8

[95] G. Ascheid and H. Meyr, "Cycle slips in phase-locked loops: A tutorial survey," *IEEE Transactions on Communications*, vol. 30, no. 10, pp. 2228–2241, 1982.

[96] S. P. Francis, T. T. Y. Lam, K. McKenzie, A. J. Sutton, R. L. Ward, D. E. McClelland, and D. A. Shaddock, "Weak-light phase tracking with a low cycle slip rate," *Optics Letters*, vol. 39, no. 18, pp. 5251–5254, 2014.

[97] P. F. McManamon, T. A. Dorschner, D. L. Corkum, L. J. Friedman, D. S. Hobbs, M. Holz, S. Liberman, H. Q. Nguyen, D. P. Resler, and R. C. Sharp, "Optical phased array technology," *Proceedings of the IEEE*, vol. 84, no. 2, pp. 268–298, 1996.

[98] A. Yariv, "Dynamic analysis of the semiconductor laser as a current-controlled oscillator in the optical phased-lock loop: applications," *Optics Letters*, vol. 30, no. 17, pp. 2191–2193, 2005.

[99] Z. Xu, K. Huang, and X. Lu, "A digital optical phase-locked loop based on field programmable gate array and its applications," *Information Science, Electronics and Electrical Engineering (ISEEE), 2014 International Conference on*, vol. 2, pp. 795–799, April 2014.

[100] D. M. Wuchenich, T. T.-Y. Lam, J. H. Chow, D. E. McClelland, and D. A. Shaddock, "Laser frequency noise immunity in multiplexed displacement sensing," *Optics Letters*, vol. 36, no. 5, pp. 672–674, 2011.

[101] T. T.-Y. Lam, B. J. Slagmolen, J. H. Chow, I. C. Littler, D. E. McClelland, and D. A. Shaddock, "Digital laser frequency stabilization using an optical cavity," *IEEE Journal of Quantum Electronics*, vol. 46, no. 8, pp. 1178–1183, 2010.

[102] E. Balestrieri and S. Rapuano, "Defining dac performance in the frequency domain," *Measurement*, vol. 40, no. 5, pp. 463–472, 2007.

[103] Opto-Link Corporation Ltd., "Polarization maintaining couplers," 2008. [Online]. Available: http://www.optolinkcorp.com/pdf/CPL_PM.pdf

[104] JDSU, "Continuous Wave (CW) Single Frequency IR Laser: NPRO 125/126 Series," 2006. [Online]. Available: http://www.lightwavestore.com/product_datasheet/OTI-Tx-L-100C_pdf1.pdf

[105] National Instruments Corporation, "NI PXIe-1085 Series User Manual," 2015. [Online]. Available: http://www.ni.com/pdf/manuals/373712f.pdf

[106] Coherent, Inc., "Mephisto Ultra-Narrow Linewidth CW DPSS Laser," 2013. [Online]. Available: http://www.coherent.com/downloads/Mephisto_DS_1013revA_2.pdf

[107] Stanford Research Systems, Inc., "Operating Manual and Programming Reference: Model SR785 Dynamic Signal Analyzer, Rev. 1.4," 2006. [Online]. Available: http://www.thinksrs.com/downloads/PDFs/Manuals/SR785m.pdf

[108] Stanford Research Systems, Inc., "Model SR560 Low-Noise Preamplifier, Rev. 3.0," 2013. [Online]. Available: http://www.thinksrs.com/downloads/PDFs/Manuals/SR560m.pdf

[109] RP Photonics Consulting GmbH, "Fiber joints," Website, 2015. [Online]. Available: https://www.rp-photonics.com/fiber_joints.html

[110] P. Wardill, Gooch & Housego, "Personal discussion on viability of LMA fiber couplers," 2015.

[111] A. Kobyakov, M. Sauer, and D. Chowdhury, "Stimulated Brillouin scattering in optical fibers," *Advances in optics and photonics*, vol. 2, no. 1, pp. 1–59, 2010.

[112] B. Anderson, A. Flores, R. Holten, and I. Dajani, "Comparison of phase modulation schemes for coherently combined fiber amplifiers," *Optics Express*, vol. 23, no. 21, pp. 27 046–27 060, 2015.

[113] A. Flores, C. Robin, A. Lanari, and I. Dajani, "Pseudo-random binary sequence phase modulation for narrow linewidth, kilowatt, monolithic fiber amplifiers," *Optics Express*, vol. 22, no. 15, pp. 17 735–17 744, 2014.

[114] A. R. Wade, S. S. Chua, M. S. Stefszky, D. A. Shaddock, and D. E. McClelland, "Path length modulation technique for scatter noise immunity in squeezing measurements," *Optics Letters*, vol. 38, no. 13, pp. 2265–2267, 2013.

[115] S. Gross and M. Withford, "Ultrafast-laser-inscribed 3d integrated photonics: challenges and emerging applications," *Nanophotonics*, vol. 4, no. 1, 2015.

[116] A. Arriola, S. Mukherjee, D. Choudhury, L. Labadie, and R. R. Thomson, "Ultrafast laser inscription of mid-ir directional couplers for stellar interferometry," *Optics Letters*, vol. 39, no. 16, pp. 4820–4822, 2014.